UNIVERSITÉ DE SHERBROOKE
Faculté de génie
Département de génie mécanique

TLC: Une architecture photovoltaïque concentrée (CPV) au potentiel d'efficacité élevé à faible coût

TLC: A CPV Architecture with the Potential for high efficiency at very low cost

Thèse de doctorat
Spécialité : Energie Solaire Concentré

Richard Norman

Sherbrooke (Québec) Canada 31 Août 2022

*** TLC et Trough-Lens-Cone sont marques du commerce de Terra Firma Innovations Inc.**

MEMBRES DU JURY

Luc FRECHETTE
Directeur

Vincent AIMEZ
Codirecteur

Simon FAFARD
Évaluateur

Laurent BECHOU
Évaluateur

Richard ARES
Évaluateur

RÉSUMÉ

La civilisation humaine est devenue de plus en plus dépendante d'un accès facile à une énergie à faible coût, mais les combustibles fossiles qui répondent actuellement à la majeure partie des besoins énergétiques de l'humanité causent la destruction de l'environnement, y compris un réchauffement climatique potentiellement catastrophique. L'énergie solaire a le potentiel d'arrêter le réchauffement climatique et, si son coût est suffisamment bas, d'amener également la population mondiale entière à un niveau de vie du premier monde. Les coûts de photovoltaïque (PV) à base de silicium ont été considérablement réduits en grande partie en diminuant le prix et en augmentant l'efficacité des cellules en silicium, cependant l'utilisation de silicium a ses limites d'efficacité théoriques, et même si les cellules étaient gratuites, la PV à base de silicium serait encore trop chère pour atteindre ces objectifs. Les cellules de photovoltaïque concentré (CPV) Tandem sont environ deux fois plus efficaces que celles à base de silicium, mais malgré l'avantage de leur efficacité, les architectures des années précédentes de CPV n'ont pas été en mesure de rivaliser avec le silicium en termes de coût.

Une nouvelle architecture CPV, appelée TLC (*Trough-Lens-Cone*) utilise la concentration initiale par un miroir parabolique à faible coût combiné avec un module CPV de 40X et ainsi réduire les coûts globaux du module.

Avant ce projet de recherche de doctorat, TLC n'était qu'une étude sur papier. Cette thèse a pour but de répondre à la question de savoir si l'approche TLC pouvait fonctionner aussi bien qu'elle était apparue, ou s'il y avait des défauts cachés qui empêchaient de battre le silicium PV sur le coût, ou pourrait même empêcher la TLC de fonctionner. Ce travail comprenait la construction d'un modèle de tableur unifié qui reliait les aspects connus de la conception TLC et les coûts estimés pour une variation de conception donnée. Nous présentons également la construction de modèles 3D-CAD pour raffiner la conception TLC, puis le prototypage de pièces individuelles et de processus, et enfin la construction d'un prototype physique d'un mini-module TLC qui est mis au soleil. Cette validation physique était nécessaire car même après que TLC ait été théoriquement et numériquement « construit » à plusieurs reprises soit, en visualisation, sur papier, sur des feuilles de calcul, puis dans COMSOL, avant que TLC soit physiquement construit, des défauts cachés pouvaient survenir à tout moment.

La mise en œuvre de ce projet a réussi, produisant une conception TLC cohérente qui avait un rendement élevé avec un coût des matériaux très bas et des faibles coûts estimatifs de processus, avec un potentiel de battre même l'objectif du département américain de l'énergie pour la tarification du silicium photovoltaïque en 2030. Le suivi de raies (*Ray-tracing)* avec un modèle 3D a montré que la conception pouvait atteindre une concentration élevée avec des angles d'acceptation adéquats. Les tests ont également montré que les cellules de prototypage ont été bien adaptées à la nouvelle configuration de TLC de récepteur à matrice de microcellules massivement parallèle. Le projet a également testé avec succès le processus de fabrication proposé pour le moulage de réseaux semi-denses d'éléments optiques tertiaires à l'arrière d'un carreau de lentille. Le projet a également réussi à assembler un mini-module TLC et à tester sous le soleil avec le focus d'un miroir parabolique. Quatre articles ont déjà été publiés, avec un cinquième article accepté, à la suite de ce travail.

Mots-clés : CPV, Photovoltaïque à Concentrateur, Concentration Multi-Étages, Photovoltaïque à Faible Coût, LCOE, Miroir Parabolique, Énergie Solaire

TLC: A CPV Architecture with the Potential for high efficiency at very low cost

Abstract

Human civilization has grown dependent on ready access to low-cost energy, but the fossil fuels that currently meet the bulk of humanity's energy needs are causing environmental destruction, including potentially catastrophic global warming. Solar energy has to potential to halt global warming, and, if low enough in cost, to also bring the whole world's population to a first world living standard. Silicon PV has dramatically reduced costs largely through decreasing the cost and increasing the efficiency of the silicon cells, but silicon is nearing its theoretical efficiency limits, and even if the cells were free, silicon PV would still be too expensive to meet these goals.

Tandem CPV cells are roughly twice as efficient as silicon, but previous CPV designs have been unable to compete with silicon on cost in spite of the efficiency advantage. A new CPV architecture, called TLC for its trough, lens and cone concentration stages, proposed using initial concentration by a low-cost trough mirror to shrink the rest of an CPV module by 40X and thus reduce overall module costs. But before this PhD research project, TLC was only a paper study.

This PhD research project was started to answer the question of whether TLC could work out as well as it appeared, or whether there were hidden flaws that precluded beating silicon PV on cost, or possibly even precluded TLC from working at all. Thesis chapter 3 details the main optical design aspects, and chapter 4 covers the design of the rest of the TLC module, including leading variations where there is more than one plausible way to achieve low cost and high reliability.

The work included building a unified analytical model spreadsheet that linked known aspects of the TLC design together and estimated costs for a given design variation. Thesis chapter 5 covers the economics of the proposed design, with a focus on materials costs since these dominate PV overall costs, and a section on reliability since product lifetime strongly influences life-cycle cost.

The work included building 3D-CAD models to refine the TLC design, and then the prototyping of individual parts and processes, and finally building a physical prototype of a TLC mini-module and putting it in sun. This physical confirmation was necessary because even after TLC has been "built" many times, in visualization, on paper, on spreadsheets, and then in COMSOL, until TLC was physically built, hidden flaws could arise at any time. Chapter 6 of this thesis covers the simulation and validation carried out to show that it is plausible that TLC can meet its cost targets.

The conclusion of this thesis summarizes the overall project. The project was a success, producing a TLC design with high potential efficiency, very low materials cost, and low estimated process costs, with the potential to beat even the US Department of Energy's goal for PV pricing in 2030. Ray-tracing a 3D model showed that the design could achieve high concentration with adequate acceptance angles, and tests showed that the prototyping cells were suitable for TLC's massively parallel microcell-array receiver configuration. The project also successfully tested the proposed manufacturing process for molding semi-dense arrays of tertiary optical elements on the back of a lens tile and assembled a TLC mini module which was tested on sun at the focus of a trough mirror. Four papers have already been published, with a fifth paper accepted, as result of this work.

TLC and Trough-Lens-Cone are trademarks of Terra Firma Innovations Inc.

Keywords: CPV, Concentrator Photovoltaics, Multi-stage Concentration, Low-cost Photovoltaics, LCOE, Parabolic Trough, Solar Energy

Dedicated to all who strive to make the world a better place, and to the family, friends, mentors, investors, business associates, professors, students, and professionals who have helped me pursue my dreams.

Acknowledgment

I thank the University of Sherbrooke for guidance from my supervisors, for access to expertise and test equipment and for professional support (especially Etienne Léveillé). I thank Programme Innovation, IRAP, MITACS and Terra Firma Innovations Inc. (TFI) for funding this research and I thank TFI for access to its TLC design work and use of overview illustrations and descriptions for TLC. TLC is a trademark of, and the TLC design is the intellectual property of TFI.

I also thank LN2 - a joint International Research Laboratory (IL 3463) funded and co-operated in Canada by Université de Sherbrooke (UdeS) and in France by CNRS as well as ECL, INSA Lyon, and Université Grenoble Alpes (UGA) and is supported by the Fonds de Recherche du Québec Nature et Technologie (FRQNT).

TABLE OF CONTENTS

TABLE OF CONTENTS ____ *i*

CHAPTER 1 ____ *1*

1. INTRODUCTION ____ *1*

1.1 Context and Issues ____ **1**
1.1.1 Market Need ____ 1
1.1.2 Trough-Lens-Cone (TLC) concept ____ 1
1.1.3 TLC Potential ____ 2
1.1.4 Itemized Requirements ____ 3
1.1.5 Single-axis or Dual-axis Tracking ____ 3

1.2 Research question ____ **4**

1.3 Objectives of the research project ____ **5**

1.4 Original contributions ____ **5**

1.5 Outline of the document ____ **6**

CHAPTER 2 ____ *7*

2. STATE OF ART ____ *7*

2.1 Overview of TLC (Trough-Lens-Cone) CPV ____ **7**
2.1.1 Tracking ____ 7
2.1.2 Low-cost Focusing Principles ____ 7

2.2 Initial Product ____ **11**
2.2.1 Summary of Inputs and Outputs ____ 11

2.3 Business Overview ____ **12**
2.3.1 Value Advantage ____ 12
2.3.2 Scaling Production ____ 14

2.4 Prior Art ____ **15**

2.5 Competitive Architectures ____ **17**
2.5.1 Standard Silicon and thin-film Flat Panels ____ 17
2.5.2 Standard Fresnel HCPV ____ 18
2.5.3 Discrete Cassegrain Mirror Flat-Panel HCPV ____ 19
2.5.4 Other Single-Axis-Tracked High-Concentration Trough Work ____ 20
2.5.5 Two-axis-tracked Troughs with Secondary Concentrators ____ 20
2.5.6 Primary Optics Sheet with Small Cells ____ 21
2.5.7 Two-Axis-Micro-Tracked System ____ 23
2.5.8 Big Dish CPV ____ 24
2.5.9 Power Tower CPV ____ 26
2.5.10 Emerging technologies ____ 26

2.6 Advantages and Issues Comparison Table ____ **29**

2.7 Conclusion ____ **30**

CHAPTER 3 ____ *31*

3. TLC CONCEPTION & DESIGN 31

3.1 Context and Issues 31

3.2 Introduction 31

3.3 Multi-stage Concentration 31

3.3.1 Initial Concentration 31
3.3.2 Achieving High Concentration 33
3.3.3 Impact of the Optics on the rest of the Module 37

3.4 Lens Details 39

3.4.1 Lens Glass Type 39
3.4.2 Cell Width and Lens Focal Length 42
3.4.3 Trade-offs in optics optimization 43
3.4.4 Lens Sheet Manufacturing 43
3.4.5 Molding Arrays of Tertiary Optical Elements for Microcell Receivers 49
3.4.6 Alternative Cone Array Manufacturing Methods 60
3.4.7 Lens Sheet Cost: 61
3.4.8 Acceptance Angles 66
3.4.9 Evening-out the Focal Intensity 68

3.5 Design Scaling 71

3.5.1 Tracker 71
3.5.2 Trough Rim Angle (Focal Length/Width) 71
3.5.3 Trough Width and Receiver Size 72
3.5.4 Cell and Lens Row Width 74
3.5.5 Scaling Down to Smaller Installations 75

3.6 Conclusion 75

CHAPTER 4 76

4. ASSEMBLY PROCESS & MATERIAL COSTS 76

4.1 Introduction 76

4.2 Cells & Electrical Interconnections 76

4.2.1 Cell Type 76
4.2.2 Cell Size 77
4.2.3 Future Cells 84
4.2.4 Cell Cost 86
4.2.5 Cell Attach Cost 88
4.2.6 Variance in Cell + Attachment Cost with Cell Size 90
4.2.7 Higher-concentration Cells 91

4.3 Receiver backplane/Spreader 91

4.3.1 Replacing DBC/AlN with an IMS 91
4.3.2 Standard IMS Process 92
4.3.3 Optimized IMS 94
4.3.4 Receiver CTE Mismatches 102
4.3.5 Receiver Material Costs 103

4.4 Receiver Electrical Connections 104

4.4.1 Cells in Parallel on a Receiver 104
4.4.2 Massively-parallel Microcell Arrays and Thermal Runaway in Unilluminated Cells 105
4.4.3 Back Contact to Substrate as Conductor 117

4.4.4 Top-Contact Conductor 117
4.4.5 Bond-wire Resistance 118
4.4.6 Bypass Diodes 118
4.4.7 Interconnecting Receivers 119
4.4.8 Module Wires 120

4.5 Heat Rejection Housing 121
4.5.1 Overview of Heat-rejection Methods 121
4.5.2 Fin Heat-Rejection Scaling Background 122
4.5.3 Monolithic Aluminum Heatsink Lid 123

4.6 Receiver Module 129
4.6.1 Lens Subassembly 130
4.6.2 Bonding the Sub-assemblies 134
4.6.3 Module Rail 135
4.6.4 Module Vertical Supports 135
4.6.5 Module Length 136
4.6.6 Module/Trough Ends 136

4.7 Dual-Axis Tracker 138
4.7.1 Mechanical Fit 138
4.7.2 Electrical Match 139
4.7.3 Dual-Axis Tracker Cost 140
4.7.4 Space Frame 142
4.7.5 Floating Azimuth Two-axis Tracking 142

4.8 Module Shipping and Installation 143
4.8.1 Shipping 143
4.8.2 On-site Installation 144

4.9 Conclusion 147

CHAPTER 5 148

5. ECONOMICS 148

5.1 Introduction 148

5.2 Assembly Walk-throughs and Costs 148
5.2.1 Components Preparation Walk-through 148
5.2.2 Receiver Assembly Walk-through 149
5.2.3 Variations on the Receivers-to-Lens Assembly Step 151
5.2.4 Module Assembly Costs 152

5.3 Reliability 153
5.3.1 Temperature Cycling 154
5.3.2 UV 155
5.3.3 Degradation 156
5.3.4 Comparison to Silicon PV 156

5.4 Summary of Costs and Value 157
5.4.1 Costs in ¢/W 157
5.4.2 Cost Reductions 159
5.4.3 Value 164
5.4.4 SunShot 2030 – Pathways to 3¢/kWh 168
5.4.5 High Efficiency 169
5.4.6 Target Market 171

5.4.7 Other Markets 172
5.4.8 Production Scaling 172

5.5 Conclusion 176

CHAPTER 6 177

6. Simulation and Analytical Validation 177

6.1 Introduction 177

6.2 Analytical models 177
6.2.1 Financial Tab 177
6.2.2 Half-Trough Shape Tab 178
6.2.3 Rim Angle and Receiver Slant 178
6.2.4 Primary Optics Tab 178
6.2.5 Lens and Receiver Size Tab 179
6.2.6 Lens Curve Tab 180
6.2.7 TIR Cones Tab 183
6.2.8 Substrate Tab 185
6.2.9 Cell & Attach Tab 185
6.2.10 Electrical Tab 185
6.2.11 Thermal Conduction Tab 186
6.2.12 Housing and Heat Rejection Tab 186
6.2.13 Module Assembly Tab 186
6.2.14 CTE Stress Tab 187
6.2.15 Equipment Tab 187
6.2.16 Tracker Tab 187

6.3 Multi-Physics Modeling and Analysis 187
6.3.1 Modeling and Simulation Overview 187
6.3.2 Simple Heat Flow 190
6.3.3 Full Thermal Stress Model, Flexible Adhesives 199
6.3.4 Thermo-Optical Model 201
6.3.5 Electrical Model 201
6.3.6 Heat-Rejection Model 201

6.4 Final Design and Conclusion 201

CHAPTER 7 203

7. Prototyping and Experimental Validation 203

7.1 Details of TLC Prototype/test at the University of Sherbrooke 203
7.1.1 PDMS optical silicone RI ~ 1.415: 203
7.1.2 Dispersing blocks: 203
7.1.3 Heatsink stock: 204
7.1.4 Pre-existing molds for shrinkage and surface roughness testing: 204
7.1.5 QDEC oversized CPV microcells: 205
7.1.6 Ultra-low-iron glass: 206
7.1.7 Three-wavelength laser: 207
7.1.8 TIR-versus-smoothness tester: 207
7.1.9 Silicon oversized CPV microcells and receiver-sized cells: 208
7.1.10 Thermally conductive prepreg and copper laminate: 209
7.1.11 Ordinary-IMS receiver substrates: 210
7.1.12 Alignment test stand for testing with laser 210

7.1.13 Temporary Mirror: 211
7.1.14 Brazed-flange heatsink lid: 212
7.1.15 Lens and thermal bridge mold: 213
7.1.16 Cone mold mold: 214
7.1.17 Self-aligning cone mold: 216
7.1.18 First lens tile: 217
7.1.19 RP-3 Trough mirror: 218
7.1.20 Receiver substrate populated with cells and diode, interconnected: 218
7.1.21 Receiver-sized silicon cell test set up: 219
7.1.22 Artificial Sun: 219
7.1.23 First electro-optical mini-module: 221
7.1.24 Instrumentation and Data acquisition system (DAQ): 222
7.1.25 Mirror and prototype holder mounted on tracker: 224
7.1.26 First mini-module (populated lens tile with heatsink) 225
7.1.27 Mini-module on sun 226
7.1.28 Well-polished cone-plate 230
7.1.29 Camera Image Sensor: 230
7.1.30 Laser Scan on Silicon-Cell Receiver 232
7.1.30.1 Two-cell Scan 234
7.1.30.2 Scan at a Rim-angle Slant 235
7.1.30.3 Lens-to-Lens Transition 236
7.1.30.4 Other Optical Losses 237

7.2 Conclusion 237

CHAPTER 8 238

8. CONCLUSION 238

8.1 Summary 238

8.2 Original Contributions 240

8.3 Research Perspectives 240

8.4 Sommaire 241

8.5 Contributions originales 243

8.6 Perspectives de recherche 244

APPENDIX A 245

Additional TLC Details and Variations 245

A.1 TLC Durability Compared to Silicon PV 245

A.2 Other Variations of TLC 253

A.3 Alternate Aluminum-Fin Heatsinks 263

A.4 CTE-matched-Heatsink Lid (backup) 265

A.5 Heat Pipe with Fin Tube: 270

A.6 Pumped Coolant with Fin Tube: 271

APPENDIX B 272

TLC ON SINGLE-AXIS TRACKERS 272

B.1 Single-axis-tracked TLC without Micro-tracked Receivers 272
B.2 Single-Axis vs. Dual-Axis TLC 272
B.3 TLC Micro-Tracking 276
B.4 Micro-tracking Goals 283
B.5 Micro-tracked Receiver Cooling 292
B.6 Micro-tracked Receiver Interconnections 297
B.7 Stripped end of Terminal Wire 298
B.8 Module (4060 mm) that includes Micro-tracking 299
B.9 Micro-tracked Module Assembly Walk-through 312
B.10 Module Assembly Cost 312
B.11 Single-Axis Tracking Requirement 313
9. LIST OF REFERENCES 315

LIST OF FIGURES

Figure 2.1: *TLC Overview Illustration (TFI illustration, first published in* (Norman et al., 2018)*).* 8
Figure 2.2: *The value ($/W) of efficiency and durability (overlay on graph from* (DOE, 2018)*).* 13
Figure **3.1:** *TLC 3-Stage Focusing.* 33
Figure 3.2: *Basic Roll-formed TLC Lens Sheet.* 34
Figure 3.3: *Lens Panel x-section showing TIR Cones.* 34
Figure 3.4: *Achieving High Concentration onto Microcells (right half first published in* (Norman, et al., 2019)*)* 35
Figure 3.5: Lens *Cross Section ALONG Trough Axis, ~3X Scale.* 36
Figure 3.6: *Re-Concentration ACROSS Trough Axis, ~3X Scale.* 36
Figure 3.7: Overall module (cross-section across trough's focus). 38
Figure 3.9: *Self-aligning Cone Mold (showing cone cavities).* 47
Figure 3.10: *Cone Mold Flexes to Mate to Thermal Contacts.* 48
Figure 3.11: *A Single Cone Plate has One Cone from Each Row.* 48
Figure 3.12: *Cone Plates packed to form a Cone Array.* 48
Figure 3.13*: (a) Visual diagram. (b) Energy flow through optics and receivers.* 53
Figure 3.14: *polyethylene Cone Mold can stretch between cavity rows to align to features on lens tile* ... 55
Figure 3.15*: Precision Parts of Mold for Cone Molds and resulting Cone Mold.* 56
Figure 3.16: *(a) Cone Mold molded in polyethylene, (b) Cone array over-molded on low-iron glass in optical silicone.* 57
Figure 4.1: *Cell Comparison: (a) Azur 3 mm, Semprius, TLC cells ~20x scale, (b) TLC cell detail* 78
at ~60x. 78
Figure 4.2: *Two bond-wire cell optimization.* 80
Figure 4.3: *Compressed Cone Coupling.* 80
Figure 4.4: *Cell Size Impact on V_{OC} (Fidaner).* 81
Figure 4.5: *Standard IMS layout for Thermal Optimum (2X scale, unpopulated).* 93
Figure 4.6: *Standard IMS layout for Electrical Optimum (2X scale, populated).* 93
Figure 4.7: *Punch Press Pattern Plate substrate.* 96
Figure 4.8: *Place, Reflow, Wire-bond.* 97
Figure 4.9: *Region of panel with strips of top copper removed to expose prepreg.* 98
Figure 4.10*: Substrate-sized region with narrower cuts exposing backplane.* 99
Figure 4.11: *Substrate-sized region with photoimagable soldermask.* 100
Figure 4.12: *Assembled receiver on sawn substrate.* 101
Figure 4.13*: (a) Visual diagram, (b) Energy flow.* 108
Figure 4.14: *With no wind, COMSOL calculates a maximum TLC cell temperature of 50°C above ambient (300K).* 110
Figure 4.15: *Injected current in increases rapidly with both voltage and temperature.* 112
Figure 4.16*: Approximate trend of injected current versus temperature when attack-V_{OC} is temperature-corrected* 113
Figure 4.17: *Exposed Conductors ~2.5x scale (backplane is orange, power-plane is gold).* 119
Figure 4.18: *Electron Flow through Receivers (thickness exaggerated).* 120
Figure 4.19: *Sketch of Monolithic Extruded Heatsink Lid.* 124
Figure 4.20: *Extrusion Profile for Monolithic Heat-sink Lid.* 126
Figure 5.1*: Effect of Die Size on Solder Joint Reliability.* 155
Figure 5.2*: Detailed Costing Pie Chart.* 158
Figure 5.3: *17.1 ¢/W Cost Pie Chart at 100 MW.* 163
Figure 5.4: *Cost Charts 15.8 ¢/W at 1 GW (left) and 14.2 ¢/W at 10GW (right).* 163

Figure 5.5: *Costs at 100 GW (Left: without substrate reuse, right: with reuse).* 164
Figure 5.6: *Value of Modules - **Maximum selling price that would meet the DOE 2030 goal**.* 165
Figure 5.7: *Path Toward 3¢/kWh* (Woodhouse, et al., 2016). 168
Figure 5.8: *CPV Losses (from S. Kurtz), annotated for TLC.* 170
Figure 6.1: *Figuring out the Exact Rim Angle.* 178
Figure 6.2: *Lens Width from Trough Focal Width.* 179
Figure 6.3: *Calculations for the lens curve.* 182
Figure 6.4: *Curve of half-lens (Both axes in millimeters).* 183
Figure 6.6: *Calculating the first TIR facet.* 184
Figure 6.8: *Thermal Analysis of TLC Unit Cell (heatsink set to 300K).* 192
Figure 6.9: *Zoom of Cell Temperature.* 192
Figure 6.10: *TLC Cooling in COMSOL model: Overall (top), zoom on cells (bottom).* 194
Figure 6.11: *Ray-traced focus **across** the trough's focus* 197
Figure 6.12: *Ray-traced focus **along** the trough's focus* 198
Figure 6.13: *Ray-traces at the center wavelength of each junction* 198
..... 204
Figure 7.1 Dispersing block shows the path of light 204
Figure 7.2: Confirming parallel microcells' resistance to thermal runaway 206
(published in (Norman, et al., 2020)) 206
Figure 7.3: One of the three TIR tester setups (by N. Caillou) 208
Figure 7.4: Silicon CPV cell with bus bar on one side 208
Figure 7.5: Standard IMS substrate for 6 x 16 array of cells 210
Figure 7.6: Alignment test stand (left), and prototype-holder base(right) 211
Figure 7.7: Temporary parabolic mirror to test mounting and tracking 211
Figure 7.8: Lens and thermal bridge mold 213
Figure 7.9: Polished surface for lens molding 213
Figure 7.10: Glass block with silicone lenses and thermal bridges 213
Figure 7.11: Cone mold mold 215
Figure 7.12: Self-aligning Cone Mold (showing cone cavities) (EL & NC) 216
Figure 7.13: Cone Mold Flexes to Mate to Thermal Bridges 216
Figure 7.14: Array ~3 mm silicone cones with mold from unpolished master array (NC) 217
Figure 7.15: Side view of cones with 1-mil (25 µm) gold wire for scale 218
Figure 7.16: The lenses and cones guide the light to microcell-sized focal spots 220
Figure 7.17: (left) Polished cone at top has bright tip and polished cones leak less light (right) 221
Figure 7.18: Roughly to-scale sketch of cone and bondwire on cell 221
Figure 7.19: Receivers mounted on a lens tile (N. Caillou, 2021) 222
Figure 7.20: Sliding mask for safety and for short-duration tests 225
Figure 7.21: Two-receiver mini-module mounted in heatsink (N. Caillou, 2021) 225
Figure 7.22: Optical Efficiency vs Optical Energy (Watts) 228
Figure 7.23: Altitude Acceptance versus miss-tracking (degrees) 229
Figure 7.24: Altitude Acceptance versus miss-tracking (degrees) 229
Figure 7.25: Polished cone (left) and unpolished cone (right) (from Etienne Léveillé) 230
Figure 7.26: Uncollimated light through polished (left) and unpolished (right) cones (Pictures from Etienne Léveillé) 231
Figure 7.27: Sensor fits multiple cones for comparison under same conditions (from N. Gariepy video) 231
Figure 7.28 Scan of a bare cell to profile the laser itself 232
Figure 7.29 Optical efficiency map through a lens tile with an unpolished cone 233
Figure 7.30 Optical efficiency map with a Parylene-polished cone 234

Figure 7.31 Optical efficiency map with a hand-polished cone 234
Figure 7.32 Combined optical efficiency map with adjacent Parylene-polished (top) and hand-polished (bottom) cones 235
Figure 7.33 Combined map of the same lenses/cones with the laser at a 25-degree slant. 236
Figure 7.34 Combined map of two lenses. 236
Figure A.1: *Glass-cored over-molded cone.* 262
Figure A.2*: Fin-pair Profile.* 267
Figure B.1: *The need for multi-row lenses. Individual simple lenses collide (left), Multi-row lens tiles don't collide (right).* 277
Figure B.2: *Lens Curve Shifted for Clearance.* 282
Figure B.3: *Ideal Lens Height with Tilt.* 284
Figure B.4: *Off-Center Axis of Rotation.* 285
Figure B.5: *Off-center Axis Result.* 287
Figure B.6: *Receiver Cross Section with Micro-Tracking (~1.6x scale).* 289
Figure B.7: *Coolant flow in receiver, only 3 rows shown.* 294
Figure B.8: *Manifold/Lens Glass-Glass Sandwich* 295
Figure B.9: *Coolant Loop.* 306
Figure B.10: Receivers on Rod 310
Figure B.11*: Dual-off-axis Slanted Modules.* 311

Lexicon / Glossary / Definition of Terms

Term	Definition
2nd 10 MW	The second 10 megawatts refers to introductory production volume after the initial NRE and tooling costs are covered. The goal is to have TLC have high enough margins in its launch markets to fund growth, starting from the 2nd 10 MW.
Aberration	The deviation of an actual focus from a perfect focus (e.g., chromatic aberration arises from lens glass having a different index of refraction, and hence a different focal length, at different wavelengths).
Acceptance Angle	The acceptance angle is the angular range of light that the optics can funnel to the CPV cell. This largely affects manufacturing and tracking tolerances; systems with larger acceptance angles are more forgiving. TLC so far has concentrated on minimizing the acceptance angle needed rather than on maximizing the acceptance angle, which will be done once the optics are fully modelled including chromatic aberration.
Active area	The power-producing area of a solar cell. The top of a solar cell typically has an active area that has a bus bar on one or more sides of the active area. Also called "Photovoltaic Area".
Air-gap lens	An air-gap lens refers to a lens that has the refractive front-surface curves and the reflective back-surface cones separated by an air gap rather than by solid glass.
AlSiC	Aluminum (Al) Silicon (Si) Carbide (C), a tough, low-CTE material made from silicon carbide (SiC) particles embedded in an aluminum matrix. The SiC provides high stiffness and low thermal expansion (CTE), which the aluminum holds the composite together with the strength of a metal.
Aluminum Nitride, AlN	A hard low-CTE ceramic with high thermal conductivity and excellent electrical insulation, commonly used CPV and TPV receivers to electrically isolate cells or cell arrays from the rest of the receiver. ALN has been eliminated from the TLC design in favor of AlSiC//LCP, which has, in turn been replaced by copper.
AR Coating, ARC	Anti-Reflection Coating. Since sharp changes in index of refraction produce reflections, and AR coating can be one or more thin layers of material with an intermediate index of refraction, reducing the amount of reflection. Some newer AR coatings have layers with structures smaller than a wavelength of light, but these are still generally expensive or fragile.
ARPA-E	Advanced Research Projects Agency for Energy. A U.S. government program aimed at potentially revolutionary but high-risk technologies in the energy production space (modelled after DARPA).
At latitude	At latitude refers to a trough or a single-axis tracker whose axis of rotation is parallel to the earth's axis of rotation, or slanted to the nearest pole at the same slant as the site's latitude.
Azur	The main European marked of tandem cells for CPV, and currently the company with the highest-efficiency commercial cells.
B270	A relatively low-cost instrument-grade glass
Backplane	An electrically conductive plane that electrically connects the backs of all of the solar cells in a TLC receiver. Also called a 'Heat Spreader' for its thermal function.
Bare Die	A chip that is not packaged, saving space.

Term	Definition
Beyond-SunShot	Once it became clear that the U.S. DOE's SunShot goal of $1/W installed cost would be reached, the DOE announced a SunShot 2030 goal of 3¢/kWh, unsubsidized, in regions of moderate sunlight (to allow for grid integration costs while still beating fossil fuels). The DOE's target is 2030; TLC aims to beat that by five years.
Binning	A process for sorting components into various bins, in this case sorting cells into performance (efficiency or photocurrent) bins.
BK7	A common instrument-grade glass that has a low CTE (like Pyrex)
Bond pad	A small pad of metal, usually rectangular and typically on a chip (in this case a CPV cell) for a wire-bonder to bond one end of a bond wire to.
Bond wire	A very fine wire, typically uninsulated, used to interconnect components on a sub-millimeter scale. A wire bonder can place many wire bonds per second, with an accuracy of a few microns, including bonding both ends and trimming the wire.
Bus bar, busbar	A relatively large conductor that gathers current from or distributes current to a number of other (typically smaller) conductors.
Bypass diode	A diode that allows current to bypass a section of a module (in this case a receiver), when that section is performing poorly enough that it would diminish the output of the whole module.
Cassegrain	A system with a reflective optical path folded to keep the exterior envelope short, originally used for compact telescopes but adopted by PARC (and SolFocus) for reflective CPV systems that look like thick flat panels.
Cell	A photovoltaic cell.
Cold plate	A cooled plate to which cells are attached to keep the cells cool. Unless otherwise specified, a micro-channel cold plate, which is a cold plate through which water is pumped in sub-millimeter channels to more rapidly absorb heat. See also Receiver and micro-channel matrix.
COMSOL	A multi-physics program that can simulate complex interactions between components, such as the change in the focusing of a lens as the light passing through it heats it.
Cover glass	A protective transparent cover over the optical entrance to the modules (used only for TLC on single-axis trackers). A cover glass will be AR-coated on both faces with a non-haze AR coating.
CPC	Compound Parabolic Curve. A parabola focuses parallel rays to a point; a compound parabolic curve focuses rays over an angular range to a short line. Versions focusing in two dimensions are also called 'Winston cones' (after Roland Winston who pioneered their use in CPV).
CPV, HCPV	**C**oncentrated Photovoltaic. A branch of solar energy in which light is concentrated onto PV cells. Although other types of CPV exist (e.g., LCPV), CPV is often used synonymously with HCPV in which the light is Highly concentrated onto the cells.
CPV cell	A high-efficiency cell that can =handle high concentration. While many types of cells could be used for CPV and even for TLC specifically HCPV, all 'CPV cells' discussed in this document are tandem cells.
CSOC	Concentrator Standard Operating Conditions. The conditions under which CPV systems are tested for comparative evaluations, of 900 W/m^2 DNI at AM1.5, and 20ºC ambient with 2 m/s wind speed.

Term	Definition
CSP / CST	**C**oncentrated **S**olar **P**ower / **C**oncentrated **S**olar **T**hermal. The concentration of sunlight to produce heat that is then converted to electric power. While TLC is NOT solar thermal, it can borrow infrastructure from solar thermal troughs.
CTE	**C**oefficient of **T**hermal **E**xpansion – the amount by which a material expands when heated, typically in units of parts per million (ppm) per Kelvin (the same a ppm per degree Celsius). Unless otherwise specified, in this document CTE refers to linear expansion.
Curamik	A German company that makes DBC/AlN and cooling systems that use DBC/AlN. Curamik is extremely cautious, so if something fits within Curamik guidelines it can be relied on to work.
Current-limiting	When electrical components are in series, current through any must flow through all. If one component cannot carry as much current as the others, it is referred to as current-limit. For CPV cells this can apply to either cells in series or to junctions in series within a cell. Since TLC has cells in parallel on a receiver and the receiver-to-receiver series is protected by bypass diodes, TLC is not sensitive to current-limiting (e.g., shaded or low-performance) receivers. For tandem cells typically the bluer junctions are current-limiting and the redder junctions are not current limiting; while Spire 3J cells had a current-limiting 3rd-junction, this was an exception and the 3rd-junction usually has excess photocurrent.
Current Matching	When electrical components are in series, current through any must flow through all. Getting the currents to match without impeding some components is referred to as current-matching.
Cyrium	A Canadian CPV cell company that designed high-efficiency tandem cells that could be made at contract fab rather than requiring a dedicated solar cell fab.
DAT	Dual-axis tracker. Unless otherwise specified, a dual-axis tracker stiff enough and accurate enough for HCPV.
DBC/AlN	Direct-Bond-Copper on Aluminum Nitride. A standard highly-thermally conductive electrical insulator (AlN) clad with electrically-conductive copper layers that can be patterned and soldered or wire-bonded to.
Declination	The angle of the sun to the north or south of the celestial equator. See also 'Tilt'
Dicing	Cutting of a large sheet into smaller units. Typically used for the sawing of a wafer into individual solar cells, but can also be used for other types of cutting (e.g., TLS) or for cutting lens sheets or backplane sheets into lens tiles or backplanes.
Dicing Street	When a wafer is saw-cut into cells, the saw blade can wobble slightly. A 25µ saw blade is this typically given a 40µ to 50µ 'street' to saw within.
Diffuse	Diffuse refers to sunlight that is not directly from the solar disk but is scattered by the air, clouds, dust, etc. before reaching a surface. See 'GHI'.
Diode, Bypass Diode	If a receiver can't generate enough power to contribute to a module (such as if a cell is defective or bird droppings block its light), the other receivers could try to drive current through it, which could further damage the receiver and could cause the whole module to become non-productive. The bypass diodes are one-way valves for electricity that will let electricity flow around that receiver rather than through it during such conditions (but not flow in the other direction in normal operation).
Direct	For light, direct refers to sunlight that comes directly from the solar disk without being scattered by the air, dust, clouds, etc. See DNI.

Term	Definition
Dispatchable	An electrical generation source that can be turned on or up as needed is referred to as dispatchable. For example, batteries supply dispatchable power, as does solar thermal with storage. In contrast PV is not dispatchable because the power is produced when the sun is shining and not on demand, and is a form of intermittent power.
DNI	Direct Normal Insolation – the strength of the highly focusable sunlight direct from the disk of the sun. This differs from global insolation, which also includes scattered light.
DOE, DoE, D.O.E.	"The DOE" refers to the U.S. Department of Energy, a major sponsor of solar energy R&D. 'DoE' can also refer to "Design of Experiment".
Embossing	A process for pressing a pattern into a (typically flat) substrate. See Hot Embossing.
End loss	End loss refers to the light lost off the end of a trough due to the slant of the sun.
EVA	Ethylene Vinyl Acetate, a clear, soft-but-tough plastic use in silicon solar modules and also occasionally in hot-embossed CPV Fresnel parquets.
FEA	Finite Element Analysis – a computational tool for simulating what happens (heat flow, stress, optical paths, etc.) in complex parts by dividing those parts into small elements and analyzing each element and its interactions with its neighbors.
Feina	A tracker company that makes two-axis CPV tracker, single-axis CPV trackers and single-axis flat panel trackers, which has proved useful in comparative costing.
Fin tube	Pipe (typically metal) that has fins around it to increase the surface area. Fins can be pressed disk or can be a continuous spiral; spirals can either be threads on the parent tube or can be spiral-wound around a tube. L-foot spiral-wound fin-tube is currently specified for TLC its durability.
Flange	A relatively narrow edge of a larger structure, typically bent at 90º to the rest of the structure. The flanges of an I-beam, for example, are the narrow extensions at the top and the bottom of the 'I'. See also Web.
Flowing film	SAT TLC uses an unusual type of micro-channel cold plate in which a manifold forces coolant to flow in a thin film under each row of cells. To attain rapid heat transfer, the flowing film is very thin/shallow (about 60 μ). See 'Micro-channel'.
FR-4, FR4	The green fiber-reinforced material that most circuit boards are made from.
Fraunhofer	A Germany-based research institute active in CPV research.
ga	Abbreviation for gauge. Generally refers to thickness of steel or size of wire, can also refer to size of needle for dispensing adhesive. High gauge numbers are thinner/smaller.
GaAs	Gallium Arsenide. The standard workhorse III/V semiconductor used for high-efficiency single-junction CPV cells, and as the basis for some multi-junction cells, and thus the lowest-risk and least-expensive to prototype with.
Galvanic corrosion	Corrosion between dis-similar metals in contact with a liquid can be greatly accelerated if they are in electrical contact, allowing electron flow between them to match ion flow within the liquid.
GHI	Global Horizontal Insolation – the total sunlight, both direct and diffuse, reaching a flat surface at a given site.
Glasshouse, Glass house	A glass-roofed, and typically glass-walled, structure (such as a greenhouse).

Term	Definition
Glasshouse solar	Solar energy systems where the collector is in a glasshouse.
GlassPoint	A CST company building large arrays of troughs in glasshouses.
Grid line	A narrow conductive line on top of a cell, a few microns wide by a few microns thick.
GW, gigawatt, GigaWatt	A billion (10^9) Watts, about the size of a typical coal-fired or nuclear power plant.
HCPV	High-Concentration PV. CPV at high concentration, often defined as over 300x.
Heat pipe	A heat transportation component that evaporates a coolant as the heat source and condenses the coolant where the heat can be gotten rid of. The latent heat of evaporation is typically very high so this can move heat very effectively. Heat pipes can use gravity return or wicking return.
Heat spreader	A thermally-conductive sheet that spread the heat from the solar cells in a TLC receiver. Also called a 'Backplane' for its electrical function.
Hot Embossing	A process for pressing a pattern into a (typically flat) substrate where the substrate is heated to soften it (or a layer of it) before pressing. Hot embossing is used to make the lens parquets in typical Fresnel-lens HCPV.
IR-matched, Refraction-matched	If two materials have different indices of refraction, light is reflected from the interface between them. Materials are said to be IR-matched when their indices of refraction are similar. See also AR Coating.
I_{MPP}	Current at maximum power point – the current that a photovoltaic cell produces at its maximum efficiency under a given illumination.
Intermetallics	Compounds of mixed metals, typically with less-desirable properties that either of the parent metals.
Inverter	A device for converting DC power (from photovoltaic cells) to AC current to feed the power grid. Matching the input voltage from a solar inverter allows solar inverters features, such a maximum-power-point tracking, to be used to maximize power output (however if the core's power output is steady then such features could be omitted for volume production when an inverter can be customized).
I_{SC}	Short-circuit current – the maximum current that a photovoltaic cell can produce.
J, Jct.	Used with a numerical prefix (e.g., 4J) to refer to the number of junctions in a tandem cell. See 'Junction'.
J_{SC}	Short-circuit current per area – the maximum current that a photovoltaic cell can produce per area, typically shown as per square centimeter.
Junction	A power-producing layer of a photovoltaic cell. Each layer adds efficiency at the expense of complexity, so using a high-efficiency single junction is ideal in prototyping and probably even in production.
Junction box	A box that protects the connection between two solar modules or between solar modules and an inverter.
Kansas City	A site that the DOE uses to represent a typical U.S. location with moderate sun quantity and moderate sun quality.
Kerf	The part of something being cut that is removed by the cutting, typically becoming sawdust.
kW	1000 Watts of power.

Term	Definition
kWh	The energy supplied by a kilowatt of power in one hour.
Laminar flow	Flow can be laminar or turbulent. In very fine channels, there is no room for turbulence to develop and the coolant flows smoothly and with little friction.
Laser-cut	A fairly low cost, fairly high accuracy cutting process that uses a laser to cut a flat sheet of material. Laser cutting is practical for very thin sheets to moderately thick sheets of many materials, including glass and steel. Good laser cutting has an accuracy of ~5 mils (~125μ), with typical inaccuracies of 2 mils (50μ). Steel sheets up to 5/8" (16 mm) can be laser cut. See also 'Plasma-cut' and 'Water-jet cut'.
LCOE	Levelized Cost of Energy. The cost of the energy output, typically per kilowatt-hour, when all planning, permitting, capital, financing and operational costs are included.
LCP	Liquid-Crystal Polymer. A low-CTE plastic whose properties are great for molding and machining.
Lens tile	A glass tile the size of a module (DAT) or a single receiver (SAT), with a lens pattern formed in it, Each DAT lens tile in the current design converters an RP-3 trough mirror's focus into roughly 3000 small high-intensity focal spots.
Light-guide optics	A flat alternative to a lens that uses internal reflective surfaces to guide the light to a cell. This could allow somewhat larger cells (making cooling rather than lens thickness the limit).
Low-iron glass, solar glass	Glass that has a low iron content. Iron absorbs longer wavelengths of the CPV-useable spectrum (giving glass a greenish tint when seen edge-on), and low-iron glass absorbs less and is thus used for solar panel module cover, CSP mirrors, etc.
Manifold	A manifold distributes a large flow into numerous smaller flows. When used without qualification, it refers to a glass manifold tile within a SAT receiver (although SAT TLC's coolant pipes feeding and draining numerous tubes are also manifolds).
Maxeon	The highest-efficiency commercial silicon solar panels, or a silicon cell thereof.
Micro-channel	A single few-millimeters-long channel through which coolant flows. In SAT TLC there is a flat micro-channel under each row of cells. To attain rapid heat transfer, the micro-channel is very shallow (about 60 μ), so to keep the pressure drop low the channel is very broad (about 560 mm), and fairly short (about 10 mm). The coolant flow as a film through this channel, producing a 'flowing film' cold plate.
Micro-channel cold plate	A cold plate that forces coolant through very narrow channels in laminar flow to absorb heat. The narrower the channels, the faster the heat transfers, so channels are typically comparable to the thickness of a sheet of paper. Narrow channels have high pressure drops so a manifold is typically used to divide the total flow into numerous parallel small flows that each flow only a short distance in the narrow channels. DAT TLC avoids micro-channels entirely. SAT TLC uses a micro-channel cold plate with a single flat channel under each cell instead of a series of vertical channels. With the semi-dense array of tiny cells this produces excellent cooling at very low cost.
Micro-channel matrix	Most micro-channel coolers use a matrix of fine channels in formed by diffusion-bonding stacks of etched sheets of copper. *TLC avoids the cost of a micro-channel matrix* by taking advantage of the narrow cells to use a heat spreader to gain area for heat transfer to heat rejection.

Term	Definition
Micro-tracking, μ-tracking	While most PV tracking moves a whole system of many modules by several meters each day to follow the sun, micro-tracking tracks the sun by moving much smaller units (in this case, receivers) a few centimeters per day within a module. TLC uses standard single-axis trackers for the first axis, increasing energy capture (especially in the key late-afternoon period), while micro-tracking on a second axis to avoid the high cost of a two-axis tracker.
Micro-transfer printing	A process for transferring large numbers of tiny units (such as tiny CPV cells) in parallel. Micro-transfer printing can have better than 2-μm placement control.
Mirror-Module	A TLC module plus its mirror (useful in cost comparisons because a mirror is part of the module's optics).
Module	The field-replaceable unit containing 32 receivers (DAT) or ~ 100 receivers (SAT)
Morgan Solar	A pioneer company in light-guide optics, which are very flat means for concentrating light. Morgan is used both to set a bound on TLC lens cost (similar molding requirements as Morgan but 25 times less area), and as an example of a n alternative type of lens that would allow slightly larger cells while still being able to micro-track.
Mrads	Milliradians, or thousandths of a radian. A very useful unit for measuring small angles, especially in solar where the sun's angular diameter is close to 10 mrads. For small angles it is very easy to calculate angles in mrads - an angular error of a micron per millimeter (or a millimeter per meter) is one mrad of error.
MW	MegaWatts, or millions of Watts. A common unit for measuring power plants. Utility-scale solar farms typically range from tens to hundreds of megawatts.
Nano-porous	Having holes or a porous texture on a scale of much less than a micron. Since the wavelengths of light of interest for CPV are a few hundred nanometers, a surface (glass, for example) textured on a nanometer scale can provide a 'gradual' change (from air to glass) from the perspective of light, which leads to much lower reflection loss. CST trough receiver tubes typically have nano-porous surface anti-reflection coatings; however, these are 'hazy' enough to not be suitable for subsequent concentration and hence are NOT used in TLC.
Nano-textured	Having a surface texture on a scale of much less than a micron. This includes both nano-porous 'holes' and nano-protrusions. Nano-protrusion surfaces ('moth-eye' coatings) can produce very low reflection over broad angular and wavelength ranges and would significantly increase the efficiency of TLC. However, it is not known if they are cost-effective yet for the areas needed, so they are not currently used in the TLC design (but are expected to be used in the future).
Off-axis	An off-axis parabolic section has a focus that is not directly in front of the parabolic section but is off to one side (in other words, the part of the parabola in line with the focus is not part of the section). This allows a module at the focus to not shade the parabolic surface.
Optically coupled	Joined with an interface that minimizes light loss, typically through matching indices of refraction.
Optical epoxy	An extremely clear epoxy whose index of refraction is well-matched to glass.
Optical silicone	An extremely clear silicone whose index of refraction is well-matched to glass.
Overmold	A process in which something is molded on top of a solid core or a solid insert. For examples, Silicone-on-Glass lenses are made by molding silicon onto a glass core.

Term	Definition
Pad	A small pad of metal, usually rectangular and typically on a chip (in this case a CPV cell) for a wire-bonder to bond one end of a bond wire to.
PARC	Palo Alto Research Center (formerly Xerox PARC). A well-known silicon-valley research company (responsible for such innovations as the graphical user interface and the computer mouse), with a strong optical systems group.
Parquet	A molded or over-molded lens tile for Fresnel CPV.
Parylene	A conformal coating that can be deposited in extremely thin, smooth layers and acts as a chemical barrier and permanent mold release to which silicone will not stick.
Performance binning	A process for sorting components into various bins based on their performance.
Petzval, Petzval effect	The change in the focusing of a lens when light reaches the lens at a slant. The Petzval effect broadens approximately with the square of the angular range of the light, which drive TLC to low rim angles (e.g., a single RP3 inner mirror, whereas a CST RP3 trough uses two inner and two outer mirrors all focusing on the same receiver trough.
PEX	Poly-Ethylene (PE) cross-linked (X). A low-cost plastic pipe installed with compression bands. While PEX itself is not used in TLC, the compression band style is used as one of the coolant connection options.
Pick-and-place, Pick-n-place	Pick-and-place refers to the process or the equipment used to place numerous small components on circuit boards. Pick-and-place equipment is used to place over 3 *trillion* components per year that are roughly the size of TLC cells.
PMMA	Poly-methyl-methacrylate. A fairly durable, very clear plastic used in some CPV systems. Plexiglas is the best-known brand of PMMA.
PPA	Power-Purchase Agreement. Many large utility-scale solar installations have power-purchase agreements where a power company guarantees to by the output at a given price (usually with an inflation escalator) over a long period (often 20 or 25 years).
PPG	A large U.S. solar mirror company.
PSI, psi	Pounds per square inch. A unit of pressure, or of material strength.
PVRD	PhotoVoltaic Research and Development, a series (two so far) of early-stage DOE funding opportunities for solar energy R&D.
Receiver module	The field-replaceable unit containing 32 receivers (DAT) or ~100 receivers (SAT)
Receiver	The small unit within a module whose cells are in parallel. In SAT TLC, the receivers tilt to track the sun on the second axis.
Receiver Tilt	The tilt (slant) of a receiver relative to the axis of the trough, typically to the north or south relative to a vertical east/west plane. For a receiver to focus properly, this must match the Solar Tilt. For a typical trough (horizontal and tracked east/west), this is the angle of the sun to a vertical east/west plane (more generally the angle relative to a plane perpendicular to the axis of the trough). This depends not only on the sun's declination, but also on the latitude and time of year and time of day; for an E/W trough it reaches its maximum one direction at solar noon on the winter solstice and its maximum in the other direction at dawn and dusk on the summer solstice. See "Solar tilt'.
Relative Declination	The angle to which the micro-tracking must rotate the receivers to point them at the sun. Now called 'Receiver Tilt'.

Term	Definition
Rim Angle	In TLC, rim angle is defined in the optical sense as half the angle that the receivers see from one trough rim to the opposite trough rim (when the receiver is at zero tilt (sun due east/west) in SAT). While for an off-axis trough this introduces slight differences if one uses the rim angle in some calculations (for example, using the rim angle and focal length to calculate trough depth for 'smile aberration' calculations, for the modest rim angles used in TLC the differences are not significant.
RioGlass	A large European solar mirror company with U.S. manufacturing.
Roll-formed	Formed or patterned by passing a deformable sheet through rollers – a process that is low cost even in moderate volume and is very low cost in high volume.
RP-3, RP3	A standard size of trough and trough mirror for solar thermal (CST/CSP). An RP-3 trough normally has two rows of RP-3-inner and two rows of RP-3-outer mirrors for a total width of 5.76 meters, all focusing on a single receiver-tube at a 1.71-meter focal length. This produces a much higher rim angle than is suitable for TLC, so TLC uses a single row of RP-3 inner mirrors per trough focus for an initial product.
RP3-inner	Just an inner mirror, or a row of such mirrors, from an RP-3 trough.
SAC, SAC Solder	SnAgCu solder, a common lead-fee replacement for tin/lead solder.
SAM	System Advisor Model, an NREL program for planning and/or evaluating various types of solar installations
SAT	Single-Axis Tracker, generally with a horizontal N/S tracking axis for E/W tracking.
Semprius	A CPV company that has pioneered micro-transfer printing for handling tiny cells.
SiC	Silicon (Si) Carbide (C). A very hard semiconducting ceramic with a low coefficient of thermal expansion. While SiC is used in some power applications, in TLC it is merely a strong, high-thermal-conductivity, low-CTE filler in AlSiC.
Singulate	Singulation is the process of cutting individual units from something fabricated as a large block or sheet. Dicing of cells and cutting lens tiles from a lens sheet are examples of singulation.
SIPS	Small Innovative Projects in Solar. The smallest, earliest-stage division of the DOE's PVRD funding program.
Skiving	A process by which thin layers are sliced into or from a workpiece (much as a plane shaves wood from a surface. Skiving can be used to raise cold-plate fins from a block of copper, producing an aspect ratio suitable for TLC fins.
Smile aberration	The smile-shaped broadening of the sun's image on a lens, perpendicular to the trough's focus, from the part of the trough that the lens sees as illuminated not being on a plane perpendicular to the lens (due to the depth of the trough).
SOE	Secondary Optical Element – typically a small lens or TIR cone in Fresnel CPV, or the curved lens fronts in TLC.
Solar Tilt	(SAT-only) The tilt (slant) of the sun *relative to the axis of the trough*, which is or the north or south relative to a vertical east/west plane for a typical tracker that tracks E/W around a north/ south axis. This depends not only on the sun's declination, but also on the latitude and time of year and time of day and the orientation of the trough's axis; for an E/W trough it reaches its maximum one direction at solar noon on the winter solstice and its maximum in the other direction at dawn and dusk on the summer solstice. See "Receiver Tilt'.

Term	Definition
Solarphire	A PPG brand of very low iron glass, the best Solarphire grade is currently the best cost/performance balance for TLC. Solarphire has been rebranded Starphire for the architectural-glass market.
SolFocus	A CPV company (deceased) that tried to commercialize the Cassegrain design of PARC. SolFocus is used to set a bound on TLC lens cost since SolFocus also used a focusing sheet (although a mirror sheet rather than a lens sheet) and had a roll-formed sheet for their Gen-2 product; TLC used 40 times less sheet per Watt.
Spectral response	The power or current output of a solar cell as a function of the spectrum of the light that it is illuminated with.
Spectrolab	A Boeing subsidiary the manufactures tandem cells. A terrestrial concentrator version of Spectrolab's 5J space cell would set the CPV efficiency record, and is the reference cell expected for product introduction.
Specular	Reflecting without scattering, like a mirror (as opposed to like snow).
Starphire	A PPG brand of very low iron glass – a rebranding of Solarphire.
Stepper motor	A motor that can be turned in small precise steps as opposed to turning continuously. Stepper motors are much slower and are more expensive than rather than continuous motors, but provide much higher accuracy for micro-tracking (and still cost far less than a large macro-tracking motor).
Street	See 'Dicing street'.
SunLab	A university-of-Ottawa lab specializing in solar energy, including continuous 1000x concentration onto CPV cells.
SunPower	A U.S. solar company that produces the highest-efficiency commercial silicon modules (see Maxeon). SunPower also has LCPV systems that concentrate 7x onto silicon cells. Most SunPower manufacturing is in Malaysia.
SunShot	The U.S. DOE's main solar R&D support program. SunShot funding start at Tier-0 (comparable to 2nd-stage PVRD funding), and progresses through at least tier-3, with only 20% non-federal matching required at Tier-0 and Tier-1. TFI/xVI intends to file for either PVRD or SunShot funding each year.
Tandem cell	A type of solar cell in which several sub-cells are stacked so that each one can covert a narrower wavelength range that it is optimized for rather than having to handle a broad range. Tandem cells have reached 46% efficiency in the lab and 44% cells are available commercially.
TCV	Through-cell Via. While most CPV cells have one contact on back and one on front, TCV's can be used to re-map a contact to put both on front or both on back. TLC does not currently do this but could if TCVs become cost-effective.
Temperature Coefficient	The sensitivity of a cell's output to the cell's temperature.
Terminal wire	The interconnect wire from a receiver. The terminal wire is connected to the next receiver during module assembly.
Thermal conductivity	The ability of a material to conduct heat, typically measured in Watts per meter Kelvin (w/mK). Most plastics have a thermal conductivity of 0.1 to 0.3, glass is ~1, stainless steel is ~25, mild steel is ~50, AlSiC and aluminum nitride are ~170, aluminum is ~220 (although many alloys are significantly lower), copper is ~400, and diamond is around 2000.

Term	Definition
Thermo-optic effect	The change of the refractive index of a material with temperature. Glass has a very low thermo-optic coefficient of typically less than 10 ppm/K (with lens glasses being lower at a few ppm/K), while hard optical plastics such as PMMA are typically around 1000 ppm/K and silicones are typically around 350 ppm/K. Glass lenses are thus relatively immune to changing focal length with temperature while silicones are sensitive enough that this causes significant performance degradation in some Fresnel/Box CPV systems.
Thermosyphon	A heat pipe that uses gravity (as opposed to a wick) to return the working fluid. The TLC heat pipe is a hybrid, using gravity to return the fluid from the fin-tube condenser and using a wick to reach the high edge of the evaporator.
Tile	A lens tile. Although the backplane and the glass manifold are similar-sized and similarly produced as larger sheets that are then singulated, 'backplane' and 'manifold' are unambiguous whereas lens could refer to an individual lens, a lens row, or a whole lens tile. 'Tile' is thus used exclusively for the lens tile.
Tilt	Slant in a northerly or southerly direction. The slant of something relative to an east/west line or a vertical east/west plane. See "Solar tilt' and "Receiver tilt".
TIR	See 'Total Internal Reflection'.
TLS	Thermal Laser Separation. A process in which a laser is used to propagate a crack through a brittle substrate in a controlled manner, cleanly cleaving it. This is both lower cost and higher yield for small dies than diamond-sawing.
TOE	Tertiary Optical Element – A few CPV designs (especially semi-dense arrays like REhnu and TLC) have three optical stages. In TLC the TIR Winston cones on the back of the lens is an array of third-stage optical elements, or TOEs.
Total Internal Reflection	One of the few near-perfect things in nature, when light hits a surface to a lower-refractive-index material at a shallow angle it is essentially completely reflected.
Trackers Feina	A tracker company that makes two-axis CPV tracker, single-axis CPV trackers and single-axis flat panel trackers, which has proved useful in comparative costing.
Value (economic)	The value of a given solar module depends on many factors. For example, all else being equal a module of twice the efficiency takes half the installation labor for a given output, as well as half the shipping, half the tracker space and half the land area. Longer-lasting modules also have higher value. In the context of this document, value is defined as 'what a module could sell for and still produce energy (kWh) at the same cost (LCOE) as the lowest-cost alternative. It should be noted that value is thus site-dependant (temperature, latitude and diffuse light) and well as partly determined by the selling price of competitor's modules, which will continue to come down over time. Value sets an upper bound on selling price and thus on margins; especially initially, a newcomer typically has to offer a considerable discount relative to value to entice customers to try something new.
V_{MPP}	Voltage at maximum power point – the voltage that a photovoltaic cell produces at its maximum efficiency under a given illumination.
V_{OC}	Open-circuit voltage – the maximum voltage that a photovoltaic cell can produce.
W	Watt(s). Unless otherwise specified, Watts (and kW, MW and GW) refer to peak watts DC. See W_{ac}.

Term	Definition
W_{ac}	Watts AC – power output of the inverter, which can be significantly less that Watts DC due to over-provisioning (10% to 30%) and conversion loses (a few percent).
Wart	A box on top of a SAT TLC module, such as a junction box or a cover for a micro-tracking motor or for a pump. These have to be sealed, they interfere with the nice packing of modules in a rack, and they are also prone to being bumped. In reality even SAT TLC will have several times fewer 'warts' per kW than silicon PV, but warts are still aesthetically offensive even if economically insignificant.
W_{dc}	Watts DC. See 'W'.
Web	A relatively large center of a structure, typically with flanges attached. The web of an I-beam, for example, is the tall central part of the 'I'. See also Flange.
Winston cone	A 2-axis Compound Parabolic Curve (CPC) focusing arrangement pioneered by Roland Winston. TLC uses a row of Winston cones to guide the light from each lens onto a row of small cells, homogenizing and further concentrating the light.
W_{peak}	Watts peak. See 'W'.

LIST OF TABLES

Table 2.1: Cost Comparison between Silicon and TLC. 18
Table 2.2: Cost Comparison between Fresnel Lens HCPV and TLC. 19
Table 2.3: Advantages and Issues of PV types. 29
Table 3.1: Absorption of Optical Silicone. 64
Table 4.1: Voltage Matching to Inverter for 3J cells. 140
Table 4.2: Voltage Matching to Inverter for 5J cells. 140
Table 5.1: Table and Chart of Cost Estimates for TLC. 157
Table 5.2: Table of source and cost reductions used in cost model. 162
Table 5.3: Table of Value Adjustments for TLC. 166
Table 5.4: Table of tooling costs versus production per year. 174
Table 6-1: Module parameters if module is fabricated as designed 202
Table A.1: Table of Reliability / Durability Risks. 252
Table A.2: Table of TLC variations and effects. 263
Table B.1: Comparison table of DAT vs SAT TLC. 276

CHAPTER 1

1. INTRODUCTION

1.1 Context and Issues

1.1.1 Market Need

Solar energy has to potential to curtail global warming, and, if low enough in cost, to bring the whole world's population to a first-world living standard. However, in spite of the dramatic drops in solar costs in recent years (particularly in photovoltaics, or PV), solar is only competitive with fossil fuels in limited cases, and additional major cost reductions are needed.

Silicon PV has dramatically reduced costs largely through decreasing the cost and increasing the efficiency of the silicon cells, but silicon is nearing its theoretical efficiency limits, and even if the cells were free, silicon would not meet the DOE's beyond-SunShot targets without significant other improvements. This leaves a need for a way to reach much higher efficiency than silicon PV.

Tandem Concentrated Photovoltaic (CPV) cells are already roughly twice as efficient as silicon and have headroom for further improvement, and tandem cells are also durable for long product life. But previous CPV products have been unable to compete with silicon on cost in spite of the massive efficiency advantage, so there is a need for a much-lower-cost high-concentration photovoltaic (HCPV) architecture.

Trough mirrors are very inexpensive and using a low-cost trough for primary optics can shrink the rest of an HCPV module by 40X and thus dramatically reduces module costs. However, a trough alone cannot reach anywhere near high enough concentration, so there is a need for a low-cost way to reach very high concentration using a parabolic trough as the primary concentrator.

1.1.2 Trough-Lens-Cone (TLC) concept

As shown in Figure 2.1, the proposed concept consists of focusing a trough mirror onto a compact module whose cover is a 'lens tile' that further concentrates the light. The proposed lens tile

consists of a linear array of linear lenses to further concentrate the trough's focus into a linear array of lines, followed by an array of total-internal reflection cones to further concentrate each line into multiple spots. This configuration is called *Trough-Lens-Cone, or TLC* and was initially proposed and patents applied for by Richard Norman at Terra Firma Innovations. **See Figure 2.1**.

1.1.3 TLC Potential

TLC has the potential to meet the market needs by using a compact module at the focus of a low-cost parabolic trough mirror and reconcentrating the trough's focus to reach high concentration. TLC can potentially do this with low-cost optics that can reach the CPV cells' limit of 1500X, reducing cell cost, with high optical efficiency, and while thoroughly protecting the cells for long life. TLC can potentially do this while using high-volume electronics-assembly infrastructure for the rest of the module, and while producing voltage and current suitable for standard inverters.

TLC also can have high local manufacturing content, and the hardest component to localize (the cells) is also one where North America and Europe lead. TLC should take far less energy to produce and thus generate far fewer greenhouse gases in production, and hard-to-recycle materials are an even smaller fraction than for regular PV. And with less area per kWh, TLC is more amenable to being raised high enough above the ground for land co-use as well.

While TLC's advantages come with trade-offs such as small cells and finned or active cooling, the compact module makes these affordable, and no technical barriers have been found. Pursuing TLC is thus well justified, and TLC stands a good chance of being a key technology in bringing solar energy to where it can beat fossil fuels, not just in sunny places with subsidies and low grid penetration, but in marginal-sun areas, un-subsidized, and with enough cost advantages to fund storage and long-distance transmission to support grid integration at high penetration.

The primary requirements for TLC are that TLC must be able to produce electricity at very low cost, and that TLC must be able to be delivered as a launchable product. It would be of no use to just beat other forms of CPV on cost – the cost of power from TLC must also beat silicon PV. And not just where silicon PV is today but beat *what silicon PV will cost* when TLC is delivered.

Furthermore, TLC must be able to do this when it is made at an introductory scale, and with enough profit margin to fund growth. Since silicon PV makers have historically been willing to adopt

unsustainably low margins in a quest for market share, TLC made at tens of megawatts must have a significant cost advantage over future silicon PV made at hundreds of gigawatts.

While that sounds like a tall order, there are some ameliorating factors. The target is utility-scale, where customers care about LCOE, or ¢/kWh, rather than ¢/W, and TLC's high efficiency reduces installation labor, shipping, land acquisition and permitting costs per watt and makes two-axis trackers affordable. TLC can build volume in markets with high sun quality, and push into lower-sun-quality markets as manufacturing scale reduces TLC costs.

1.1.4 Itemized Requirements

- Produce electricity at low cost:
    - Very low production cost in high volume
    - Long product life
    - High efficiency
    - Easy to install
    - Low maintenance cost
- Launchable Product:
    - Reasonably low cost in low (10 MW) volume
    - Low capital cost for scaling
        - Fits current infrastructure such as trackers and inverters
        - Partners in related fields bear scaling costs as normal production
    - Enough protectable intellectual property to allow fundraising
- Nice-to-have:
    - High local manufacturing content and high U.S. / Canada content
    - Land co-use, Low embodied CO_2, Recyclable materials, Heat byproduct option

1.1.5 Single-axis or Dual-axis Tracking

TLC can be used on two-axis trackers, allowing the TLC module itself to be static, or on one-axis trackers, with 2nd-axis micro-tracking within the TLC module itself, with these designs being roughly equal in LCOE (Levelized Cost Of Energy = lifetime cost per kWh).

The dual-axis-tracked version offsets the significantly higher cost of two-axis tracking by having extremely low module cost; this produces higher margins for TLC itself even if the price is lower to compensate for needing two-axis trackers. Use of trough mirrors on two axis trackers makes both heliostat companies and trough mirror companies potential partners.

The single-axis-tracked trough version uses low-cost single-axis trackers and capitalizes on a CSP trough form-factor to use CSP infrastructure and gain CSP trough companies as natural partners, but it adds the cost and complexity of micro-tracking as well as increasing assembly and cooling costs and having lower optical efficiency and a lower capacity factor. However, a recent design advance micro-tracks only small mirrors within the module rather than micro-tracking the mini-lens-tiles with their receivers, so the new design avoids micro-tracking part of the cooling system.

Which version will ultimately have a lower LCOE depends largely on tracker costs. The two-axis-tracked version is simpler, so it is generally discussed first, with the single-axis version mentioned where there is a significant difference. Given the roughly equal LCOE, the two-axis version is likely to be preferred even at large scale, and its simplicity and more-forgiving tolerances mean that it should be built first even if the single-axis version ultimately is lower cost. However, if heat is a useful byproduct (e.g., for desalination), then the single-axis-tracked versions gain a significant advantage, especially with the new micro-tracking design that does not micro-track the cooling.

1.2 Research question

However high TLC's potential, until the start of this PhD research project TLC only existed as a paper study. Would it work out as well as it appeared, or were there hidden flaws that preclude beating silicon PV on cost? Many designs look good on paper, but as Admiral Hyman Rickover said of nuclear reactors (Rickover, 1953):

"An academic reactor or reactor plant almost always has the following basic characteristics: (1) It is simple. (2) It is small. (3) It is cheap (4) It is light. (5) It can be built very quickly. (6) It is very flexible in purpose ('omnibus reactor'). (7) Very little development is required. It will use mostly off-the-shelf components. (8) The reactor is in the study phase. It is not being built now. "

"On the other hand, a practical reactor plant can be distinguished by the following characteristics: (1) It is being built now. (2) It is behind schedule. (3) It is requiring an immense amount of

development on apparently trivial items. Corrosion, in particular, is a problem. (4) It is very expensive. (5) It takes a long time to build because of the engineering development problems. (6) It is large. (7) It is heavy. (8) It is complicated. "

1.3 Objectives of the research project

The objectives of the research project are to:

- build a unified analytical model and 3D-CAD models to refine the TLC design;
- evaluate its feasibility experimentally by prototyping individual parts and processes;
- demonstrate its capabilities by building a physical prototype and putting it on sun.

This experimental approach is necessary because even after TLC has been "built" many times, in visualization, on paper, in analytical model spreadsheets, and then in COMSOL, until TLC is physically built TLC is just a paper CPV system and hidden flaws could arise at any time. The only way to ensure that TLC is feasible is to build it.

1.4 Original contributions

The initial TLC concept, the unified analytical model to guide the design from a cost perspective, the ray-traceable 3D model that allowed refining the TLC design from an optics perspective, and the proposed manufacturing techniques were all the work of the current PhD candidate. During the research project, the detailed TLC design has been extensively refined and many issues have been resolved along the way; while a team of students, professors and research professionals has been involved, this design work has been led by the current PhD candidate, who also led the writing of the published papers on the work (under the guidance of his supervisors). With regard to the physical prototyping and characterization of the prototypes, medical issues have limited the PhD candidate's current involvement to defining the experiments, designing the test setups, as well as analyzing the results.

Overall, the key original contributions that will be presented herein are:

- The lens-tile optics, and potential low-cost, high-accuracy production methods for these

- The gap-free flexible-cone coupling of the lens tile to the microcells
- The compact massively-parallel microcell-array receivers
- The receivers interconnecting in series as they are bonded to the lens tile

1.5 Outline of the document

With its three-stage optics and its massively parallel microcell receivers, TLC is different enough from traditional CPV that Chapter 2 starts with an overview of TLC as a whole. Because mainstream CPV designs are much costlier than silicon PV, many have dismissed all CPV as an expensive technology, so the overview includes a brief economic and business case overview to show that this novel form of CPV has the potential to be competitive with silicon PV.

Chapter 2 then has a summary of competitive architectures to TLC. This is after the TLC overview so that the strengths and weaknesses of each competitive architecture can be compared to TLC.

Chapters 3 and 4 then follow the flow of energy through the TLC design, through the multiple stages of the optics, through the cells and their electrical interconnections, and through the thermal path to cool the cells and the module. This is followed by details of the mechanical aspects of the proposed TLC module, and the proposed assembly process for the module. The TLC design featured in this document requires dual-axis tracking, so this document then discusses the suitability of various dual-axis trackers for TLC. Multiple solutions are presented when relative costs can change enough that what is less favored now can become more favored later.

To have an impact on the world, TLC modules will need to be sold to customers, shipped to installation sites, installed on trackers, and remain operational for decades, so Chapter 5 discusses module shipping and installation, and cost, value and reliability relative to silicon PV. Chapter 5 then discusses scaling factors and TLC variations to guide future versions of the TLC design.

Chapter 6 describes work done to validate the TLC concept and the basic TLC design, starting with the unified analytical model that guides the design and estimates the production cost of a given design. **This model (spreadsheet) is archived in *Savoirs*, at the Université de Sherbrooke (https://savoirs.usherbrooke.ca/).** This is followed by a summary of early multi-physics modeling and analysis done by the current PhD candidate. Chapter 7 summarizes the physical prototyping and testing done as part of the research project, and Chapter 8 concludes the thesis.

CHAPTER 2

2. STATE OF ART

TLC is a novel architecture for CPV. It is proposed that three-stage concentration can produce a lower overall cost and the traditional one-stage and two-stage concentration architectures used in CPV, and this chapter outlines how using initial concentration by a low-cost trough mirror to shrink the rest of a CPV module by ~40X can potentially dramatically reduce overall CPV module costs. This chapter proposes an initial TLC product and outlines a business case for that product as a viable competitor to silicon PV. (The models and validations for this are covered in Chapter 6.)

This chapter then covers the prior art of architectures related to TLC, as well as the context of numerous competitive architectures currently on the market or proposed as market contenders. This chapter then compares TLC to these prior architectures and competitive architectures, and shows that if TLC can achieve its goals, it will be better than any of the known competition.

2.1 Overview of TLC (Trough-Lens-Cone) CPV

2.1.1 Tracking

TLC can be use either a dual-axis-tracker, just as traditional CPV modules do, or TLC can use a single-axis-tracked trough with internal micro-tracking. With current tracker and cell prices, the dual-axis-tracked version of TLC will produce electricity at lower cost. Dual-axis-tracked TLC, or DAT-TLC is also a simpler system, and many of the principles are the same, so outside of a few references to SAT-TLC the single-axis-tracked version is relegated to an appendix.

2.1.2 Low-cost Focusing Principles

TLC uses novel three-stage '**TLC**' optics: a parabolic trough (**T**) focuses onto long, narrow modules. Each module has a lens tile with lenses (**L**) on the front that concentrate the light on the

second axis, and cones (C) on the back of the lens that further concentrate the light onto micro-cells on small receivers. Troughs are low cost and the trough's focus reduces the lens area by ~40x.

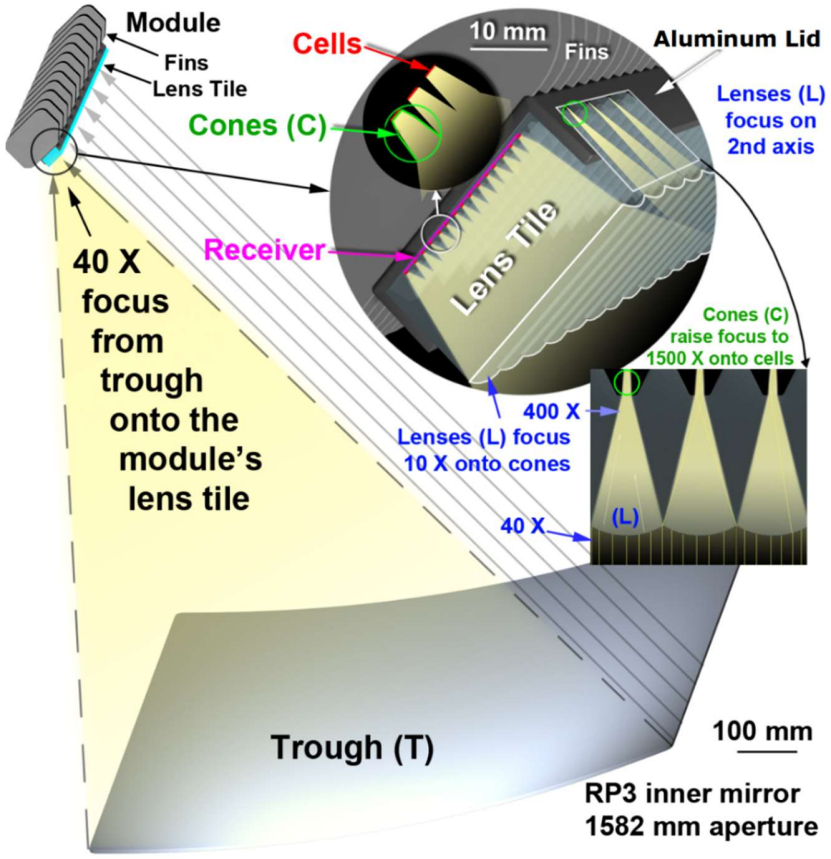

Figure 2.1: *TLC Overview Illustration (TFI illustration, first published in* (Norman et al., 2018)*).*

In this overview illustration, a trough mirror (T) focuses sunlight ~40X on one axis onto a compact module. The module's front is a 'lens tile' that has linear lenses (L) on front that focus ~10X on the *second* axis. Lens focusing is shown in the lens-tile zoom at the top right of the illustration, where part of the aluminum-finned steel lid has been cut away to show the focusing of the lenses.

Each lens focuses on a row of cones; each cone (C) further concentrates the light as it funnels light onto a single cell, reaching 1500X. The lens-focusing detail view on the middle-right of the illustration is a slice ***along*** the trough's focus: each of the three lenses shown is focusing onto a row of cones seen end-on. The main lens-tile zoom show a slice ***across*** the trough's focus; here the 16 cones of one row of cones can all be seen.

- ***Lens Tile:***

The light from the trough is re-focused by lenses roll-formed in a glass lens tile. Large sheets of glass can be roll-formed at low cost, and each square meter of lens sheet will produce almost 15 kW of lens tiles, making the cost per Watt extremely low; with the trough having reduced the lens area needed by 40x, the glass lens tiles will cost less per Watt than the *cover glass* of a silicon panel. The lens tile is thin enough to keep the optical path short and optical losses low when using low-cost low-iron glass (the same type of glass as flat-panel module covers).

In single-axis TLC each receiver has its own lens tile, while in dual-axis TLC the lack of moving parts allows all lenses for a whole module to be merged into one module-sized lens tile. In either case the lenses on the front of the tile focus onto an array of small total-internal-reflection (TIR) Winston cones (over-molded on the back of the lens tile that further concentrate the light while aligning it to the cells. The cones can be formed with silicone-on-glass over-molding similar to the Fresnel-lens industry, but at much lower cost due to ~40 times less area per Watt.

- ***Receivers:***

  - Cells:

Each TIR cone guides the light onto a small tandem solar cell; these cells are commercially available at 44% nominal efficiency and have reached 47% in the laboratory. The small lenses and their high concentration into narrow focal lines makes the cells ~0.65 mm wide by ~0.9 mm long. While cells less than 1 mm^2 sounds like too many cells to be practical, at 1500x each small cell produces 10% as much power as a silicon wafer that is over 25,000 times larger. And the small cells can be placed with standard high-speed electronic-component placement equipment that is currently used to place over 3 *trillion* small components per year, so placing TLC cells costs *several times* ***less*** *per Watt* than placing fragile thin-wafer silicon cells in a silicon module.

  - Electrical Connections:

Within a receiver, each cell has two wire-bonds, or 192 wire-bonds per substrate. Wire bonds are very cost effective, and 15 trillion wire-bonds are made on similar-sized substrates each year.

All cells on a receiver are in parallel, eliminating current matching issues within a receiver. The intensity along the focus of a trough is naturally fairly even, and a receiver averages out six rows for a very even total current. The receivers are in series (as wafers in a flat panel are) along the

trough, so the voltage builds rapidly. A module of many receivers looks electrically much like a regular solar panel, so TLC uses the same solar connectors and inverters that flat panels do.

- Cell Cooling:

A receiver's cells are mounted on a ~half-millimeter-thick copper heat spreader (low cost and reasonably thermal-expansion-matched to the glass lenses and the tandem cells). Thermally conductive, electrically insulating 'pre-preg' epoxy between the heat spreader and the module back provides electrical isolation at minimal thermal cost. An aluminum-finned heatsink 'module lid' hermetically seals to the lens tile; its fins have an aggregate area significantly larger than the trough mirror, and thus provide excellent cooling through natural convection.

- ***TLC Module:***

The dual-axis-tracked module is very simple, with no moving parts. While the module could be any length, a module matching the 1.7-meter mirror segment length is expected to be optimal, being easy for a single person to handle, and producing just over 1 kW under full sun.

- Connectors:

The currents and voltages look much like strings of standard flat panels, so standard weatherproof 1500V solar connectors are used, with only 1/3 as many needed per kilowatt as flat panels use.

- Cooling:

The module back is a finned heat sink. The ~1-mm-thick aluminum fins are 140 mm wide and 75 mm tall and have an aggregate surface area almost twice the mirror aperture area. The fins are also always vertical, so this provides excellent cooling through natural convection.

- Module size:

The fins dominate the module; the PV part (the lens and cell) is only 45 mm wide by 22 mm tall. The finned module is thus 140 mm wide by ~100 mm tall, by 1720 mm long.

- Mirror:

The mirror is much larger even than the cooling, being 1582 mm wide by 1700 mm long. The mirror is generally included when module **cost** is discussed because it is an integral part of the module optics. When there might be confusion, the terms mirror-module is used when the mirror is included, and receiver-module is used when the mirrors are not included.

2.2 Initial Product

The current plan is to introduce a module designed to sit on a two-axis tracker first because the development effort is much less and there are fewer things that can go wrong. A ~1700-mm module matches the 1700 mm standard RP3 individual mirror length, and will have 32 receivers each 53 mm long Today's best commercial CPV cells have a V_{OC} (open circuit voltage, or maximum voltage produced) of 3.13 V, so this gives a V_{OC} of 100V per module, and the receivers' 12-amp I_{SC} (short-circuit current, or maximum current produced) fits a normal low-cost weatherproof solar connector; with current cells the overall DC efficiency should be ~38% for just over 1 kW per module. With the 5J cells expected by product introduction, the V_{OC} will be ~6V, so the same receiver length would produce a module V_{OC} of ~200V and have a current of only 7A. The efficiency is then expected to reach ~41%, for a ~1.1 kW module. In either case modules can be strung together as needed to reach the voltage for feeding an inverter.

2.2.1 Summary of Inputs and Outputs

- **Optical:** The TLC system provides very high concentration (1250X to 1650X, with 1500X initially targeted) from very low-cost optics.
- **Mechanical**: The mechanical form factor for a DAT TLC module is 1720 mm long by 120 mm wide by ~100 mm tall, which is ten times more Watts per volume than SunPower's best traditional flat panels. The mirrors pack densely enough to be weight-limited in shipping, and 350 kW of mirrors plus modules fit per 20-foot container, an overall shipping volume four times lower than with the best flat panels.
- **Electrical:** With the cells expected at introduction, a DAT TLC module will produce ~1.1 kW (~7A at ~150V); with today's 3J cells it would be ~1 kW (~12A at ~86.5V). Normal solar connectors connect modules in series to reach a suitable voltage for a given inverter.
- **Power:** Unless otherwise specified, power calculations are made for a DNI of 1000 W/m^2 to ensure adequate electrical and cooling performance under demanding conditions. The analytical model allows adjusting this to any value desired.

Many economic comparisons are adjusted to a DNI of 766 W/m^2 because when this de-rated DNI is used to calculate the module power, the capacity factor (and thus the annual energy output)

matches silicon flat panels on single-axis tracker and the U.S. DOE's reference location of Kansas City Missouri, making fair comparisons easy. However, that number is site dependent, and the fair-comparison-DNI to rate at is much higher in desert regions.

- ***Energy and Cost:***

For energy and cost comparisons such as LCOE, power outputs are de-rated on a site-by-site basis based on DNI. Unless otherwise stated, the cost comparisons in this document use a DNI-to-GHI ratio for Kansas City Missouri (roughly halfway between that of Northern Vermont and that of central Arizona); since utility-scale solar installations tend to be built in areas with high sun quality, this is pessimistic. Cost and efficiency numbers for comparisons are scaled to match the module being compared (e.g., NREL's base case of silicon mounted at a fixed tilt for 6 ¢/kWh, or NREL's target of flat panels on a single-axis tracker reaching 3 ¢/kWh).

2.3 Business Overview

2.3.1 Value Advantage

Not only is TLC's introductory cost below the DOE's expectation for silicon panels in 2020 (as will be detailed in Chapter 5), but the much higher efficiency of TLC also raises the ***value*** dramatically. TLC should also win on durability: CSP's glass and metal receivers at the focus of solar-glass trough mirrors is the *only* solar technology that has over 25 years in the field at utility scale (NREL, 2009), and tandem cells have very low degradation and are well cooled and hermetically sealed behind a glass lens, making 30 years a safe assumption and 50 years realistic.

A US Department of Energy chart (DOE, 2018) shows how much a panel can cost, versus efficiency and durability, and produce electricity at 3¢/kWh *unsubsidized*, in a moderate-sun place like Kansas City, Missouri. This allows comparing the *value* of panels. The chart is copied below and is overlaid with data on the DOE's target for silicon panels in 2030 (mounted on single-axis trackers), along with the cost and efficiency estimates for TLC on dual-axis trackers.

In Kansas City TLC gets a significantly lower capacity factor because not using diffuse light more than offsets the better performance in hot weather and at lower sun angles. The 15.9% capacity factor reduces the 38.5% module efficeincy to the equivalent of 29%, and raises the equivalent cost

from the expected cost of 15 ¢/W (at 1 GW) to just over 21¢/W. While this cost and efficiency provide sufficient value to cover the DOE's expected cost of two axis tracking with a small positive margin in Kansas City, in high DNI regions like Las Vegas the margin is large enough to fund rapid TLC production growth. And this comparison is based on TLC, built with cells available today, at a mere 1 GW per year production volume, while silicon is given the advantage of the Department of Energy's expected progress through 2030, at production volume 500 times larger.

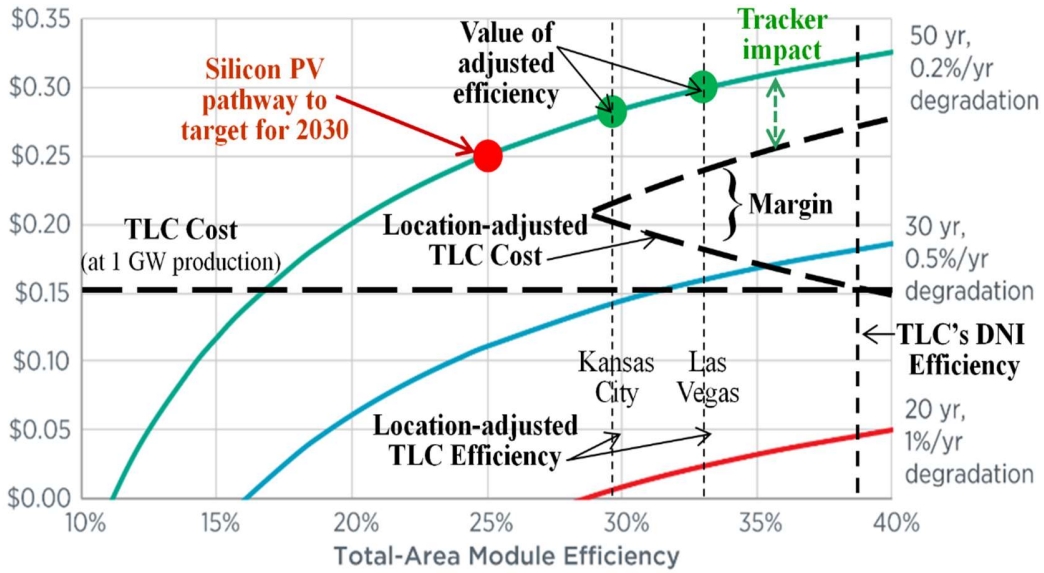

Figure 2.2: *The value ($/W) of efficiency and durability (overlay on graph from* (DOE, 2018).

The cheapest panels are currently <30¢/W, and as Paula Mints pointed out in 2017, "the lowest price in the market is believed to be the average, and any data that does not support this view is often ignored" (Mints, 2017). According to Mints, the average price was twice the lowest price, and those striving for the lowest price "make compromises on inputs, choosing the lowest price backsheet, EVA, junction box, etc."

The premium for high-quality, high efficiency panels can be considerable, with the highest quality panels selling for ~40% more than 'premium panels' and ~75% more than standard tier-1 panels (Brakels, 2016). Mints shows the highest priced panels at the start of 2017 selling for five times more than the lowest priced. And even after de-rating for not using diffuse light in a mediocre-sun area like Kansas City, TLC has a bigger efficiency advantage over SunPower's best panels than SunPower has over typical silicon panels.

Silicon panel makers have for years been willing to accept zero or even negative margins to gain market share, so one cannot count on rational pricing. However the average current (18%-efficient) silicon panels could be *given away for free* and TLC should still have a lower LCOE in Las Vegas, even at an introductory scale, with enough margin to fund growth. Thus although silicon will continue to improve, and continuing drops in the cost of silicon panels will hold TLC profit margins to far below the ludicrous margin implied by today's silicon panel costs, TLC will be able beat silicon in its high-DNI launch markets even when initially made a much smaller scale, while providing ample margins for rapid growth.

2.3.2 Scaling Production

- ***Capital Cost of Scaling:***

With TLC's expected cost advantage in its launch markets, the limit on growth will be scaling manufacturing capacity. The world already has 1 GW/year of cell and lens ovemolding capacity, and more than that of other materials and components. The capital cost of scaling to a gigawatt production is thus a few TLC-specific molds and jigs.

- ***Grid Scale:***

As detailed in Chapter 5, the TLC cost model estimates how much costs will fall if TLC is made at the scale that silicon panels are made today (~100 GW/year). For dual-axis-tracked TLC and using the NREL cell cost estimates (Horowitz, et al., 2015), the TLC cost (including the mirrors) should only be 11 ¢/W, falling to 9¢/W with NREL's substrate reuse costs. This is probably pessimistic because it does not include a "Swanson's Law" scaling of costs with volume on the rest of the system, but there is not a lot of cost left to reduce.

Costs for TLC for single-axis trackers should come down to ~17¢/W just from scale, and to ~15¢/W with substrate reuse. While tracker costs are in addition to this, silicon panels are already put on singles-axis trackers at utility scale, and even with the extra costs of making a flat-panel tracker stiff enough, accurate enough, and full tracking range, the higher efficiency of TLC reduces tracker cost by reducing module area needed. TLC module cost for single-axis tracking will also be lower if the heat byproduct is useful because gathering the heat byproduct costs less than the current cooling with heat rejection.

It is likely that two-axis trackers will still cost significantly more than single-axis trackers, making the total installed cost higher than for TLC on single-axis trackers, but based on the DOE's cost targets for heliostats for CSP, which have more-demanding requirements than TLC needs, the extra cost of two-axis trackers should fall fast enough that overall margins with dual-axis tracking should remain higher than for single-axis trackers.

2.4 Prior Art

In 1975 Miller and Stephens of NASA used an un-tracked trough primary with secondaries that move in two dimensions to achieve moderate concentration (Miller and Stephens, 1979), but did not gain the increased energy capture of primary-axis tracking, and the optical efficiency and concentration were not high. In 1977 Wells added one-axis rotation of primary mirrors to gain tracking benefit (Wells, 1981), but the optical efficiency was not high and the concentration was not sufficient to make today's multi-junction cells cost-effective. In the 1990s Yeomans achieved higher concentration with a secondary that confined light on the first axis while concentrating it on the second (Yeomans, 1994), but on-axis secondaries and numerous external reflections further limited the optical efficiency. While the early work had used linear-Fresnel primary mirrors, Airlight recently designed a fairly high concentration system (~600x) with a single-axis-tracked trough primary mirror (Pedretti-Rodi, et al., 2015). But Airlight's primary concentrator was exceedingly massive, the deep reflective secondary concentrators limited concentration to below that of optimal for today's CPV cells, and the optical efficiency was only ~75%. However theoretical work from Airlight (Cooper, 2014) confirms the high-concentration potential.

Panasonic and Penn State both have two-axis-micro-tracking projects that eliminate module tracking by doing all optical alignment with shifts within the module (ARPA-E, 2015), and Insolite is trying to commercialize such a module. But this requires complex patterns and hermetically sealed modules the size of the entire light-collecting area, and vast numbers of tiny CPV cells spread over a large area. Also, while such modules would be suitable for their rooftop target, for utility scale single-axis tracking lowers LCOE even for module efficiencies as low as 15% (First Solar), so forgoing single-axis tracking is counterproductive at utility scale.

Arizona State University is also doing work on two-axis concentration using a trough for first-axis concentration and low-cost (roll-formed) lenses for second-axis concentration (Wheelwright, et al.,

2014). However, that work uses a moderate-rim-angle trough (twice TLC's rim angle) and as a result even lenses curved on a second axis can barely reach 1000x concentration. SunOyster has a similar architecture but with a lower rim angle than Wheelwright (although still higher than TLC's), and can thus use a simpler lens while reaching 1000X (Corino, 2016), but uses only imaging optics and lacks re-concentration on both axes, the optics fit traditional cells on a traditional single-focal-spot receiver with a classic DBC/AlN substrate and classic wired interconnections, and uses a Pyrex-tube cover-glass which adds cost and lowers efficiency.

TLC overcomes these limitations of the prior art. Field-proven solar mirrors reduce the module area, allowing high-efficiency glass lenses that can be roll-formed at low cost in a full-module cover-glass. The lens further reduces the area, allowing non-imaging optics to be molded on the back of the lens at low cost; these provide higher concentration with adequate acceptance angles. The receivers use tandem microcells for high efficiency and improved cooling, but with a two-dimensional array spaced widely enough for excellent heat spreading yet closely enough to pack on the order of 100 cells onto a substrate that fits in high-speed pick-and-place equipment for low placement cost in spite of the small cells.

Instead of a substrate with a single layer of expensive highly-thermally conductive electrical isolation with a complex electrically conductive pattern on its surface, TLC uses an electrically conductive heat spreader with two electrically isolating layers; one layer is low voltage and not in the thermal path, and hence is very low cost; and the main isolation is after the heat has been spread and so can be less thermally conductive and thus also low in cost. The receivers connect themselves as they are interconnected along a single lens tile per module, the module matches a single contiguous mirror for easy alignment, and a module is size to be easy to handle while minimizing connection costs, and the module and receivers are optimized to maintain a photo-current match even while capturing most of the light in the spill zone at the module ends.

TLC is also adaptable to 1-axis trackers with internal micro-tracking; numerous additional enhancements for this include and off-set micro-tracking axle that keeps the focal width nearly constant at widely varying receiver tilts, multi-lens tiles that can micro-track without collision even when closely packed, pinch-flowing-film heat absorption, and massively redundant heat rejection without excessive replication of components.

2.5 Competitive Architectures

2.5.1 Standard Silicon and thin-film Flat Panels

Silicon PV is currently the market leader, and the safe bet ('never bet against silicon') until some disruptive technology comes along (which appears to be TLC). The DOE's 2016 Funding Opportunity Announcement (FOA) (DOE, 2016) had a good cost breakdown of a silicon module. The FOA also asked for a comparative analysis for new technologies to something known, so one was done for TLC at using the 1 GW cost estimates from this document.

The FOA uses Kansas City Missouri as a reference, so NREL SAM models were built for 10 MW installations with both silicon modules (at the reference efficiency of 16%) on low-cost single-axis trackers and TLC on CPV-grade single axis trackers from the same tracker company. This was not straight forward; for example, SAM couldn't handle HCPV (no diffuse light) on single-axis trackers, so a dense North/South packing (to not have to adjust for different Watts/Hectare) of silicon on two-axis trackers was compared to the same packing for silicon on single-axis trackers to get a 6% derating for single-axis compared to 2-axis, and this was put into the TLC SAM model as a 6% 'optical error'.

The SAM models then indicated that in Kansas City 24.5% more TLC was needed to account for not using diffuse light, but that a lower temperature coefficient gave TLC a 4% boost in output. These values were then used to adjust the cost and efficiency of TLC so that comparisons were on a same kWh/year basis. But TLC has a higher peak-to-mean output from producing less on cloudy days, which required a 2 ¢/W adjustment for inverter loading, and TLC's higher efficiency saved 3.8 ¢/W at current tracker process. This produces what is believed to be a fair comparison, the results of which are as follows (Norman, 2017):

Silicon (FOA @ 16%, 250W)	¢/W	**TLC** at **1 GW/year** w/today's best cells	¢/W	Adjust. ¢/W	Delta ¢/W
Poly -> Wafer -> Si PV Cell	28.5	Cell + trough mirror + lens	16.64	19.92	-8.58
String, tab, interconnection	2.46	Cell place, bond, diode, receiver assembly, wire	1.44	1.72	-0.74

EVA, backsheet	4.94	Spreader, pipes, seals, pump, panel, fin tube	4.18	5.00	+0.06
Glass, frame, sealant, j-box	7.44	Glass, shell, sealant, support, j-box, µ-track	0.97	1.16	-6.28
Total of items in FOA for Si	43.3	Total TLC & Total for same kWh as Si	23.2	27.8	-15.5
1-axis tracker, **same supplier**	14.55	1-axis **CPV** tracker, & adj +2¢/W for ILR	8.98	12.75	-1.8

Table 2.1: *Cost Comparison between Silicon and TLC.*

Thus, even single-axis tracked (SAT) TLC cut module cost *for the same* energy *output* **by 15.5¢/W**, **a 35% drop,** in mediocre-sun Kansas City. TLC's output equals a tracked ~30.5%-efficient flat panel, and TLC has durability for >50 years, so **TLC addresses all three SIPS goals**. (TLC also cut tracker cost 3.8¢/W but adds 2¢/W_{ac} for inverter loading ratio, saving 1.8¢/W more). Dual-axis-tracked (DAT) TLC beats SAT TLC, so the win over silicon is even greater.

TLC will benefit from the same tracker, inverter, glass, automation, and installation cost reductions as silicon PV, and tandem cells have more potential for progress on both efficiency and cost reduction than silicon does, so TLC should maintain or even increase this advantage. So even at 1/100 of the scale that silicon is made at, TLC should easily beat silicon.

2.5.2 Standard Fresnel HCPV

NREL analyzed Fresnel sheet HCPV costs (Horowitz, et al., 2015), so TLC costs (for DAT) are compared to this study, *downgraded to use the same cell* and with other TLC costs based on a 2nd **10 MW** run (with most pricing from **first-world manufacturers).** As detailed in the table below (**Figure 2-4**), TLC cuts module cost from 66.1 ¢/W to 24.8¢/W, a reduction of **41.3¢/W, or 62.5%!**

NREL's HCPV Costing	¢/W	**TLC** cost for **2nd 10 MW** with same cells	¢/W	Delta
Cell	18.0	Cell (higher concentration and optical eff.)	10.0	-8.0
Place/Diode/connector/AlN 7.5W	9.8	AlSiC, place, bond, diode, wire, lens 33W	3.4	-6.4
Mold, coat, place SOE lens	3.5	Roll, over-mold lens, coat,	0.8	-2.7

Mold Fresnel lens Parquet	16.8	Trough Mirrors	5.3	-11.5
Aluminum passive cooling	6.4	Heat pipe, methanol, fin tube	3.5	-2.8
Assemble 375W module	11.6	Supports, back glass, epoxies, connector, asm	1.9	-9.7
Total for HCPV module	66.1	**Total TLC module** (inc. trough mirrors)	24.8	-41.3

Table 2.2: *Cost Comparison between Fresnel Lens HCPV and TLC.*

So even at 10 MW introduction, TLC can beat Fresnel panel HCPV. This can be seen graphically:

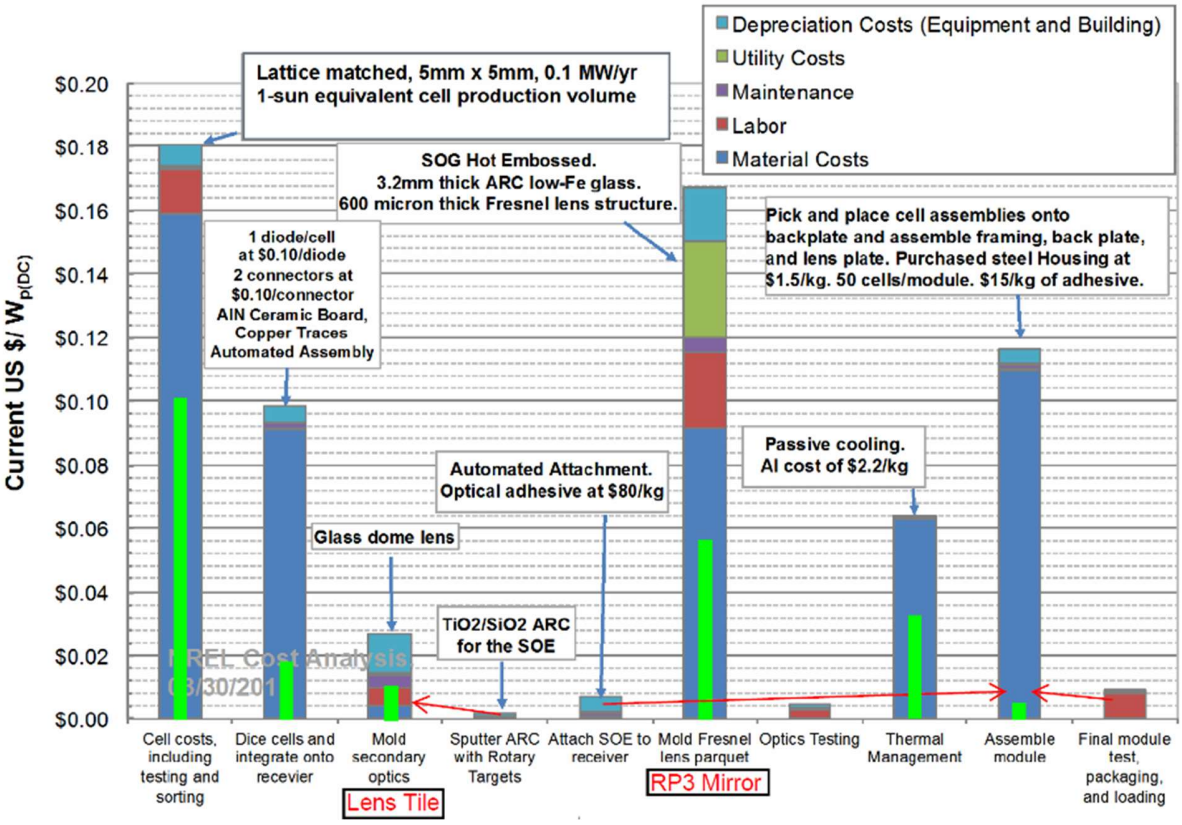

Figure 2.3: *Step-by-step CPV module manufacturing costs by NREL* **(Horowitz, et al., 2015)**, **with green bars for TLC costs added to the NREL chart**

2.5.3 Discrete Cassegrain Mirror Flat-Panel HCPV

Numerous companies (such as SolFocus for their Gen-1, and recently CGE) have used small mirrors instead of small Fresnel panels. No cost breakdown has been compared, but the costs are thought to be similar to standard Fresnel HCPV, and therefore **can easily be beaten by TLC**.

2.5.4 Other Single-Axis-Tracked High-Concentration Trough Work

Airlight recently designed a fairly high concentration system (~600x) with a tracked trough single-axis primary mirror, and mini-tracked reflective secondary concentrators (Pedretti-Rodi, et al., 2015). But Airlight's primary concentrator was exceedingly massive, the deep reflective secondary concentrators limited concentration to below that of optimal for today's CPV cells, and the optical efficiency was only ~75%. However, a PhD thesis (Cooper, 2014) by someone working with Airlight did a huge amount of theoretical work on second-axis re-concentration, so they might have an improved version in the works and are a company to watch for. In 2009 Norman taught the basics of using refractive micro-tracked secondary concentrators in a U.S. patent application (Norman, et al., 2010), and TLC is a continuation of that work with considerable improvements.

2.5.5 Two-axis-tracked Troughs with Secondary Concentrators

- ***SunOyster:***

The SunOyster is a high-concentration trough that uses a static lens to achieve over 1000x (Corino, 2016). The SunOyster is an elegant design and appear to be well made, so it should be taken seriously as a competitor. However, the SunOyster uses a DBC/AlN card 6x larger than the photoreceptive area, versus no DBC/AlN at all for TLC; the lenses will cost more due to not being roll-formable, and without the TIR cones they will not reach as high a concentration as TLC (in spite of the two-axis tracking).

TLC should be able to beat the SunOyster in cost by a sizeable margin, although SunOyster's focus on roof-tops and heat cogeneration means that there might not be direct competition for a long time since this is near the opposite end of the market.

- ***Wheelwright:***

Another high-concentration trough system that uses two-axis tracking is from Dr. Roger Angel's group in Arizona (Wheelwright, et al., 2014). This design reaches 1000 with a roll-formed refractive secondary concentrator, using a dual-RP3-inner mirror. However, using the dual inner mirror produces twice the rim angle, and thus roughly four times the Petzval effect, that TLC gets with a single such mirror-width focusing onto a row of receivers. To reach 1000x with such a high

Petzval factor, Wheelwright has to use a large cylindrical lens sheet instead of a small flat lens sheet, and even then, the design has too little acceptance angle at 1000x to be practical.

While Wheelwright's design as it stands is thus not practical, it does show that roll-formed lenses can reach 1000x at a trough's focus. And since it is in Dr. Angel's group, someone there might put those parts together and become a serious competitor.

- ***TLC on a Two-Axis Tracker:***

As two-axis tracking comes down in cost, SunOyster and Wheelright are becoming more competitive. But TLC can also benefit from 2-axis tracking, and TLC's low-cost lens sheet, low-cost substrate, and low-cost cooling give TLC a significant advantage over these other trough + secondary concentrator architectures.

2.5.6 Primary Optics Sheet with Small Cells

- ***Semprius:***

The Semprius HCPV modules (Furman, et al., 2010) are in some ways similar to a TLC module, but Semprius does not pre-concentrate so their modules are full-aperture-area. A Semprius module uses glass lenses, formed as large sheets, to focus onto tiny cells (half the size of TLC cells), with extra concentration and alignment of the light right before the cell. If anti-reflection coatings progress to where TLC uses a two-part lens, there will be even more optics similarity. Semprius also initially transfers and bonds the cells onto a small area (but then places the assemblies again on a larger area). However, comparing to TLC show numerous major advantages for TLC:

- TLC's glass lens area is ~25 times smaller due to the initial trough concentration
- TLC has 2x fewer cells to place and to wire-bond
- TLC uses only 2 wire bonds per cell – Semprius looks like 4 (Burroughs, 2010)
- Even SAT TLC has 80x fewer glass parts to place (lens + manifold vs 160 glass balls)
- TLC has 160x fewer cell+lenses+substrate assemblies to place
- TLC's assembly placement is on a ~40X-smaller area
- TLC's CTE / optics alignment is 1D and easier to match
- TLC's cooling has better performance
- TLC's sealed module area is ~12x smaller (and ~6x smaller volume)

- TLC protects the cells much better for long life
- TLC doesn't depend on novel assembly techniques (but could use them if they work)
- TLC doesn't depend on singulation after placement (but could use it if it works)
- TLC gets higher concentration from non-imaging final concentration
- TLC has a low-bus-bar-area cell design that reduces cell area

These far more than offset the few classic Semprius advantages:

- Semprius doesn't need extra heat rejection hardware
- Semprius appears to have mastered transfer printing
- Semprius has low-cost cells from substrate reuse
    - But to get this Semprius needs transfer printing and its own cell fab

However, Semprius has at least one more card that they might play. Semprius received almost $3M from the DOE to test transfer-printing onto a PV substrate (probably a silicon wafer) that would use the light that did not hit the cell – this would let them use diffuse light at silicon efficiency. However, a silicon cell (including placement and wiring) is more than half a PV-module's cost and diffuse light is typically less than ¼ of DNI. In Kansas City this would add roughly 12% to the module output but at a cost of about $1.10/W for the extra watts. Thus, although this trick might win business where the highest output per module area is critical, it will not be cost-effective for the low-cost output that is key at utility scale. So, TLC can easily beat other lens-sheet sheet HCPV architectures, including Semprius.

- ***PARC/SolFocus:***

Another optics-sheet module was the SolFocus Gen-2, a 'solid-state' version of their Gen-1 Cassegrain design. This is somewhat similar to both Semprius optics and the TLC lens, but it uses small mirrors instead of small lenses. This design was originally from PARC (Elrod, 2008), but never quite made it to commercialization. However, SolFocus touted it as $0.35/W, and PARC used a $4/cm^2 cell cost leaving little for the optics and setting an upper bound on cost. Its similarity to TLC comes in the one-piece multi-cell molded glass optics, but it used silvered mirrors, and had a longer path through the glass, and could only afford half-decent glass (further lowering optical efficiency) since the whole 40X-larger aperture area needed to be the molded/mirrored glass. Therefore, this **can easily be beaten by TLC**.

- ***Light-guide Optics:***

Another optics-sheet module is the new Morgan Solar integrated modules. Morgan is a survivor and is very good at finding new tricks. Their new modules are molded lens sheets on each side of glass sheets that encapsulate arrays of small cells. However, they still have twice the glass area of TLC, with two full-aperture molded sheets (50x the molded optics area) and have worse thermals. And while their cell size is comparable (0.8 mm on a side, 1.2 times the TLC cell area), their cells and wire-bonding are on a full-area (~40x bigger) substrate which requires special equipment. Therefore, this design **can easily be beaten by TLC**.

2.5.7 Two-Axis-Micro-Tracked System

There is also some work on flat-panel CPV systems that use internal micro-tracking on two axes. Some of this work is even DOE-funded, largely through the ARPA-E MOSAIC program (ARPA-E, 2015). Not only did Semprius receive several million dollars under this program, but two internal micro-tracking projects were similarly funded: Panasonic Boston Laboratory, for "*Low Profile CPV* Panel *with Sun Tracking for Rooftop Installation"*, and The Pennsylvania State University, for "*Wide-Angle Planar Microtracking Microcell CPV"*. There is also a commercial startup introducing a similar panel (Carron, 2016).

No information has been found on the concentration achieved or the cost of the optics, but even if the two-axis internal micro-tracking CPV projects are successful, they will still suffer the high cost of full-aperture-area two-axis optics. And while internally tracking on both axes is fine for roof-tops, at utility-scale any even a moderately-efficient panel would be thrown on a single-axis tracker anyway to increase its capacity factor, *so a second micro-tracked axis is largely superfluous at utility scale*. It would allow a slightly lower-cost non-CPV-grade single-axis tracker, but that difference is small enough that it is less than the difference in cell cost between 1000x and 1500x.

Such systems are therefore possible, but unlikely, competitors at utility scale. However, they might be allies in pushing small tandem cells forward.

2.5.8 Big Dish CPV

Big Dish CPV uses a parabolic dish, or a reasonable facsimile thereof, to concentrate onto a large-area receiver. Many companies have used this approach, from Solar Systems in Australia and Southwest Solar and REhnu in Arizona to SpaceWatts here in Quebec and Zenith Solar in Israel.

Most big-dish solar uses arrays of two-axis mirrors, and thus have the cost of two-axis molding of glass for the entire aperture area, the cost of aligning numerous mirrors precisely on two axes, and the cost of two-axis tracking of large structures. Unless these costs all drop dramatically, this cannot compete on cost with TLC with current tandem cells. However, this architecture can achieve extremely high concentration, so if, say, 5000x cells were to become available at 60% efficiency, this architecture could be a competitor. Not all big-dish HCPV has all of these handicaps. While all known big-dish CPV uses two-axis tracking, variants overcome the other cost issues:

- ***Zenith:***

Zenith Solar was a company in Israel that replaced the two-axis mirrors with an array of roughly 1000 small flat mirrors held in a molded plastic dish for alignment. However, to produce sufficient concentration for CPV cells, the mirrors were of over a hundred different easily confusable shapes. While they did deliver some systems, and used a clever trick for being able to current-match cells in an uneven focus using only two sizes of cells (Löckenhoff, et al., 2010), the design had low optical efficiency and could only compete when heat was desired, and the sea of small, confusable mirror would have been hard to produce at reasonable cost.

- ***REhnu (www.rehnu.com):***

REhnu is a solar startup in Arizona that is attacking the cost of the two-axis mirrors and their alignment. REhnu has the advantage of Dr. Roger Angel, who is probably the world's top expert in accurately shaping glass (having pioneered the techniques used to mold the world's largest telescope mirrors). And surprisingly for one used to creating near-perfect instruments, Dr. Angel *also* understands how to wring out cost from molding processes. For example, Dr. Angel's team developed the process that was a key driver in bringing parabolic trough mirror cost down by half in 5 years, while also increasing their accuracy and the range of focal lengths producible.

REhnu uses advances in mirror shaping to produce larger two-axis mirrors with shorter focal lengths than traditional big-dishes use. These mirrors are large enough that a single host a receiver

can produce enough power to be worth the overhead of active cooling. A single mirror minimizes steel structure (2-axis molded glass make a fine structural element) and alignment issues.

REhnu expects the mirrors to be ~\$40/m^2 in high volume, and 2-axis CPV trackers currently cost \$80/m^2, so that's \$120/m^2 or ~30 ¢/W. Even given the higher capacity factor from 2-axis tracking, that's equivalent to 25¢/W, which is almost 10 ¢/W higher than the TLC tracking and optics (and that's TLC at low volume – that's higher than the entire TLC cost will be at very high volume).

While two-axis mirrors still cost more than single-axis mirrors, so REhnu must be seen as a serious potential competitor, especially if future cells can handle dramatically higher concentration.

- ***SpaceWatts SW35KW:***

A third project on reducing big-dish cost is the SpaceWatts SW35kW, prototyped right here in Sherbrooke. The SpaceWatts dish attacks the optics, and the optics alignment costs, and starts attacking the receiver cost, and work is in progress on the two-axis tracking cost as well. The dish mirrors are compound single-axis mirrors and are not even molded at all.

An improve mirror shaping technique that is higher accuracy and even lower cost is ready for testing (although Dr. Angel's work on troughs means that the savings would not be as large). Mirror alignment is also simplified with 'rib modules' where a pair of ribs can be aligned with a machinist's level to then accurately position a whole bank of mirrors, and this, too, has advancements ready for testing.

The two-axis tracking can also be simplified. The current dish tracking is based on a hybrid carousel design that was expected, based on older work (Sargent and Lundy., 2003) to be lowest cost. However, the dual-dish design also lends itself to a pole-top tracker, and an innovative design that uses rigid stays to get the effect of three-point support with a single central point will allow a tripod to provide the stability leverage of a carousel without the complexity. This will reduce the cost penalty of two axis tracking, although it will still cost more than single axis.

While the win is not as dramatic as against other competitors, TLC is none-the-less expected to beat even cost-reduced big-dish CPV. But big dish CPV appears to produce an ideal high-concentration size for solar thermochemical (sunlight-to-fuel) receivers.

2.5.9 Power Tower CPV

Power towers, in which numerous heliostats focus on a central HCPV receiver on a tall tower, are another form of CPV that has been worked on. However, while this is fairly promising, it works best on a very large scale and so the number of companies working on it is much smaller.

Solar Systems (later bought by Silex) did some work in this field, and its founder is now with a new company, RayGen, pursuing power towers for CPV. Heliostats are expected to come down into the $75W/m^2 range (DOE projections), but cosine effects bring efficiency down, raising the effective primary optics/tracking cost (and TLC looks lower cost on receivers and cooling).

SpaceWatts has three main improvements that can be used independently or together:

Multi-tower fields that reduce cosine losses and increase concentration by orienting mirrors to direct light to the most effective tower at any given time, rather than always to the nearest tower; this increases peak-to-mean power and should cut overall field costs by ~15% at most latitudes.

Focusing heliostats that cost-effectively allow many fewer heliostats and much shorter towers to reach the same concentration.

Multiple receiver types with that use the same heliostats, which allows using PV receivers for peaking power when demand is high but storing heat for future use dispatchable power when sunlight exceeds immediate power demand.

2.5.10 Emerging technologies

In a fair comparison TLC appears to be able to beat all current competitors, but even if TLC meets performance and cost expectations it is far from a guarantee of invulnerability. Just as TLC is a disruptive technology that could dethrone silicon and thin-film flat panels, some emerging technologies could turn out to be even more cost-effective and dethrone TLC

Even at silicon's maximum realistic efficiency, it would have to be almost free to compete with TLC, and even then, TLC has lower use of glass, steel, and land per Watt and per kWh.

- ***Luminescent Concentrators:***

It possible to greatly concentrate even diffuse light if one absorbed it and then emits it directionally at a longer wavelength – this is how an optically-pumped laser works! However, in spite of

considerable effort, even monochromatic versions of this, which would have efficiencies similar to silicon PV, are stuck at less than 50x concentration and low optical efficiency.

Therefore, this potential game-changing technology is at least two breakthroughs short of being practical even for building-integrated PV, and three breakthroughs from competing with TLC.

- ***Tandem Flat panels:***

Multi-junction flat panels could reach 30% or higher efficiency and become competition to TLC. It would be hard for large-area tandem cells to beat TLC on overall output because the efficiency boost from concentration is greater than that from using diffuse light, except in areas with quite poor sun quality. And TLC's non-cell/non-mirror costs would still be roughly equal per watt to a tandem flat panel, so such panels would have to beat TLC on cell cost. Since TLC already beats silicon dramatically on cell cost, it thus unlikely that a tandem cell more efficient than silicon could achieve a much-lower cost per area than silicon. Thus, such panels are unlikely to challenge TLC at utility-scale, although they could be fine for rooftops.

- ***Rectifying Antenna Arrays:***

A radio antenna can convert radio waves to useable electrical energy far more efficiently that a photovoltaic cell could convert such individually weak photons. Over 30 years ago technology advanced to where this was possible for microwaves, and it was suggested as a way of beaming power from solar arrays in space down for terrestrial use (Brown, 1984). However shorter wavelengths require proportionately shorter antennae and faster diodes, and light has a wavelength roughly 100,000 times shorter. In 2007 the wavelength range was extended into the infrared (Gritz et al., 2009), closing the gap toward visible light.

Theoretically an antenna for optical wavelengths could feed an extremely fast rectifier that would convert optical photon energy to useable electrical energy with higher efficiency than a PV solar cell, and an extremely dense array of such antennae could deliver very high conversion efficiency, so research toward this has proceeded. In 2015 Georgia Tech reported a first rectenna array at optical wavelengths using carbon nanotubes for the antennae/diodes (Sharma, et al., 2015), but while their devices functions, their conversion efficiency was still below 1%.

A company called NovaSolix has recently claimed to have breakthrough technology that makes rectifying antenna arrays from carbon nanotubes practical (NovaSolix web site). NovaSolix's

claims are extremely aggressive – 45% efficiency at introduction, with 80% to 90% possible, and a cost 10 times lower than silicon flat panels.

45% efficiency would roughly match TLC efficiency, and if the cells can be made at thin-film prices and simplicity, they could conceivably come in at lower cost than TLC could even at grid-scale production. Therefore, NovaSolix is a company to watch, and outside of NovaSolix, rectennae are a possible emerging technology competitor to watch.

- ***Hot Carriers:***

Hot carriers also offer a path to much higher efficiency. By capturing the extra above-band-gap energy before it is thermalized, hot-carrier cells could theoretically reach~65% efficiency and might possibly be low cost since perovskites have long hot-carrier lifetimes (Guo, et al., 2017) and perovskites are a hot research area for low-cost flat panels. Theoretically hot carriers could therefore deliver the best of all worlds – high efficiency, low cost, and use of diffuse light, so this is an area to watch.

2.6 Advantages and Issues Comparison Table

PV Type	Sili-con	Thin Film	Fres. ~5x5	Semp rius	PARC / SolF	2A µ track	**TLC DAT**	**TLC SAT**	Other DAT Trough	Dish DRA	Power Tower
Raw Effic.	Med	Med Low	High	Very High	Med High	Med High	**High**	**High**	High	Very High	High
Diffuse Use	good	Good	None	Low	None	None	**None**	**None**	None	None	None
Temp Coeff	Med High	Med	Low	Low	Low	Low	**Low**	**Low**	Low	Low	Low
Durable	Med	Med	Med	Med	High	Unk*	**Very high**	**High**	Very High	Med High	Med High
Tracker Cost	Med Low	Med	High	High	High	Very Low	**High**	**Low**	High	High	High
Tracker Benefit	High	High	Very High	Very High	Very High	None	**Very High**	**High**	Very High	Very High	High
Cell Cost	High	Med High	Med Low	Low*	Med Low	Unk	**Low**	**Low**	Low	Low	Low
Place Cell	High	Low	Low	Med Low*	Med	Med Low*	**Low**	**Low**	Low	Low	Low
Primary Optic	Very Low	Low	High	High	High	High	**Med Low**	**Med Low**	Med Low	Med Low	Med Low
Align Primary	No Cost	No Cost	Low	Low	No Cost	Med	**Low**	**Low**	Low	Low	Low
Final Optics	No Cost	No Cost	Med	Med High	No cost	No cost	**Very Low**	**Med Low**	Low	Med	No Cost
Final Optics Place	No Cost	No Cost	Med	Med High*	No cost	No cost	**Low**	**Med Low**	Low	Low	No Cost
Isolation Cost	Low	No Cost	High	High	Med	Unk	**Very Low**	**Very Low**	Med	Med	Med
Cooling Cost	Low	No Cost	High	Very Low	Low	Very Low	**Low**	**Low**	Med Low	Med	Med
Cooling Perform	Good	Good	Not good	Med	Med	Med	**Very Good**	**Very Good**	Good	Very Good	Good
Other Receiver Assembly	Med	No Cost	Med High	Med High*	Med	Med	**Low**	**Med Low**	Unk	Low	Low
Module Frame, Assembly	Med	No Cost	Med	Med	Low	Med	**Very Low**	**Low**	Low	Very Low	Unk

* = Requires experimental technique

Table 2.3: *Advantages and Issues of PV types.*

2.7 Conclusion

This chapter has introduced a novel CPV architecture called TLC, and has discussed how TLC is related to previous CPV architectures. This chapter has shown that if TLC can be built as designed, it will beat previous CPV architectures, and potentially even silicon flat panels, on cost. This establishes that TLC is worth pursuing as a CPV architecture that has the potential to significantly reduce the cost of solar energy and contribute to raising the world's living standard while reducing global warming. This is important because it forms a justification for carrying on the research undertaken as part of this PhD work.

CHAPTER 3

3. TLC CONCEPTION & DESIGN

3.1 Context and Issues

3.2 Introduction

The next two chapters are the heart of the thesis. Chapter 3 starts with a more detailed presentation of the multistage concentration used in TLC, and an overview of its impact on the rest of the module. This chapter then goes through the lens tile details, from the glass front to the silicone cones molded on the back. (The models and validations this is based on are covered in Chapter 6.)

This chapter then briefly discusses the scaling principles of TLC modules to different sizes, and extensively covers other variations of TLC that have been considered but that are not currently thought to be optimal.

3.3 Multi-stage Concentration

3.3.1 Initial Concentration

A key to cost-effectiveness is low-cost optics; parabolic troughs are very low cost, and while lenses are needed, the concentration from the trough makes the lens area roughly 40 times smaller, which dramatically reduces the lens, substrate, and assembly cost.

TLC uses parabolic trough mirrors from the CSP industry for its light-gathering area and initial concentration. A subsequent second axis of concentration allows concentrating less on the first axis, for relaxed tolerance and low costs. All ray tracing has been done assuming that mirror and tracking inaccuracies are comparable to standard commercial troughs (in other words, commercial troughs are already good enough).

- ***Low Rim Angle:***

The second-axis concentration works best with a low rim angle for the trough, and the relaxed first-axis concentration allows a low rim angle. A single RP-3 inner commercial inner-mirror segment provides a suitable rim angle, and such mirrors are widely available at low cost. The low rim angle allows off-axis concentration, which eliminates mirror shading by the module.

- ***Efficiency:***

Mirrors for the CSP industry are typically ~94% to 95% efficient at reflecting the sun's energy onto a receiver-tube. However, CPV cells use a narrower range of wavelengths than CSP and, CSP mirrors actually work better for CPV than for CSP, typically reaching ~95% reflectivity for CPV-relevant wavelengths, and for current cells the key wavelengths are where mirrors are most reflective, typically exceeding 96% (Angel, et al., 2014). Thinner glass improves the reflectivity slightly, and the best 3 mm glass mirrors (PPG's tempered-glass mirrors) exceed 97% in the key wavelength range (from PPG datasheet, 2011).

- ***Cost***

  - Mirror Cost:

The current mirror price (RioGlass quote, 2017) of EUR 17.75/m^2 for 10 MW is roughly \$20/m^2 at current exchange rates, or ~5.2 ¢/W. Since mirror costs decrease over time, this is pessimistic for volume production. Such mirrors are already made in quantities sufficient for a gigawatt, so volume-driven cost reductions are not included for gigawatt-scale production in 2025 either; this is highly pessimistic since normal purchasing negotiations for a gigawatt can surely do better than the first quote for 10 MW that a technical person gets from the first supplier contacted.

  - *Tracker Cost:*

The dual-axis tracker cost allocated is an un-negotiated 10 MW quote from the first supplier to offer pricing. It is thus pessimistic, even for 10 MW, because a purchasing agent should be able to do significantly better than an un-negotiated quote. But even at roughly 20¢/W, it simplifies TLC and reduces other costs enough to be worthwhile. NREL Heliostat cost targets (Mehos, et al., 2016) suggest two-axis trackers will soon cost less than this even when installation costs are added, and that heliostat costs (including installation) should fall to 12¢/W by 2030.

3.3.2 Achieving High Concentration

A trough's concentration is nowhere near sufficient to make today's tandem cells affordable, so TLC adds two additional concentration stages; as shown in **Figure 3.1**, the trough focuses onto the module's lens tile, where the lenses focus on rows of cones, which, in turn, funnel light to the cells.

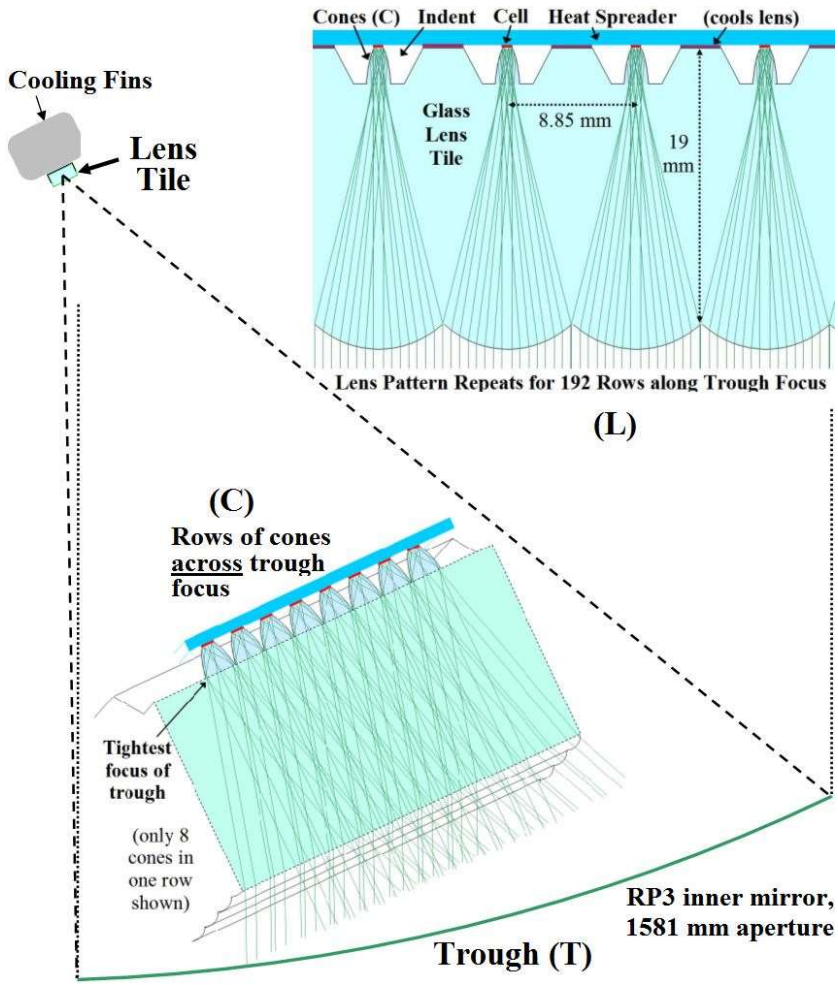

Figure **3.1:** *TLC 3-Stage Focusing.*

- ***Roll-formed Lens Sheet:***

Concentrating ~40X on a first axis greatly reduces the area of optics needed for concentrating on a second axis. The lens requirements are not demanding by lens standards, so the lenses have also been designed to be roll-formed as a sheet for ultra-low cost. The lens sheet as rolled will have the focusing curves on the front face, and channels on the back that displace the same volume of glass,

making this endless linear pattern easy to roll-form accurately (**Figure 3-2)**. The glass used is a standard ultra-low-iron glass thickness of 19 mm (used for windows in tall buildings).

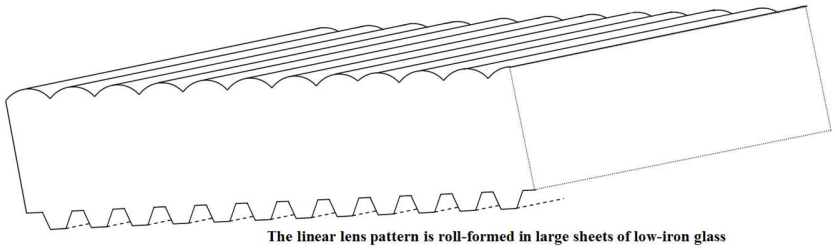

Figure 3.2: *Basic Roll-formed TLC Lens Sheet.*

- ***Over-molded TIR Cones:***

The lens sheet will then be cut into module-sized lens tiles. Silicone TIR cones will then be over-molded in the channels in the back of the lens (**Figure 3-3)**. This process is similar to that used for silicone-on-glass Fresnel lenses for traditional Fresnel/box CPV, but the primary concentration reduces the molding area needed by ~40x and thus greatly reduces costs.

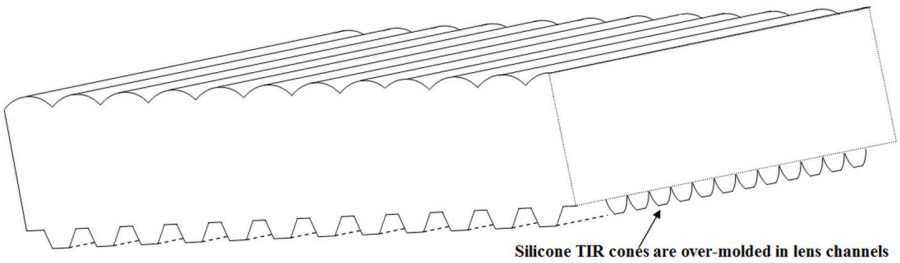

Figure 3.3: *Lens Panel x-section showing TIR Cones.*

- ***Three-stage Focusing:***

The effects of the optical stages on the light can be seen in **Figure 3.4**. A trough mirror produces a long focus a few centimeters wide (T). Linear lenses then focus on a second axis, turning the trough's focus into a series of narrow focal lines a few centimeters long (L). Cones use total internal reflection from compound parabolic curves to further concentrate the light into an array of small rectangular focal spots (C). This array of focal spots matches the cells on the TLC receivers.

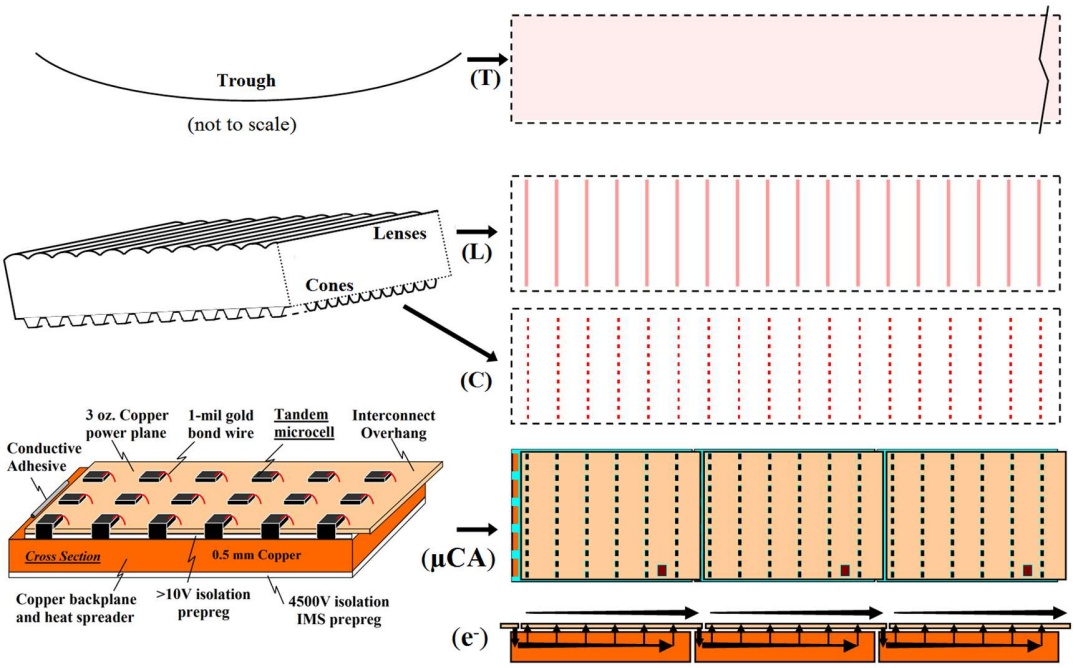

Figure 3.4: *Achieving High Concentration onto Microcells (right half first published in* (Norman, et al., 2019)*)*

- ***Lens Focusing Capability:***

The trough mirrors add ~4 mrads (milliradians) to the sun's ~9.5 mrad diameter and mounting and tracking each add another mrad so ~16 mrads is adequate. Normally a glass lens could concentrate almost 100x from a 16 mrad source, but the light coming in on the first axis spans an angle big enough that Petzval effects become significant. For the optimum lens thickness for affordable glass and not-too-small cells, the Petzval effect and chromatic aberration cut the concentration to ~9x.

- ***TIR section for Refocusing:***

Each lens has a row of compound-parabolic-curve cones at its focus that use total-internal-reflection consume the remaining angular budget from the trough and lens for additional concentration. While the lenses focus to just less than a millimeter of width, the cones are ~1.6 mm wide at the top, providing tolerance during over-molding (**Figure 3.5**). The cones concentrate the light into the cell width (currently 0.65 mm).

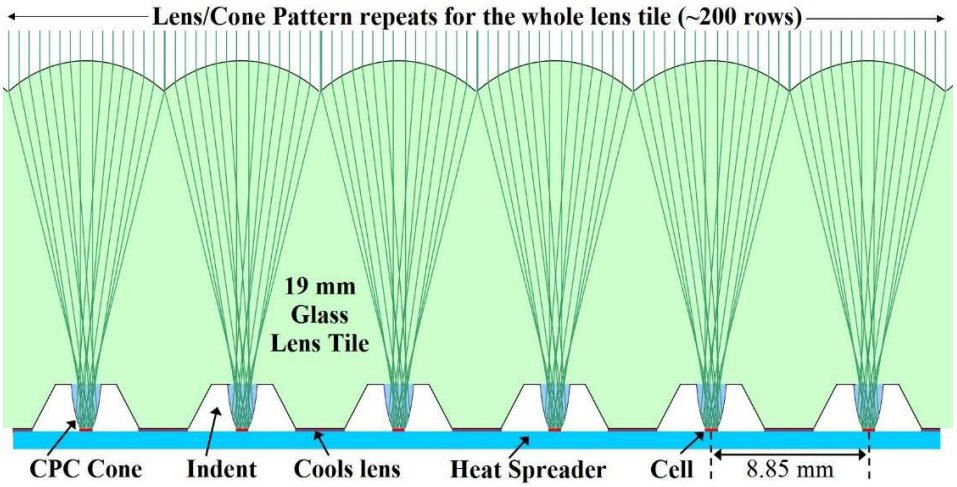

Figure 3.5: Lens *Cross Section ALONG Trough Axis, ~3X Scale.*

The TIR cones are broader in the direction across the trough, where a row of sixteen cones concentrates the light from a ~33-mm-long lens focal line into a dashed line of sixteen foci each matched to a 0.9-mm-long cell. This is illustrated in this slice down the middle of one of the lens channels (**Figure 3.6**), where the light from the width of the trough (and thus at a wide range of angles) is funneled by the row of cones onto the row of cells (12 shown to avoid crowding).

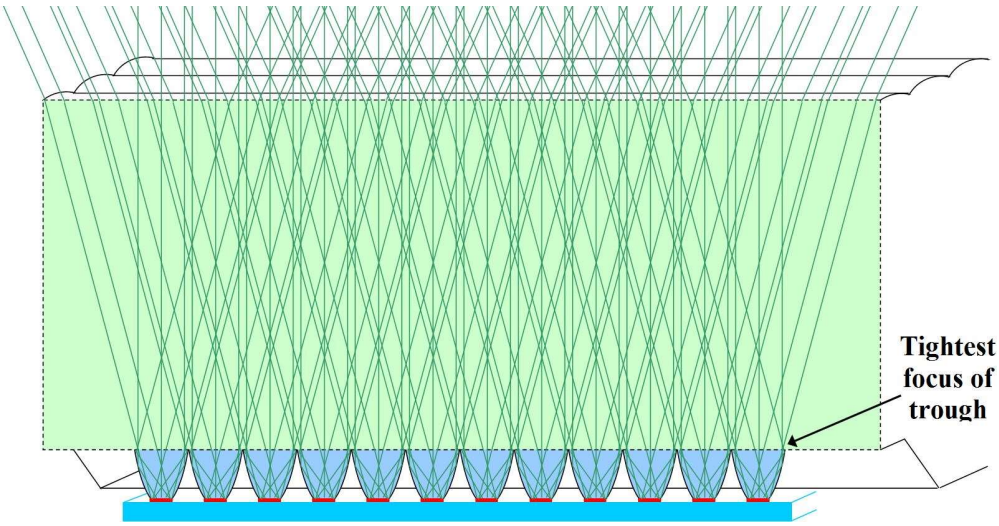

Figure 3.6: *Re-Concentration ACROSS Trough Axis, ~3X Scale.*

The lens intercepts the rays from the trough just before the trough's focus, so the rays are still converging at the surface of the lens and continue converging within the lens. The cones intercept

the light at the trough's maximum focus within the lens. Only the central rays and the extreme rays from the edges of the trough are shown in this sketch so that the focusing can be seen; the intermediate rays fall between these extrema.

- ***Limits to Re-concentration:***

Separating the concentration into stages greatly reduces the cost because the optics at each stage are very low cost, but the trade-off is a Petzval-effect impact on focusing at each re-concentration due to the angular range of the light on the other axis. A mirror with half the rim angle would reduce the Petzval effect ~4x while only reducing the primary concentration ~2x, allowing even higher concentration, but a longer focal length starts adding cost, and higher second axis concentration would make the cells even narrower. Figure 4.19 of Cooper's PhD thesis (Cooper, 2014) shows that even without first-axis re-concentration, a trough with micro-tracked secondaries could reach ~2300x at a latitude of 30º (which would be ~1900x for the 37º the single-axis TLC design is optimized for), at the rim angle of an RP-3 inner mirror, and that on a two-axis tracker even higher concentration can be reached (~4500x), so higher concentration could be achieved even with the chosen RP-3 mirror.

A flatter lens would leave more angular budget for the non-imaging CPC cones, but a flatter lens would mean either thicker glass or narrower cells. Since current cell efficiency starts decreasing significantly after 1500X, higher concentration is counterproductive with today's cells anyway.

Thus, the concentration was targeted to be 1500X (the cell size was set to be 1/1500 of the trough aperture area), and any excess concentrating capability was 'spent' to make manufacturing easier.

3.3.3 Impact of the Optics on the rest of the Module

TLC optics are designed not just to be low cost themselves, but to reduce the rest of the module cost as well. As shown in **Figure 3.7**, the optics match microcells, and TLC's primary concentration packs a ten-of-Watt cell array on a receiver sized to fit high-speed pick-and-place and wire-bond machines. A copper heat spreader serves as a common electrical backplane for a receiver's cells. Receivers are placed along the back of the lens tile (along the trough's focus), with the cells aligned to cone tips; copper strips connect the receivers in series as receivers are placed. A standard 1500V solar connector is used at each module end.

The module lid is currently an aluminum-finned heatsink:

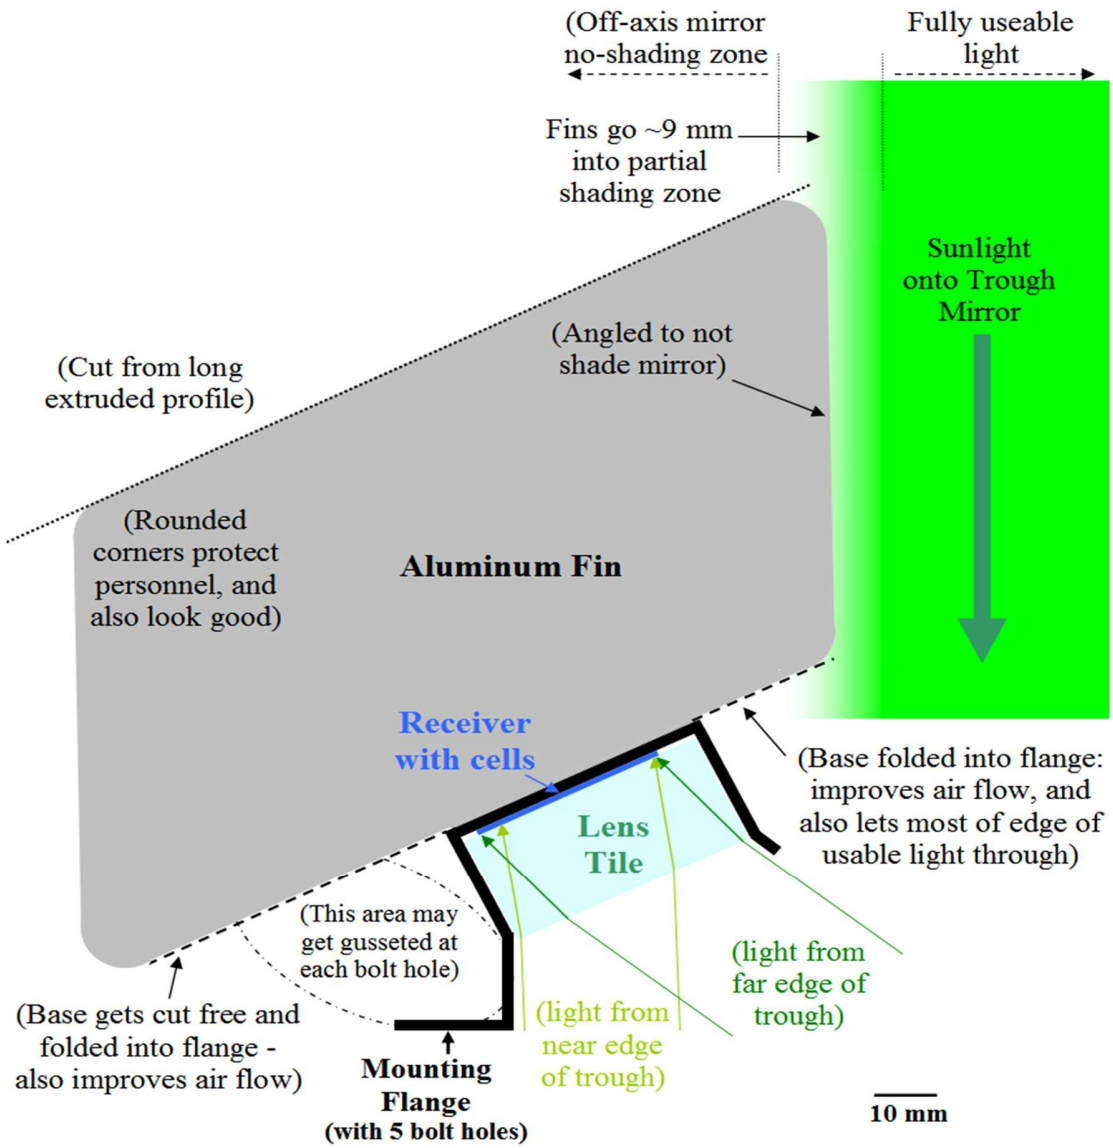

Figure 3.7: Overall module (cross-section across trough's focus).

3.4 Lens Details

3.4.1 Lens Glass Type

A trade-off between lens cost and optical loss strongly influences the type of glass to use. PARC/SolFocus Gen-2 (Elrod, 2008) covered the entire aperture with roll-formed glass, and so could only afford a cheap glass that had a loss of ~4.8%/cm. TLC uses 40x less lens area, so even with thick lenses TLC uses an order of magnitude less glass and can thus afford better glass.

- ***Relation to Cells:***

At a given 2nd-axis concentration, cell width is almost directly proportional to lens width so wider cells require wider, and thus longer-focal-length, lenses. This requires thicker glass, which increases both cost and absorption losses, so the cells are quite narrow to keep the lenses from being too thick. However, the lens tile could be twice as thick and the cells thus four times as big without raising the cost noticeably (Norman, et al., 2019).

- ***Current Cell Spectral Sensitivity:***

The absorption of commodity glass varies significantly with wavelength, and drops significantly above~850 nm, or ~925 nm for a good glass, so the spectral response of the cells is also related to allowable lens thickness. Fortunately, today's most common 3J tandem cells have excess photocurrent in their bottom junction, and that bottom junction starts at about 930 nm so essentially all of the extra loss beyond that wavelength has no impact on the cell's efficiency.

- ***Future Cell Spectral Sensitivity:***

Extra lower-junction photocurrent is likely to continue in future more-efficient cells, and the first such cell analyzed, Boeing Spectrolab's 5J cell (Chiu, et al., 2014) at AM 1.5D has ~6% extra photocurrent in 4th Junction (890 nm to 1130 nm), and 12% extra in 5th junction (1130 nm to 1600nm), so it has some tolerance for iron (although not as much as current Ge-based 3J cells).

- ***Best Commodity Glass:***

The best commodity glass appears to be Vitro (PPG) 'Starphire' for the architectural market. Starphire prices have not been obtained, but a sample has been obtained for prototyping. Similar ultra-low-iron glass is also available from Aohong Glass in China, and it is only $0.56/kg by the

20' container (enough for 5 MW). Since Aohong does not have data on absorption versus spectrum, the spread sheet models the transmission of Pilkington Optiwhite glass. This is probably pessimistic since even Aohong's low-iron glass looks just like Optiwhite (and has overall absorption to match) and Aohong's *ultra*-low-iron looks just like Starphire (seen edge on it is almost completely transparent, with the same pale-blue cast as Starphire).

The following discussion is based on Starphire because the most data is known about it (and it would be affordable even if significantly higher than Aohong's 0.18 ¢/W). Solarphire's 0.8%/cm absorption (from a 2011 datasheet) in visible wavelengths is very low, but its absorption increases significantly starting at 925 mm. From Dr. Angel's work (Angel, et al., 2014), the extra loss averages ~2% from two passes through 4 mm glass, which is 2.5%/cm.

For current Ge-based 3J cells, this means that a lens could have many centimeters of glass before the extra loss impacts performance, and other limits are hit first. In the visible range Solarphire is essentially as clear as lens glass, and so for 3J cells Solarphire is the best glass.

For Boeing Spectrolab's 5J cells (Chiu, et al., 2014) it is not so straightforward. The extra light from 890 mm to 925 mm has slightly more photons per nanometer of wavelength, so this offset ~42 nm of the 925 nm-1130 nm range in addition to its own 35 nm, so at about 35 mm of glass the photocurrent in the 4th junction is reduced to match the first three junctions. Two passes through the mirror takes up 8 mm of that budget, leaving 27 mm. The lens thickness is ~20 mm so this extra 925+ nm solar-glass absorption should leave a slight excess 4th-junction photocurrent. The 5th-junction's light is reduced by a larger percentage, but the junction has twice as much excess and hence it will not be current limiting. So ultra-low-iron glass wins again.

- *Tuning the Cell:*

Future cells may not have such fortuitous extra photocurrent, since TLC has thicker glass than most other CPV. However, it is generally not difficult to tune the junctions to match the spectrum, and a slight shift can provide a lot of tolerance for thicker glass. For example, shifting the 4th junction cut-off from 1130 nm to 1140 would provide an additional 3% excess as a cost of ~0.01V, which is 0.2% relative or 0.1% absolute, a tiny price if extra 4th junction light is needed.

- *Entry-level 'Optical' Glass:*

Would better glass be better? B270 is an entry-level instrument-grade glass that costs $15/kg in modest quantity, which at a density of 2.55 kg/m^{2}mm is $38.25 per m^{2}mm. For a module ~40x

smaller than the primary aperture one gets ~14 kW/m^2, so that is 0.27¢/W per millimeter, so lenses start to be expensive at ~5 mm thick which is much thinner than the thinnest single-tile lens under consideration (but is in the range of thicknesses for air-gap lenses).

Dr. Angel's group got Schott to work out high-volume B270 pricing for solar mirrors (Angel, et al., 2014). Schott gave a \$26.00/m^2 floor for 4 mm thick B270 sheet, corresponding to ~\$124/m^2 for 19-mm sheet, or ~0.86 ¢/W. At 4X the cost of the best low-iron glass, this would add ~0.66 ¢/W. The floor was because B270 glass is not suitable for manufacturing on a float line (which is necessary for the large smooth flat mirrors needed for thermal troughs) because B270 uses antimony to convert highly absorbing Fe^{+2} iron in the glass to much less absorbing Fe^{+3}, and the tin bath of the float process would absorb the antimony from the glass. But roll-forming glass does not use tin so an antimony-boosted low-iron glass should be useable in the future.

- ***<u>Conclusion on Glass for Lens Tiles</u>***

Wide lens rows would be thicker and thus need expensive glass to keep absorption loss low, but thick expensive glass would not be cost-effective. Before the lenses are narrow enough for expensive glass to be affordable, a short path makes premium low-iron glass useable for both 3J and 5J cells. Therefore, the clearest commodity glass wins; in the U.S. such glass is made by at least Vitro (PPG), Guardian, and RioGlass, and these should be evaluated against one another.

The sweet spot for the initial receiver appears to be six rows of cells with 8.85 mm-wide lenses, which allows using the thickest standard architectural low-iron glass (19 mm = ¾" thick). The glass only costs **~0.20 ¢/W**, and the 19-mm average optical path costs only 1.5% of the light through absorption in the key wavelength range (TBC: to be conservative 2.2% modelled).

A traditional solar anti-reflection coating (ARC) on the lens surface reduces reflection to ~2.5% of the light; this is not as good as the air-glass interface for CSP trough receiver-tubes, but the CSP tube's 1.5%-reflection nano-porous glass interface scatters the light over a few degrees, which would limit the ability of the lens to concentrate on the second axis.

Thus, the lens should cost **~4%** of the light, including both absorption and reflection losses (with 4.7% currently modelled due to the lens glass not yet being confirmed as matching Starphire).

3.4.2 Cell Width and Lens Focal Length

Narrower lens rows have shorter optical paths through the lens and thus cost less and have lower losses but require narrower cells. If the cells are not too small, narrow cells help with cooling, but when the cells get too narrow, then kerf losses rise, perimeter recombination losses rise, handling costs rise, and tolerances become too tight.

So, lens optimization starts with the cell and works backward. Due to dense packing on the heat-spreader backplane, cell size doesn't start to raise cost noticeably until after a soft minimum at around 0.5 mm to 1 mm (lower at higher concentrations). Even at 0.5 mm the cost increase from testing, dicing, placing, and bonding the small cells is only ~0.6¢/W and the kerf losses are less than 1¢/W, which are more than made up for by the lower electrical losses and higher cell yield.

The target concentration on a cell is 1500x; below that the cells cost more, and beyond that the efficiency of the cell drops. Even on a single-axis tracker the RP3-inner-mirror trough + lens can concentrate over 100x on the first axis, so roughly 15x is needed on the second axis, which drives the lens toward ~15x the cell width. While the optimum depends on the cell, the dicing and placement methods, bonding, lens glass, latitude, etc., after looking at several cases the optimum appears to be around 0.65 mm wide cells with 8.85 mm wide lenses.

- ***Future Cell Widths:***

Transfer-printing of edge-passivated laser-singulated cells that reach peak efficiency at 2000x could push the optimum down into the 0.4 mm to cell-width range, and two-part lenses with 3x longer focal lengths troughs of lower rim angle and cells with a sharper efficiency peak at below 1000x could push the optimum cell width to over a millimeter.

- ***The Concentration then sets the Minimum Lens Depth:***

The minimum lens focal length to achieve 100x and 15x from an RP3 inner mirror is about 1.6 times the width, with ~2x the width being a bit better (less reflection at the lens edge and less slant of the light reaching the cell. That drives the minimum focal length to very roughly 20 mm. Again, this is a soft minimum one can push the focal length a bit shorter, but losses increase at the lens edges as the angle gets steeper, and the TIR cones (and thus notches) get deeper and sharper and become more challenging to over-mold.

- ***The Glass Sets the Maximum Lens Depth:***

The absorption loss and the cost of the glass both increase with lens depth, which sets a soft maximum at around 25 mm even if other factors permit. Another limit in low volume is 19 mm (3/4") being the thickest standard architectural low-iron glass, so in low volume 19 mm is a practical limit at introductory volume (roll-forming glass hot from the furnace would eliminate that constraint, as would a two-part lens).

3.4.3 Trade-offs in optics optimization

The trade-offs between mirror rim angle, mirror width, lens thickness and cell size were evaluated, and the design presented above is close enough to the optimum to work quite well. However, a slightly lower rim angle would decrease cost, and, with passive cooling, a narrower mirror would be preferred. These trade-offs were presented at CPV-15 and published shortly thereafter (Norman, et al., 2019).

3.4.4 Lens Sheet Manufacturing

Roll-forming is so inexpensive in high volume that it was used to texture glass for ordinary silicon-panel module covers that needed ~100 times more area roll-formed per Watt than TLC (most texturing has now been replaced by AR coatings). The TLC lens is therefore designed to be roll-formable; it can be rolled as large sheers and then cut into lens tiles, a lens tile is expected to have 194 rows (32six-row receivers plus two module ends) and be ~43 mm wide and 19 mm thick.

- ***<u>Design for Easy Roll-forming</u>***

- *Lens Tile Front:*

The TLC lens sheet has features on both faces. The front is linear with smooth features less than 2 mm deep (which is the easiest type of pattern to roll-form with high accuracy).

- *Lens Tile Back – some Features NOT Roll-Formed:*

The final lens-tile back is also a pattern with features only a few millimeters deep, but the final features are not linear and deepest feature ends in a fairly sharp notch and runs at 90 degrees to the features on the front. This would make roll-forming both faces accurately more difficult and would especially reduce accuracy on the back-side features traverse to the sheet motion during rolling, which are the sharp notches that would already be the difficult features to roll-form. Forming the

sharp notch in glass could also serve as a crack-initiation site, weakening the glass. Therefore, the cones are not roll-formed.

- *Rolling and then Over-molding:*

The problematic sharp notch in the glass lens is easily avoided. Only the easy to roll-form features are roll-formed, and the difficult features are added by over-molding afterward. This not only avoids the difficult features but lets the lens back be designed to make the front even easier to roll-form with high accuracy.

Instead of trying to roll in difficult features, a linear indent is rolled in under each lens. This provides a place to later over-mold the sharp traverse features. This:

- avoids sharp features in rolling.
- avoids traverse features during rolling.
- allows the rollers to be tuned for front face accuracy with the back face non-critical.
- allows the cross-sectional area of the channel on the back to be matched to that of the lens on the front so that no bulk lateral material movement is needed during rolling.
- lets the flat area between channels provide non-critical-surface to support the lens sheet during cooling after rolling, without needing precision-contoured rollers for the cool-down conveyor (these flat areas also provide thermal contact for cooling the lens in operation)
- gives flexibility in materials – a cone row could, for example, be molded in an expensive (the volume is tiny) higher-refractive-index glass (for higher concentration, higher acceptance, stiffer cones, etc.) and inserted with optical adhesive.
- allows the area between the channels to be used for alignment when molding the cones.

- *Roll-forming accuracy:*

Forming the lens will leave a small manufacturing radius between lenses. With precision molding this could be ~0.5 microns (based on examples from Adaptive Optics Associates (Borrelli, 1999)); while roll-forming has lower precision, roll forming is a very accurate process when the features are designed for roll-forming, and even if this the radius is an order of magnitude large it will still

only block ~0.1% of the light. This is supported by PARC expecting the roll-formed PARC/SolFocus gen-2 design to be sufficiently accurate even with precise non-linear feature on both faces.

Dr. Angel's group also achieved 1000x (Wheelwright, et al., 2014), without the refractive index enhancement (equivalent to over 2000x for TLC), with a lens sheet that was cylindrical rather than a simple flat sheet (a far more difficult challenge), using hand-built roll-forming equipment. With its design for roll-formability, the TLC lens sheet should achieve very high accuracy.

- *Over-molding Silicone TIR Cones:*

This one-face-good/other-face-prepared sheet will then be over-molded (or hot embossed) with a refraction-matched optical silicone to form the accurate TIR cones that complete the optical surface. While other optical polymers could be used, silicone is easily moldable, very clear, and highly UV resistant (Beukema, et al., 2019). Similar over-molding/embossing is already used in making the primary aperture lenses in traditional Fresnel-lens CPV (Horowitz, et al., 2015), and in TLC 40x less area is needed. The refraction-matched optical material is hermetically sealed between the thick lens glass and the rest of the receiver and hence is also better protected than in traditional Fresnel-lens CPV.

The over-molding could be done on large sheets, but to keep the cones free from particulates it will probably be done after singulating the individual tiles. Although optical silicones can be as clear as a low-iron glass, they are much less thermally conductive, but the TIR cones are less than 4 mm tall and are conductively cooled from both faces. Silicone secondary optical elements (SOEs) have survived for years at 600x concentration with no visible degradation (Victoria, et al, 2014). Although TLC's concentration is 2.5x higher, those SOEs were 17 mm tall and were cooled on only one face (~10x farther for the heat to conduct), so TLC has over an order of magnitude less heating and ~3x less yellowing than an already adequate silicone SOE.

- *Molding Feature Size and Alignment:*

Polymer over-molding to improve the precision of pre-shaped glass optics has been used for lowering the cost of precision aspheric optics by over-molding lower-cost spherical glass optics (Wallace, 2017). TLC cones do not need the precision of optical instruments. The feature size achieved in low-cost silicone-on-glass Fresnel lenses for CPV is a ~2 μm radius (Luce and Cohen, 2010), so ~0.2% of the light will be lost at the cusps for the 900 μ length of the cone top (and none

on the width since the width cusps will be out of the optical path). It is likely that this can be improved with 40X less area to mold.

Analyzing a Fresnel lenses presentation (Jacob and Nitz, 2017) suggested a 1.5º draft angle being allowable for Fresnel facets where the other face was not steep (such as TLC's cone sides, as shown above), and at Peter Nitz's suggestion this was confirmed with Suncycle (who had molded the lenses analyzed in the paper). Suncycle also confirmed that draft angles as low as 3º should not be a problem on back-to-back facets (such as TLC's cone ends), and the current cones have an angle more generous than that.

However, Suncycle pointed out that SOG lenses need no alignment of the mold to the glass. While alignment requirements could be removed by over-molding the top lenses, that would put polymer on the lens front, which is not acceptable in the lens-as-cover-glass DAT version (unless in a glasshouse). The lens front would also be sensitive to thermo-optic effects changing the focal length (although no more than in current Fresnel/box CPV systems).

The roll-formed optics could be aligned to the mold; mold alignment has been done for centuries, and tight alignment can be met over large areas with CTE-matched molds. Alternatively molding can be done on smaller areas at a time since the cone arrays have millimeters between cones and each such molding can be aligned to the lens. The alignment tolerance is also increased by not concentrating further on the 2nd axis; the cone tops are significantly broader than the light from the lens, providing +/- 100 µm tolerance (which is huge by mold alignment standards where precision pins often align to 25 µm).

- *Adaptive Cone Mold:*

A better alignment method is to use a flexible mold that aligns to the optics. The thermal contacts roll formed into the bottom of the glass are a noncritical feature that is out of the optical thermal and electrical paths, so it can be converted to alignment feature as follows:

1) The thermal contacts are on the opposite face the lens tile from the lenses themselves, and so are produced by different roller. This roller may not be perfectly aligned to the lens roller so the thermal contact will be rolled undersized and then over molded with an accurate mold. The roller-to-roller misalignment should vary only slowly, so a large sheet of glass could be accurately aligned (using two lasers shining through the lenses as alignment means) to a CTE-matched injection mold. A hard plastic can then be injected to fill out the thermal contacts so

that their outer surfaces are accurately aligned to the lenses. This could easily handle up to a millimeter of roller misalignment (much more than is expected).

2) The large glass sheet is then singulated (cut) into individual lens tiles.

3) The cones are then molded onto a lens tile using a cone-array mold that comprises a series of cone-row molds with slightly flexible linkages between them. This flexible mold not only makes demolding easier, but it automatically aligns to the thermal contacts with no CTE match issues along the 1720 mm length of the lens tile. While thermal expansion still has to be watched across the ~ 40 mm tile width, this is a 40 X reduction in sensitivity. These cone-array molds can themselves be molded either in a thermoplastic polymer such as polyethylene (which does not inhibit the cure of addition-cured silicones), or in a setting polymer such as urethane which would then be Parylene-C coated to prevent cure inhibition of the silicone. Since the molds are themselves molded, they can be very low cost so a slow-curing silicone can be used for the optics without requiring a massive investment in cone molds.

4) Since this would provide high accuracy at low cost, and the flexible cones would adapt to variations in solder thickness and would allow much thinner optical coupling between cones and cells, this was one of the highest priority processes to prototype.

A sketch of this self-aligning cone mold is shown in the following figures:

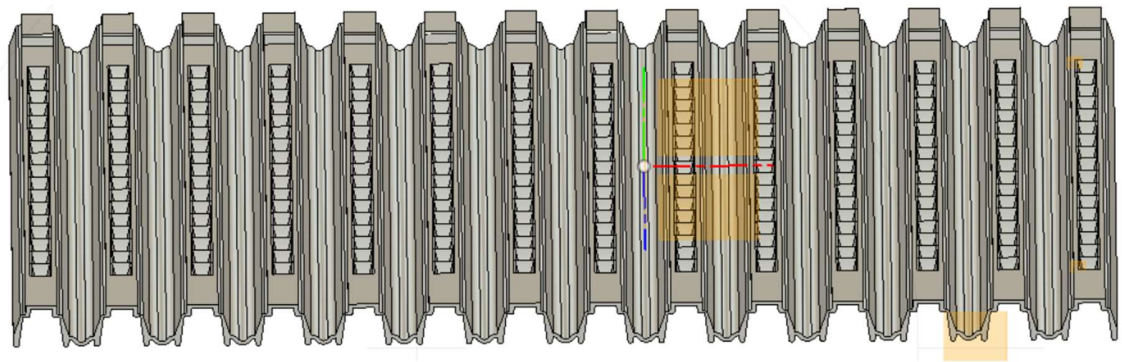

Figure 3.9: *Self-aligning Cone Mold (showing cone cavities).*

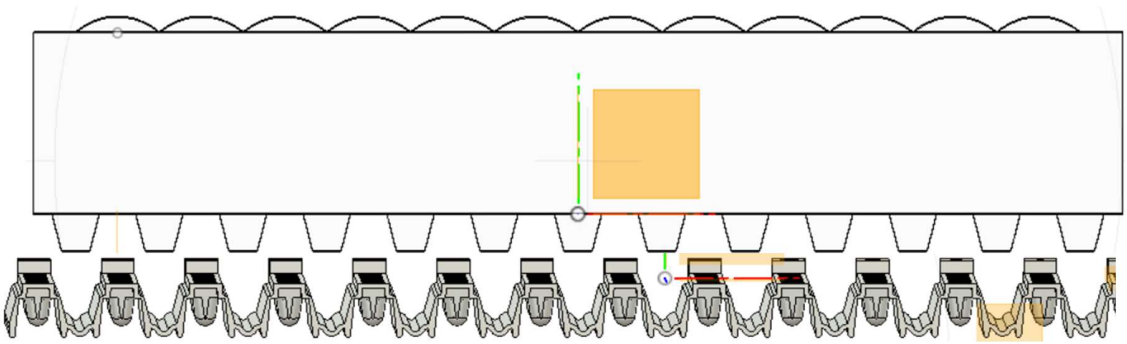

Figure 3.10: *Cone Mold Flexes to Mate to Thermal Contacts.*

Molding the Cone Mold: The basic shape of a single cone is easily formed through wire EDM, but there is no room for an EDM wire to pass between the cones of a row so simply using EDM to form an array of cones would leave many tens of microns between cones in a row, which would cost ~ 5% to 10% of the light (depending on EDM wire diameter). The solution being adopted in prototyping is to form the cone array from cone plates where each plate has one cone from each row (**Figure 3.11**).

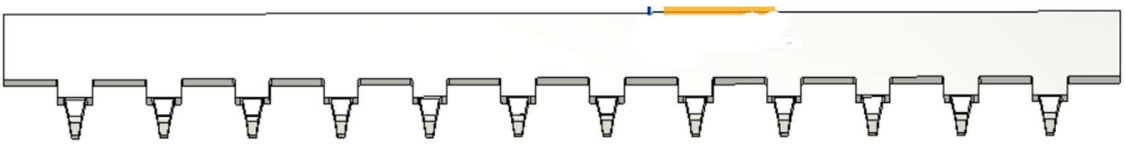

Figure 3.11: *A Single Cone Plate has One Cone from Each Row.*

The plates then pack together to form a cone array (**Figure 3.12**). Cross pieces (four shown) then hold the cone plates in an injection mold (rest of mold not shown).

Figure 3.12: *Cone Plates packed to form a Cone Array.*

While the overall concepts of the adaptive mold and the packed cone plates are from the PhD candidate, the details were worked out jointly with professional researcher Étienne Léveillé, who also did the CAD work for the mold components. A polymer that silicone does not stick to (e.g., polyethylene) is then injected into the cone-mold-mold to form a self-aligning cone mold such as the one shown earlier.

- *Confirming the self-aligning cone mold:*

The principles of the self-aligning cone mold have now been tested. The molding process worked well, and the main issue found was that EDM did not produce anywhere near as smooth a cone surface as expected. One cone plate has now been laboriously hand-polished to a near-mirror finish, and comparative tests between a polished cone and an unpolished cone have been used to confirm the improvement (see chapter 7).

Although hand polishing is exceedingly tedious, it would be affordable in production because even a short master cone array such as that used in prototyping could produce enough cone molds to mold cones for on the order of a gigawatt of TLC modules. Single-point diamond milling could also probably produce a sufficiently smooth finish without requiring hand polishing at all, and it would also be highly affordable for a production master cone array.

A summary of the developed and tested process and the results of using it has been accepted as a paper for CPV-17. The CPV-17 paper is included here because it contains additional details (and further details can be found in the Masters' thesis of Nathan Caillou).

3.4.5 Molding Arrays of Tertiary Optical Elements for Microcell Receivers

- ***Foreword:***

This paper contains additional details of the tested cone molding process. This paper is kept separate from the rest of this thesis because in addition to the work of the present PhD candidate, the paper also includes extensive work performed by people other than the present PhD candidate.

Authors: Richard Norman, Etienne Léveillé, Nathan Caillou, Jade Blais, Stephan Rosa, Vincent Aimez and Luc G. Frechette

All authors are associated with the University of Sherbrooke. Richard Norman is a PhD student, Etienne Léveillé is a research professional, Nathan Caillou and Jade Blais are Masters' students, Stephan Rosa is an undergraduate student, Vincent Aimez and Luc G. Frechette are professors.

Date of acceptance: September 21, 2021

Acceptance Status: Final Publishable Version Accepted

Reference: R. Norman, E. Leveille, N. Caillou, J. Blais, S. Rosa, et al. Molding Arrays of Tertiary Optical Elements for Microcell Receivers. 17th Conference on Concentrated PhotoVoltaics Systems (CPV17), Apr 2021, Freiburg (On Line), Germany. [Molding Arrays of Tertiary Optical Elements for Microcell Receivers]

Contribution to the thesis:

This article contributes to the thesis by detailing the development of a molding process that would be practical for producing cone arrays on the back of lens tiles. The article shows that densely packed arrays of transparent cones of sufficient flexibility can be cast in optical silicone using an adaptive, self-aligning mold, and that such molds could themselves be molded in low-cost plastic using a steel master cone array.

French title: Moulage de Matrices d'éléments Optiques Tertiaires pour Récepteurs Microcellulaires

French abstract:

Trough-Lens-Cone PV, ou TLC, est une architecture de module CPV qui utilise une optique à trois étages pour permettre une concentration élevée à faible coût. Un miroir creux concentre ~40X sur un axe sur un module long et étroit. Le couvercle du module a des lentilles linéaires à l'avant qui se recentrent ~ 10X sur un deuxième axe, produisant une série de lignes focales étroites ~ 400X. L'étage optique final est un réseau de cônes CPC (courbe parabolique composée) qui augmentent la concentration à environ 1500X. Un récepteur a un réseau de microcellules en parallèle pour gérer la concentration inégale du creux, et chaque cône dirige la lumière vers une microcellule. Les creux paraboliques appropriés sont déjà peu coûteux, les lentilles sont conçues pour être roulé à très faible coût et les réseaux de microcellules peuvent être assemblés à faible coût. Cependant, la formation des matrices de minuscules cônes est un obstacle potentiel. Cet article décrit les critères auxquels les cônes doivent répondre, les défis de la production de cônes appropriés, le processus de moulage

des cônes proposé et les premiers résultats du prototypage de validation de principe du processus proposé.

Note: No corrections have yet been requested by the members of the jury.

- ***Article:***

Abstract: Trough-Lens-Cone PV, or TLC, is a CPV module architecture that uses three-stage optics to enable high concentration at low cost. A trough mirror concentrates ~40X on one axis onto a long, narrow module. The module cover has linear lenses on front that reconcentrate ~10X on a second axis, producing a series of narrow ~400X focal lines. The final optical stage is an array of CPC (compound-parabolic-curve) cones that raise the concentration to ~1500X. A receiver has an array of microcells in parallel to handle the trough's uneven concentration, and each cone funnels light to a microcell. Suitable parabolic troughs are already low cost, the lenses are designed to be roll-formed at very low cost, and the arrays of microcells can be assembled at low cost. However, forming the arrays of tiny cones is a potential obstacle. This paper describes criteria that the cones should meet, the challenges of producing suitable cones, the proposed cone-molding process, and first results from proof-of-concept prototyping of the proposed process.

- *Background:*

While CPV once claimed the potential to lower the cost of photovoltaics through replacing expensive semiconductor area with inexpensive optics and metal, entire silicon PV panels now cost less per watt than the optics and metal of the most common type of CPV. A detailed 2015 cost study by NREL (Horowitz, et al., 2015) shows the challenge facing this common CPV architecture, hollow-box Fresnel CPV. At 100 MW/year production, the 67 ¢/W CPV module manufacturing cost would be higher than the then-current average flat-panel selling price of 57 ¢/W (Munsell, 2016).

NREL's analysis found potential pathways to reduce the ~18 ¢/W CPV cell cost by >95% (per area) at very-high-volume production (tens of GW/year), but only pathways for ~20% reduction in the ~35 ¢/W full-module-area component costs (initially comprising ~17 ¢/W for the Fresnel lenses, ~12 ¢/W for the housing and ~6 ¢/W for the aluminum-plate thermal management). With ~25% improvement in the ~10 ¢/W non-cell receiver costs, NREL's cost reduction pathways still left a ~38 ¢/W CPV module manufacturing cost, and thus, with NREL's estimate of a minimum sustainable 15% margin, a ~44 ¢/W selling price (USD). But while 44 ¢/W would have been

competitive in 2015, silicon PV costs have continued to fall through ever-greater economies of scale and cumulative learning.

NREL found that if the module efficiency could also be improved by 50% (e.g., a 60% efficient cell), a minimum sustainable selling price of 29 ¢/W could be reached for Fresnel CPV. Since neither 60% cells nor a 95% drop in cell cost per area have been achieved, and silicon PV is already below 29¢/W (~25 ¢/W European spot market price in Sept. 2020 (Schachinger, 2020)), it is hard to see how Fresnel CPV can compete with silicon PV on cost.

Fresnel CPV may win on other attributes; for example, the University of Sherbrooke campus has ~250 kW of STACE Fresnel CPV, and the partially transparent modules let most diffuse light through, in contrast to the oppressively-dark silicon PV flat panels nearby, and a much more attractive installation might command a significant price premium in a market like community solar. However, by far the largest solar market is utility scale solar (Feldman and Margolis, 2020), and if CPV can gain a toehold at utility scale it can benefit from the virtuous cycle of economies-of-scale lowering costs which in turn grow the addressable market and bring further economies of scale. Cost is critical at utility scale, so CPV growing beyond niche markets depends both on achieving low cost at modest volume and having the potential for very low cost at the ~ 100 GW/ year scale at which silicon PV is made.

- *TLC Overview:*

Trough-Lens-Cone PV, or TLC, is a CPV architecture introduced at CPV-14 (Norman et al., 2018) that uses three-stage optics to enable high concentration at low cost. As shown in Figure 1, in TLC a trough mirror (T) concentrates ~40X on one axis onto a long, narrow module at the trough's focus. The module cover has linear lenses (L) on front that reconcentrate ~10X on a second axis, producing a series of narrow ~400X focal lines (**Figure 3.13b**). The final optical stage is an array of CPC (compound-parabolic-curve) cones (C) that use non-imaging optics to raise the concentration to ~1500X, producing an array of small focal spots. Each cone funnels light to a microcell, and a receiver has an array of microcells in parallel to handle the trough's uneven concentration. From the mirror, glass, antireflective coating, and optical silicone data sheets, optical efficiency should be ~90%, and ray tracing shows acceptance angles of ± 0.425° on the trough-focusing axis and ± 0.69° on the lens-focusing axis (Norman et al., 2018).

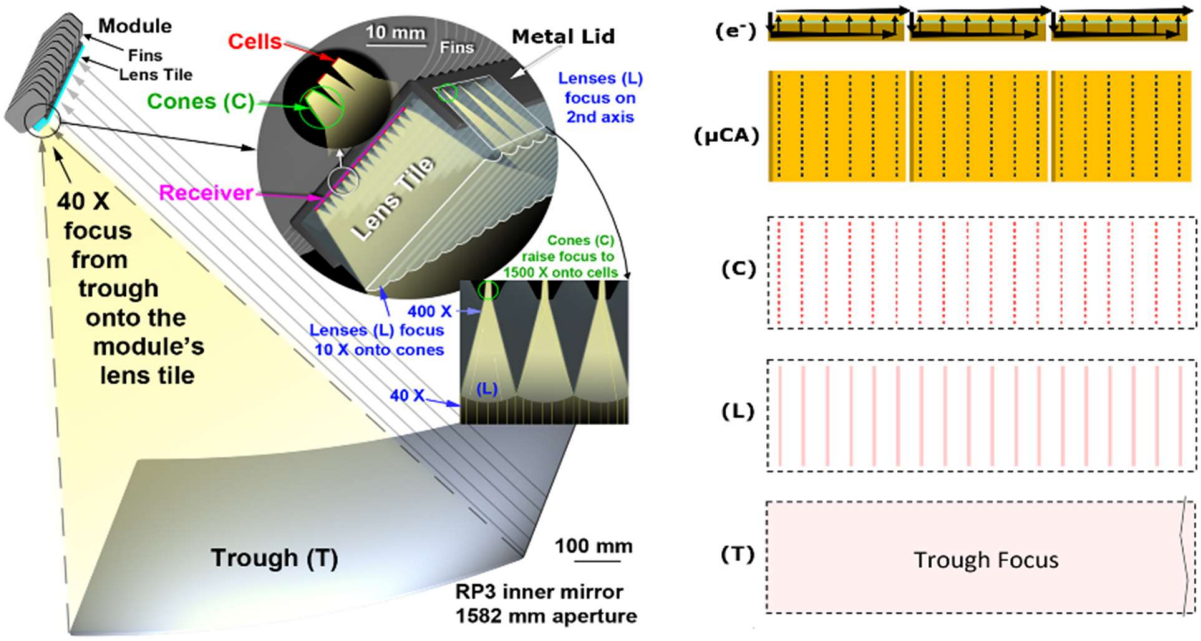

(a) (b)

Figure 3.13: *(a) Visual diagram. (b) Energy flow through optics and receivers.*

T: A **T**rough mirror focuses ~40X on the first axis onto long, narrow module (2-axis tracking).

L: Linear **L**enses on module front focus ~10X on the 2nd axis, forming short ~400X focal lines.

C: CPC **C**ones funnel the short 400X focal lines into arrays of small 1300X-1500X foci.

μCA: Microcell arrays on small receivers mate to the cones. A receiver's 96 cells are in parallel.

e-: Electron path to the next receiver is short for low resistance. A ~1700-mm-long module with 32 receivers in series (3 shown) will produce ~1 kW (3J cell MPP ~11.5A, ~87V).

A receiver with tens of Watts of microcells fits high-speed placement and bonding equipment for low assembly cost, and receivers connect in series as they are mounted on a lens tile. Suitable mirrors are already produced at low cost for solar thermal (Rioglass, 2017), and the linear lenses can be roll-formed in sheets of low-iron glass at low cost, but forming the cones is a potential obstacle to reaching the module cost target of ~16 ¢/W in gigawatt volume (Norman, et al., 2019).

- *Cone criteria and proposed cone-molding process*

In TLC, each linear lens focuses light onto the back of the lens tile, where it is intercepted by a row of cones. To efficiently guide light from the lenses to the cells, the ~3.7 mm tall cones should be clear and UV-resistant, accurately shaped, smooth surfaced, aligned to the lenses, tightly packed

within a row, and be compressible by a few tens of microns to handle cell and solder thickness variations and to accommodate bond-wire tails on cell surfaces.

Optical silicone can meet these requirements (S. Li, 2012). Optical silicone is also already used in CPV; it is over-molded on low-iron glass for Fresnel lenses, and it is used to optically couple glass secondary optical elements to cells. Due to the trough's 40X concentration, a TLC lens tile has only 2.5% the area to over-mold (per Watt) as Fresnel CPV lenses (Norman et al., 2018), and the cones use 2% as much silicone, so even high-grade optical silicone can meet TLC's cost target.

Optical silicones can be injection-molded, but the cone array would be challenging. While light through a lens on each end of the module could be used to precisely align a roll-formed lens tile to an injection mold, the ~100°C temperature change involved in optical silicone injection molding (Hopmann and Röbig, 2016) would make maintaining alignment challenging; even the 1.2 ppm/K Coefficient of Thermal Expansion (CTE) difference between the glass at 8.7 ppm/K (Vitro, 2020) and a low-CTE stainless steel (410 stainless is 9.9 ppm/K) would shift the alignment by twice the ±50 µm cone-to-lens tolerance budget (Norman, et al., 2019) at each end of a 1700-mm-long lens tile, so multiple optical alignments per lens tile would be needed with a steel mold. Alternatively, a mold could be made of titanium (8.6 ppm/K), but even free-standing titanium is not easy to polish to a mirror finish, and polishing the insides of the more than 3000 tiny cone cavities (which do not have round cross-sections), would be exceedingly tedious at best.

While injection molding might be practical for large-scale production, even a small such mold would be prohibitively expensive for proof-of-concept prototyping. The much deeper aspect ratio of the cones also makes injection molding harder for TLC than for Fresnel CPV. While the challenges could probably be solved for production, they make injection-molding the silicone difficult for prototyping.

Optical silicone can also be cast in a no-to-low-pressure process, which is easier for prototyping. However, silicone's long cure time at lower temperatures means that in production inexpensive molds would be needed to avoid high capital costs, and our goal is to prototype (when practical) with processes that can scale to production.

- *Molding the Cone Molds:*

The solution adopted is cast the silicone in cone molds that are themselves molded from polyethylene, a low-cost, easy-to-mold plastic that silicone does not stick to, and that doesn't

inhibit the curing of optical silicone. A full-module polyethylene mold (1700 mm long) would be hard to handle, and consistent lateral alignment is needed within a receiver, so shorter molds that are each a multiple of the receiver length are used. Along the lens tile, alignment is critical (Norman, et al., 2019), so the cone molds have flex zones between rows of cone cavities, and features on the back of the glass will align a cone mold at every row (**Figure 3.14**). In production these alignment features (which will also mechanically support the receivers and help conductively cool the glass) would be roll-formed in the glass itself at the same time as the lenses, but in prototyping they are molded in epoxy.

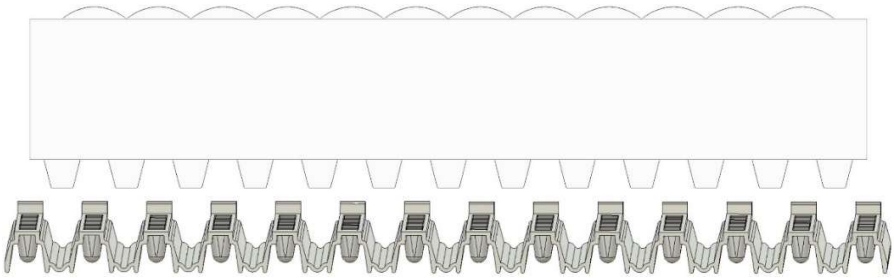

Figure 3.14: *polyethylene Cone Mold can stretch between cavity rows to align to features on lens tile*

The cones must be tightly packed within each row to prevent loss of ~400 X light between them. Because the cone-to-cone pitch within a row is only 2.09 mm, each 20 µm of gap would cost ~1% of the light. The master cone array used to mold the cone molds must therefore itself be made with tightly packed cones. A cone array is the intersection of two orthogonal curves, so wire-EDM (electrical discharge machining) was used to cut a hard steel (HRC 55) master array. To prevent a >100 µm inter-cone gap from the EDM wire kerf, and to allow polishing all cone surfaces, the master cone array comprises a series of plates which each contain one cone from each row, as shown in **Figure 3.15**. Interlocking cross-parts align the cone plates and lock them together to form a top insert.

A bottom insert was also cut with EDM (with minor post-processing on a milling machine), and these precision inserts (unpolished, so the resulting cone array is more visible) were then bolted into a cavity milled into a pair of 24-mm-thick steel plates, forming a sturdy steel injection-mold suitable for molding polyethylene cone molds.

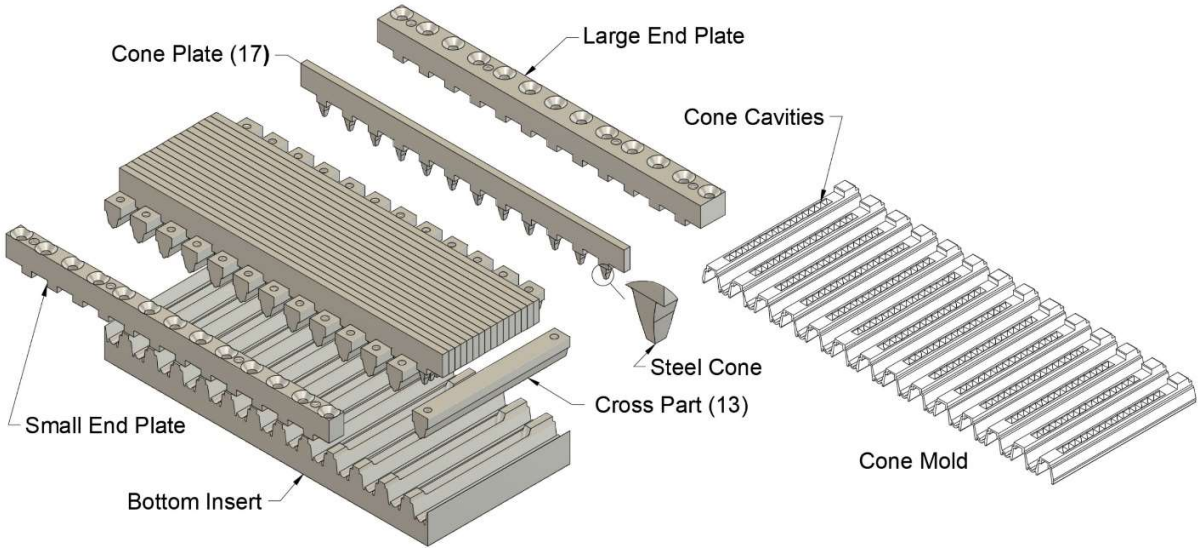

Figure 3.15: *Precision Parts of Mold for Cone Molds and resulting Cone Mold.*

- *Results:*

Molten polyethylene was injected using a Benchtop Injection Molder (Medium Machinery, LLC). Cartridge heaters were added to the platen to control the mold temperature, and with the polyethylene reservoir at 215°C, the injection nozzle at 130°C, and the mold at 105°C, the polyethylene flowed well into the mold corners and the narrow gaps between cones, without visible flowlines or merge lines (**Figure 3.16a**). This polyethylene cone mold was then used to mold a cone array on a 19-mm low-iron glass block (**Figure 3.16b**) using degassed optical silicone (Sylgard 184), producing rows of tightly packed cones, ~ 3.7 mm tall, in flexible optical silicone.

Cone compressibility was measured with a TA.XTplus texture analyzer, and averaged ~0.8 mN/µm. The variation in cell and solder thickness is expected to total less than 50 µm, and even to compress all 96 cones for a receiver by this amount would only require 4 N of force, which is well within the range of pick-and-place machines (Assembleon, 2021). Once placed, a receiver can be held in place with drops of instant adhesive while the various epoxies cure.

After the first few cone arrays, a prototyping recipe was developed. A low-iron glass block (19-mm Vitro Starphire glass) is prepared with Sylgard 184 lenses and with Epotek 301 epoxy for the alignment features (in production the lenses and alignment features would be roll-formed in the glass itself). The glass is then primed with DOWSIL 1200 OS adhesion promoter where the cone

rows will go and coated with Ease 200 Release (sprayed through a mask) where adhesion is not wanted.

The cone mold cavities are filled with Sylgard 184 (10:1 mix) optical silicone and then slowly degassed under increasing vacuum. The rows of cone cavities are then topped up with silicone to obtain a meniscus, and the cone mold is placed in a jig.

A bead of silicone is run along the exposed glass where each cone row will go. The glass block is then carefully placed on the cone mold, starting from one edge and lowering the other edge until the glass rests on the cone mold. A metal plate is then clamped to close the jig. After curing the silicone for 60 minutes at 100ºC, the lens tile is de-molded by prying carefully and slowly between the mold and the alignment features.

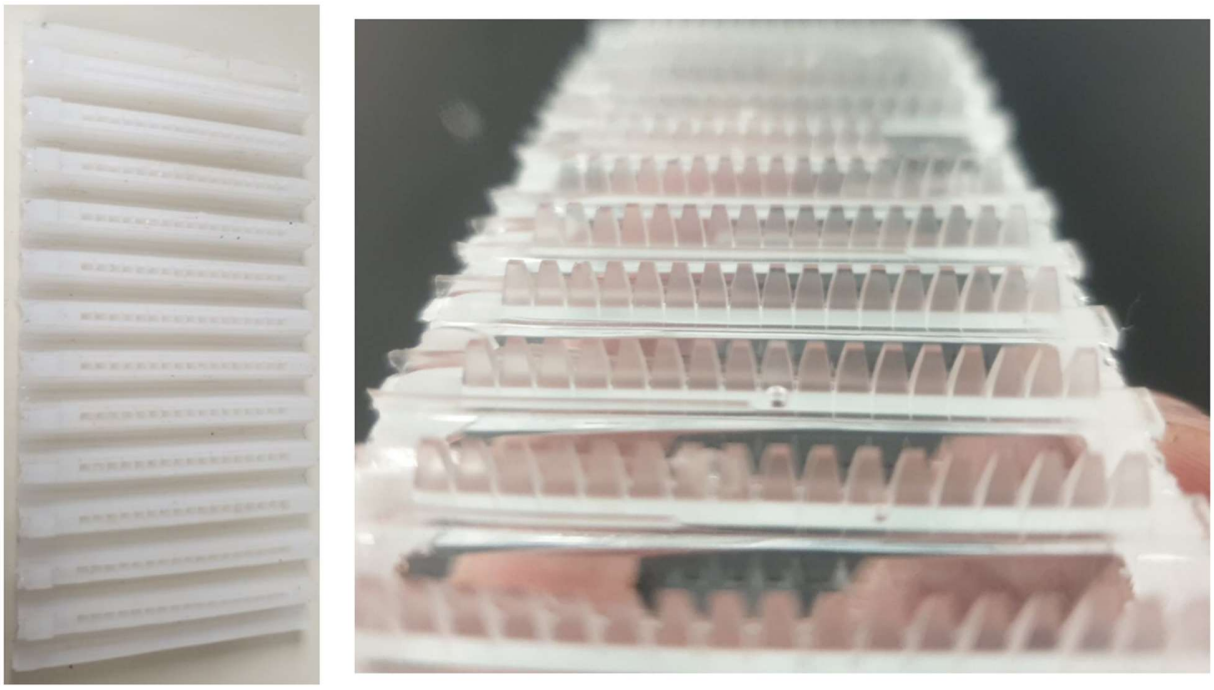

(a) (b)

Figure 3.16: *(a) Cone Mold molded in polyethylene, (b) Cone array over-molded on low-iron glass in optical silicone.*

- *Next Steps:*

The first complete lens tiles will now be used to build TLC minimodules, testing the receiver-to-cone alignment process, and then testing the receiver-to-receiver interconnection via copper strips

on the lens tiles. After that the first functional TLC mini-modules will be assembled and, if all goes well, tested on sun.

Meanwhile cones plates with polished cones will be made. Tests with optical silicone cast on polyethylene that had been molded against a surface roughness specimen card showed that surface roughness greater than 100 nm Ra affects the optical quality of the final silicone surface, so the final cone surfaces should be close to this smooth. A wire-EDM recipe has been found that is supposed to produce ~0.22 µm Ra surface finishes (Y.S. Liao, 2004); this will be used to greatly improve the starting surface smoothness, and electropolishing should be able to improve this to ~0.1 µm Ra to avoid (or at least greatly reduce) laborious hand-polishing of the multitude of tiny cones of the master cone array. New cone molds will then be molded and used to mold cones on lens tiles for final characterization.

- *Conclusion:*

In production the lenses would be roll-formed in low-iron glass as an integral part of the lens tile, providing a durable, cleanable module surface. In the interior of the module, the flexible silicone cones will be protected from dirt, moisture, and shorter ultraviolet wavelengths by the thick glass front and the metal heatsink lid.

Sylgard 184 optical silicone cones are compressible enough to overcome receiver cell-thickness and solder-thickness variations, and Sylgard 184 casts well in polyethylene molds. A hard steel master mold can typically produce over a million polyethylene parts (Borgerson, 2009), and each polyethylene cone mold should be reusable many times, so even a 12-cone-row master mold (cones for ~66W of microcells) should mold enough cone molds to mold cones for hundreds of megawatts to potentially gigawatts of TLC modules. Molding of the cone molds is thus a potential low-capital way to produce flexible optical silicone cones accurately aligned to roll-formed lenses on a glass lens tile.

- *Acknowledgments:*

The authors thank the University of Sherbrooke for access to expertise and test equipment, Programme Innovation, IRAP, MITACS and TFI for funding, and TFI for access to its detailed TLC design and use of overview illustrations and descriptions for TLC. TLC is a trademark of, and the TLC design is the intellectual property of Terra Firma Innovations Inc. (TFI). The authors thank LN2 - a joint International Research Laboratory (Unité Mixte Internationale UMI 3463)

funded and co-operated in Canada by Université de Sherbrooke (UdeS) and in France by CNRS as well as Université de Lyon (UdL, especially including ECL, INSA Lyon, CPE) and Université Grenoble Alpes (UGA). LN2 is also associated to the French national nanofabrication network RENATECH and is supported by the Fonds de Recherche du Québec Nature et Technologie (FRQNT)

3.4.6 Alternative Cone Array Manufacturing Methods

- ***Plunge-EDM Cone Mold:***

It would also be possible to make a flexible metal cone mold through plunge EDM. Either 400 series stainless steel or titanium would have a good CTE match to the lens tile. A much faster silicone cure would be desirable to keep the capital tied up in molds low, and a CTE-matched mold would allow curing at an elevated temperature. Even a $100,000 mold would be highly affordable if it could be used several times per hour.

- ***Alternative to Over-molding:***

An alternative to over-molding would be to mold rows of cones and insert the cone rows into the lens glass. A cone row would have a 'root' that would tie its cones together and anchor the row in the lens, with an optical adhesive refraction-matched to the lens glass holding the row in place.

The volume is small enough that an instrument-grade glass would be affordable; if the root has as much volume as the cones it would still be roughly 50 times less glass than the lens itself. The NREL costing does not state the SOE size, but a dome secondary for a 5-mm cell has roughly enough glass to cover the cell 25 mm deep (Espinet-Gonzalex, et al., 2012), and even doubling the cone volume for the root it would only cover the cells 15 mm deep, and adjusting for the higher concentration (1.5x) and higher optical efficiency using the same cells (1.2x), TLC would use roughly 1/3 the glass volume.

While a cone row is only 5.5 Watts instead of the 7.5W in the NREL study, inserting the cone rows before sawing the individual tiles would dramatically reduce the parts handling. Precision molding can start with glass rod as the blank, so long rows can be molded - a meter-long insert in a meter-wide sheet would produce ~132W per insert, or 17.5 times as much per part handled as the Fresnel costing. Since the insert volume is ~1/3 of the cost Fresnel secondaries, the insert length for TLC costing is chosen to be such that the handling is also 1/3 that of the 5-mm-cell Fresnel reference case. No ARC is used (since the row is optically coupled), so the Fresnel costing is 2.7¢/W for molding and 0.7 ¢/W for attachment, or 3.4 ¢/W, and 1/3 of that is 1.13¢/W.

While this is 0.55 ¢/W higher than the cost of over-molding silicone, it is an affordable alternative. Glass cones would CTE-match the lens, and instrument-grade glass is extremely clear for a slight

performance boost. A higher refractive-index material could also be used to increase acceptance angles, and, with re-tuning the cell AR-coating, also increase transmission into the cell.

3.4.7 Lens Sheet Cost:

- ***Basic Roll-forming Cost:***

Roll-forming is so low cost that it is used for texturing of low-cost silicon-module cover sheets that cost \$4/m^2 (including the glass). For silicon flat panel cover glass, roll-formed textured sheet actually cost less than flat sheet since roll-formed glass costs less that float glass.

Since the 19-mm glass used is made on float lines anyway, at least in the short-term rolling cost must be added to the glass costing. Based on a \$4/m^2 total minus the raw cost of the glass, roll-forming is allocated \$1/m^2 for 3.2 mm glass. At low volume this is increased proportionately to the thickness (since more glass has to be reheated), but in high volume this would be done on raw glass while the glass is still warm so the cost will be invariant for thickness. Even \$6/m^2 for 19 mm glass in low volume is insignificant – at ~40% efficiency and 40x initial concentration at the front of the lens, the 14 kW from each square meter of lens brings this to 0.04 ¢/W, and this falls to less than 0.01¢/W in very large volume when the roll-forming is done in-line.

An upper bound on the cost of lenses in high volume can be obtained by studying similar lenses in solar. The most similar is the PARC (Palo Alto Research Center) monolithic-optics Cassegrain design that SolFocus called "Gen-2" (Elrod, 2008). The cost of molding the optics was not shown separately, but SolFocus gave a whole-product target cost of \$0.35/W, including the raw glass, roll-forming the glass, mirroring the glass, the cells, and assembly. At the 475x concentration and \$4/cm^2 cell price shown in a PARC presentation, the cells alone would be ~27 ¢/Watt, leaving 8¢/W for everything else including the lens. TLC will get ~50 times the Watts per lens area, so roll-forming a TLC lens can be at most ~0.17 ¢/W, but that is highly pessimistic since SolFocus's cost includes the glass itself, mirroring, cell placement, diodes, etc. as well as roll forming.

- ***Chemical Strengthening:***

19-mm glass is pretty strong, so strengthening should not be needed. If strengthening is desired, tempering would probably not be suitable because it would have to be done after cutting the individual tiles, so the AR coating and the cones would have to be done on individual tiles.

However chemical strengthening is also low enough cost that Flabeg uses it on trough mirrors and claims a result as strong as tempered glass (at 120 mPa surface fracture resistance), and chemically toughened glass can be cut afterward. If it is comparable in cost to tempering it would be ~~0.1 ¢/W and would be worth it for further insurance (good marketing for a 50-year product).

- ***Anti-Reflection Coating Cost:***

AR coatings are cost-effective enough to be affordable on $1/W total-installed-cost 16% solar panels. The difference is price between solar module covers with and without AR coatings is only 15¢ per square meter; this is multiplied by 10x to allow for the handling cost of the thicker glass. With the ~40x primary concentration onto the lens this is still only ~0.01¢/W.

- ***Over-molding Cost:***

The amount of silicone used is tiny. Even with fully-silicone cones, the cones bases are only <20% of the lens area and they average ~2 mm deep, or ~0.25 mm averaged over the whole lens. This is thinner than the 0.3 mm average silicone depth in the NREL Fresnel costing study, where the 9.1 ¢/W material cost is $27.30/m^2. But $4/m^2 of that is the low-iron glass, leaving $24.30/m^2 for the silicone cost, or $17/m^2 for the slightly thinner (average) TLC silicone. However, TLC's lens uses 50x less area per Watt; at over 14 kW/m^2 of lens with current (Azur 44%) cells, the silicone cost is only 0.11 ¢/W. The lower purchasing volume at for 10 MW of TLC increases this to $0.17 ¢/W.

This is pessimistic because in an SOG lens the average thickness is significantly more than half the deepest thickness because the thickness between the facets (teeth) does not go to zero, but to a base thickness that connects the facets. While TLC will probably also have a contiguous base, the cones are much narrower and taller so the base will be relatively less significant. Even if the base extends a bit up the sides of the indent (for larger area for adhesion), TLC will use a fraction of the base volume. However, the savings are only ~~0.02 ¢/W so it is not worth including them.

Using half of the ~100 μm alignment tolerance should allow molding the whole lens length at a time, for an area at least as large as Fresnel SOG lens sheets, so the non-material costs should be the same per area. The NREL cost study shows non-material costs of 7.6¢/W, which is $22.8/m^2, which the concentration onto the lens reduces to a small 0.26 ¢/W at comparable production scale. At a smaller production volume, costs will be larger, increasing this to 0.42 ¢/W at 10 MW.

- ***Lens Tile Cutting Cost:***

It will be far easier to cut a 19-mm-thick sheet of glass into sturdy 43-mm wide bars with millimeter accuracy than it is to cut a 150 mm silicon boule into fragile 0.15 mm thick silicon slabs with tens-of-micron accuracy. Each meter of lens also only requires 200 cm^2 of cutting for over 650W, while a polysilicon cell producing only 3.6W requires 225 cm^2 of cutting. Thus, dicing the lens takes 200 times less cutting per Watt than silicon PV, and with less-demanding accuracy.

In silicon each meter of diamond wire at 0.58¢ produces 80 cm^2, or 0.0074 ¢/cm^2. That is 0.87¢ for each meter of lens tile (Goodrich, et al., 2013), and that is pessimistic because wire is the main cost and quartz cuts ~20% faster than silicon (South Bay Technology, undated), and the form factor is much easier to work with. That is a tiny 0.002¢/W. For the initial 10 MW this is increased 10x due to using a silicon carbide saw or diamond saw rather than bulk wire sawing.

- *Sanity-checking the Lens Cost:*

NREL's Fresnel costing uses low-iron glass 1/6 as thick, and AR coats it with a CPV-compatible weather-proof AR costing, and overmolds silicone of comparable average thickness on it. In the TLC costing 1/6 the glass plus the coating plus the cones at 100 MW is 0.51¢/W. Fresnel CPV needs 48 times more area per Watt; that 25¢/W is 50% higher than the 16.7¢/W shown by NREL, this difference matches lower cost from the 500x more area used in 100 MW of Fresnel lenses than 10 MW of TLC. at the typical 5% per doubling cost reduction with production volume.

- ***TLC Lens Temperature***

- Absorption in Glass:

Each cell's worth of lens absorbs 2% of 1.5W/mm^2, or 30 mW/mm^2, for 0.78 mm^2, or ~23 mW.

- *Absorption* of silicone:

1 mm of refraction-matched optical silicone, AngstromLinks's Al-3252 (Elgin, Undated), absorbs:

Band Center nm	Band Width nm	%/mm	Spectrum W/m^2nm	mW/mm	Loss in ~ 3 mm path in cone	% of junction
375	50	2	0.3	0.45	1.35 mW	0.3% 1st
900	50	0.3	0.8	0.12	0.4 mW	0.1% 2nd

1175	100	4	0.5	2.0	6 mW	1.5% 3rd
1400	100	4	0.05	0.2	0.75 mW	0.4% 3rd
1500	100	2	0.25	0.5	1.5 mW	0.8% 3rd
1650	100	12*	0.25	3	9 mW	No Jct
Total				Heat = 6.3 mW/mm		3rd Jct = ~2.7%

* (DeGroot, undated)

Table 3.1: *Absorption of Optical Silicone.*

This absorption totals a bit more than 6 mW/mm deposited in the silicone cones, but since there may be some after 1700 nm (silicone is pretty clear after 2500 nm but an absorption spectrum in the 1700-2500 nm range has not been found), so 10 mW/mm = 30 mW is allowed for.

- Silicone refractive index:

RI 1.52 phenyl silicones used to have unacceptable susceptibility to UV and thermal yellowing, and while new phenyl silicones our improved, and while TLC filters the harsh UVB and cools the cones well, suitability of a phenyl silicone would need to be verified experimentally. Methyl silicones are clearer than phenyl silicones, reducing absorption losses and also cone heating. Methyl silicone can handle the expected heat and UV (Blais, 2022), but have a low refractive index (RI ~1.41); the maximum concentration scales with the square of the refractive index so lower RI means higher cell cost. Nanoparticle loaded methyl silicone should provide the best of both worlds with RI as high as 1.6 (Huang, et al., 2017).

- Steady State Temperature:

With each cell's worth of lens absorbing ~23 mW, and roughly half of the heat from the silicone cone, or ~15 mW, also flowing through the lens, the total in a cell's worth of lens is ~38 mW. Half of this, or ~19 mW, must be conducted in each direction to reach the thermal contact pad between the cell rows. This path is 2.77 mm wide, and it starts 18 mm tall and narrows to ~2.5 mm over roughly 10 mm of travel. Even if the whole path were 2.5 mm x 2.77 mm in cross section, the thermal flux would only be ~3 mW/mm^2, and glass has a thermal conductivity of ~1 mW/mmK so

the maximum temperature would be 30K over the substrate temperature, and the much larger area for most of the path cuts this to ~~10K, which is insignificant for glass.

These estimates are confirmed by COMSOL finite-element analysis of a 3D-CAD lens model (see **Figure 6.8**).

- Initial Heating:

The lens tile absorbs ~2% of the incoming energy (1.2% of the usable energy, plus additional infrared). 2% of 3.5 W/cm^2 is 0.07 W/cm^2, or ~0.04 W/cm^3, so with a density of 2.53 and a specific heat of 0.88 J/gK, heating takes 2.2J/K and the glass warms at ~1º per minute. With the mirrors having already absorbed wavelengths that are strongly absorbed by glass, the energy is fairly evenly deposited as a function of depth. Horizontally the distribution is less even: the highest concentration that an area of lens sees is ~400x at the top of the TIR cones (12º/min), but the heat only has to be conducted an average of 2 mm to even out the temperature. The copper heat spreader warms faster than this, so *initially* the lens gain heat from the receiver. A low thermal conductivity adhesive heats the lens more gradually for lower stress, so a cheaper adhesive is actually better.

- Silicone TIR Cone Temperature:

The lens near the cones also influences the cone temperature. If the bottom 2 mm of glass carries the cone's heat as well as its own heat, that's ~5 mW through a 5 mm path 5.4 mm in cross section, or about 5K (COMSOL has been used to confirm this), so the lens near the cone is about 15K above the heat-spreader temperature.

Roughly half of a cone's ~30 mW of heat flows into the lens and half into the cell. 15 mW into the lens flows ~1/4 of the height of the lens, or 0.65 mm, on the average. That part of the cone averages ~ 3 mm^2, and silicone is typically 0.2 mW/mmK so the middle of the cone will be ~16K warmer than the lens and thus ~30K warmer than the heat spreader. Silicones can easily handle such temperatures, with most being able to take ~200ºC. Again, this is confirmed by COMSOL analysis of the 3D-CAD model (see **Figure 6.9**).

Midriff Bulge: The most noticeable effect of heating will be a slight bulge in the middle of the cone that slants the cone sides by less than a degree. This is related to the facet distortion discussed

in Jacob and Nitz for SOG Fresnel lenses. While the bulge from heating while on sun varies little (because the effect of the much cooler off-sun temperature has no effect on the light gathering), the on-sun range is slightly larger than the ambient range, so the cone will be tuned for 40ºC to 80ºC average cone temperature (ambient range of roughly 10ºC with low sun to 40ºC with full sun).

With a CTE of ~200 ppm/K, in three directions at an average of ~20K, the silicone's volume will expand by ~1.2%. It will have two directions to expand in, but will be pinned on two faces, so the middle will bulge ~0.8%. On the entrance to the cone this will actually improve cone shape (by decreasing the draft angle left from the molding), and it has no effect in the middle of the cone, which is free to move outward.

Near the cone tip there is a slight distortion that is most significant across the trough, where the cone expands by ~6 µm in each direction. This is ~0.4% of the 1600 µm half-height of the cone, or an average of ~0.2º distortion, but it will be higher near the cone tip where the cone is partially constrained by the optical coupling. There is more tolerance for angular change near the tip so there should be no adverse impact (which should be confirmed through modeling).

Thermo-Optic Effect: Silicone's thermo-optic effect is ~$4x10^{-4}$/K, so a swing of +/- 20K is a +/- ~0.008 change in silicone's refractive index, which is negligible (in SOG Fresnel this is significant because the silicone is responsible for much more of the concentration and concentrates through refraction).

3.4.8 Acceptance Angles

- ***Across the Trough:***

The 1st-axis acceptance has a sharp cut-off: ray tracing (section 6.3.1) showed no misses in 10K rays spread over a 0.425º half-angle, while light beyond that angle simply missed the lenses or the cones. The sun's size consumes 0.27º, and CSP troughs routinely achieve mirrors, mounting and tracking within the remaining 0.15º, so this acceptance angle should be adequate. TLC's single inner mirror will also be easier to align than the two inner and two outer mirrors of a CSP trough, and 2-axis tracking accuracy targets are even tighter than troughs (Mehos, et al., 2016).

With a 41 mm lens front, shifting by 20.5 mm costs half the light; at a ~1900 mm average distance to the trough that is 10.7 mrads or 0.61º, so miss-tracking by ~0.61º will cost half the light and

roughly half the power. Optical acceptance angles include the size of the sun, which adds 0.265º. Thus the 50% optical acceptance angle should be just less than 0.9º.

The drop-off is low until the edge of the sun from most mirrors starts to miss the lens at about 10 mm off focus, so the 90% tracking acceptance angle is very roughly at 10 mm, or 5 mrads or ~~0.3º. This makes the estimated 90% optical acceptance angle ~~0.55º.

- ***Along the Trough:***

Cone-to-lens mold alignment influences 2nd-axis acceptance. With +/-100 µm alignment modelled, ray tracing showed no misses in 10K rays within 0.69º, with a gradual cut-off after that. Trough mirrors are less accurate on this axis, so the trough needs 0.22º, the sun needs 0.27º, and dependency of the lens glass's index of refraction on wavelength consumes the equivalent of 0.2º, so this acceptance should also be adequate.

Along the trough the acceptance angle is limited by the light reaching the tops of the TIR cones. A cone top has a half-width of 0.55 mm, so when the focus is shifted by 0.55 mm roughly half of the light is lost. At 18 mm from the top of the lens, this is ~30 mrads or a tracking 50% acceptance angle of ~1.7º and an optical 50% acceptance angle of ~2º. The cut-off is gradual, so the tracking 90% acceptance angle is estimated to be about sqrt(5) times less beyond the all-light angle than the 50% angle is, or roughly 1.2º

- ***Acceptance Angle and Sun-Overhead Tracking:***

The sun's azimuth suddenly changes by 180º if the sun passes directly overhead, and an altitude /azimuth dual-axis tracker (the most common type) cannot spin infinitely fast to keep up. However, the tracker can be ahead or behind by the 90% tracking acceptance angle with a <10% temporary impact on output. When the sun is nearly overhead, a big change in the azimuth also creates only a small change in the angle of the light to the trough, so a good acceptance angle allows a slower maximum tracking speed.

The maximum miss-track is proportional to the tracker speed, at UdeS Capteur (big dish) tracking speed (not fast at ~13 minutes for 180º), the maximum miss-pointing would be only ~0.5º. The UdeS capteur would start tracking at its top azimuth speed at 6.6 minutes before solar noon, and at N-5.5 (5.5 minutes before solar noon) it would get far enough ahead to start missing rays, missing

~10% at N-4.5 and ~15% at N-3.5, at which point the closeness to a straight overhead sun starts to decrease the effect to ~10% at N-2.5 and a few % at N-1.5.

A similar pattern would occur in the minutes just after noon, so the system would lose a total of roughly one minute of output spread over the ten minutes around noon. Since this would only occur on those days where the sun passes almost straight overhead, this predictable loss of a minute of output should be acceptable (and it can be eliminated with a faster tracker if needed - the UdeS Capteur has a home-built tracker and was not designed for speed).

3.4.9 Evening-out the Focal Intensity

A trough naturally has an even focus along the trough, and the centimeter-scale variations from mirror imperfections are averaged out by the six rows of cells per receiver. Across the trough the focus is naturally somewhat uneven, and a generous tolerance budget actually increases the unevenness by adding typically unilluminated lens width. The focus is still much more even than two-axis Fresnel CPV is without a homogenizing secondary optical element, so it won't have hot spots that endanger the cells, but a more-even focus typically makes a module more efficient. The effect is not huge because some areas get more efficient as others get less efficient, and the effect can be reversed in less-intense light, but it is significant enough to address.

There are several ways to even the focus, none of which will be used in early prototyping. Once the actual focus is profiled and the initial acceptance angle budget is confirmed, the focal evenness will be optimized. These are discussed in the order that they fit into the optical path.

- ***Mirror Modifications:***

The focus from the middle of the mirror is tightest, and thus has excess tolerance budget. The mirror shape can be modified to slightly flatten the middle of the mirror to direct some light from the center of the mirror away from the center of the lens. If the system is scaled to two mirrors wide, this can be done simply by slightly slanting the mirror segments outward relative to each other. Alternatively, the mirror shape can be adjusted, either with a custom mandrel or with the proprietary folding shapers. Any of these would be practical at roughly the 100 MW level.

- ***Lens Curved in Linear Direction:***

While a purely linear lens is the easiest to roll-form, this is a nicety rather than a necessity, and most roll-formed shapes have a 2D pattern.

- The light from the far edge of the trough spreads more, so one end of the lens only receives light from the far edge of the trough. If this end of the lens is slanted down, it compresses that light, which is the weakest part of the focus anyway, away from that end of the cone row. This has not been modelled for ray tracing, but it looks like the cone row could be shrunk by about 5%. This would increase the average concentration as well as decrease the variance in concentration, and the higher average could be spent to raise acceptance angles or even widen the cells slightly – probably 15 cells per row, each 0.91 mm x 0.68 mm.
- The top of the lens can be depressed in the center (across the trough) to divert light from the more intense center of the row toward the less intense ends. This is less powerful than modifying the mirrors because it moves the light less far, but it could be implemented at lower volume (even 10 MW).

The focal length of the lens surface would be adjusted slightly to compensate for the thinner regions of the lens produced.

- ***Different Cone Widths:***

The cones on the ends of the row can also have larger mouths to catch more light while the ones in the more-intense middle of the row have smaller mouths. To avoid cones much taller, this can be done by allowing the light to hit the cell at a flatter angle on the end cones; the slight loss of efficiency for this is made up for by a corresponding increase in efficiency on the narrower central cones.

The ends of the lens row can also be curved down, while this diverts light away from the least-intense area, it diverts it to an under-illuminated area, and it reduces the angular range seen by the end cones enough to allow them to be much broader and thus capture more light.

- ***Utilizing Uneven Intensity:***

There are also several ways to tolerate or even benefit from uneven intensity. If the cells are performance-binned during testing, then the best-performing cells can be used in the central region where they will receive the most of light, and the marginal cells can be used at the margins of the

focus. This increases the average efficiency, increases cell yield by allowing lower-performing cells to be used in low-priority locations, or a combination of these.

Finally, stacked-junction cells with multiple copies of each junction would increase efficiency at higher concentrations, taking maximum advantage of the more-intense central illumination. The Université de Sherbrooke is a pioneer of such cells, which would be an excellent fit for TLC.

- ***Thermal Runaway from Uneven Focal Intensity:***

A Fraunhofer paper (M. Steiner, G. Siefer, and A. W. Bett, 2013) detailed how fully illuminated cells can gang up on under-illuminated cells and drive enough current through them to trigger thermal runaway, but stated that the susceptibility depends upon the cell cooling. This has been analyzed and tested; QDEC microcells have been found to be resistant to this, and experiment has confirmed that TLC with QDEC cells is not susceptible under expected operating conditions. A paper on this was published in the CPV-16 conference proceedings (Norman, et al., 2020).

However, repeating the experiment with AZUR 3C44 cells showed that these have much less safety margin than QDEC cells. On-sun testing with 3C44 cells is planned to verify whether the safety margin is adequate for 3C44 cells in an operating system.

- ***Summary on Focal Intensity Evenness:***

Having the cells of a row in parallel means that the focal intensity across the trough does not have to be even, and several approaches can even take advantage of this. The approaches are generally cooperative, and the optimal blend of approaches to focal evenness across the trough's focus will be evaluated once the actual focus is profiled and the tolerance budgets are known. Until then, however, this remains a significant source of uncertainty; a ball-park estimate is that if no approaches are used, the natural unevenness will reduce power output by very roughly 5%, but that mirror shaping, and performance binning could turn this into a net gain and stacked-junction cells plus binning could turn this into up to a 5% boost.

3.5 Design Scaling

3.5.1 Tracker

- ***Tracker Length***

Trough length, and thus tracker length, is freely scalable in DAT TLC; in SAT TLC end losses (<1% at 100 meters) grow on shorter troughs but become arbitrarily small with very long troughs. If the low-grade heat is desired as a co-product (TLC-T), then the coolant flow scales much like a CSP trough in that twice the length requires twice the coolant pumped twice as far. However, in TLC-T the pipes are low temperature and low pressure, and the off-axis focus lets large distribution pipes be below the modules where they are not limited in diameter and can use ordinary glass wool rather than glass/vacuum insulation, so the cost of scaling is still low.

- ***Tracker width:***

Tracker width must be an integral number of half-trough widths but can otherwise scale upward or downward independently based on tracker economics. Single-axis flat panel trackers 3.3 meters (two common solar panel lengths) wide seems to be standard, and 3.3 meters is a near-perfect match for two inner-RP-3 half-troughs and hence are a good choice for SAT TLC.

- ***Heliostats for DAT TLC:***

Heliostats for CSP Power Towers are bringing two-axis tracking costs down rapidly. While not a technical scaling issue, these are trending to ~150 m^2 so that is an economic scaling sweet spot.

3.5.2 Trough Rim Angle (Focal Length/Width)

A lower rim angle does not allow a trough to concentrate as much due to a longer focal length F still spreading the sun's image, so the receiver width grows proportionately (for small RA) to F and the primary concentration scales by ~1/F. However, a lower rim angle allows the lens tile to concentrate more on both axes, so the maximum lens concentration grows by roughly ~F^2 and the achievable concentration grows by ~F. This holds until the Petzval effects become insignificant and sun's diameter and system imperfections limit further concentration.

The cells are by far the most expensive component and their cost is directly affected by concentration, so up to the limit of cell's ability to use higher concentration, cell cost drives rim angle lower and thus focal length longer. However, going much beyond 1500x could damage today's cells, and the cell efficiency is already dropping noticeably by that concentration, so with current cells higher concentration (beyond 1500X) is harmful rather than helpful.

Once the cells' concentration limit is reached, a lower rim angle concentrates less initially, requiring wider lenses, and the higher secondary concentration produces narrower cells, and the cells are so narrow that further narrowing would both increase kerf losses and mechanically require short (and thus smaller) cells. Therefore, further lowering the rim angle below a single RP-3 inner mirror is thought to be counter-productive until cells can survive higher concentration (and by then cell costs may be low enough that higher concentration is simply not necessary).

3.5.3 Trough Width and Receiver Size

In most places trough-width scaling has minimal effect. Module supports get longer as well as wider and thicker so steel per watt increases, but this cost is small. Receiver handling costs drop slightly with wider troughs, but with DAT TLC this is not significant. For SAT TLC most micro-tracking costs and plumbing connection costs are cut in half at twice the trough width, so SAT TLC can benefit greatly from custom mirrors.

However, the simplest cooling for DAT TCL, pure conductive cooling, shows a strong scaling effect. Wider troughs require larger fins which require thicker aluminum to conduct heat farther, and larger fins also need larger gaps for air-flow, requiring even larger fins to make up the lost area from fewer fins. Scaling the fins 2x in each dimension provides $2^{3/2}$ more heat rejection while consuming 2^3 more aluminum, so conductive cooling reaches a limit and other cooling is needed. However, a single RP-3 mirror, while wider than optimum for the cooling, is still within the range where pure conduction should offer the lowest cost overall.

- ***Receiver Width:***

If the trough width and receiver width are scaled, but not the receiver length, then the length of the interconnection between receiver does not increase so a receiver W times as wide has W times the

current and W times the power. For a DAT receiver, the conductor is also W times as wide, so nothing electrical changes except that W times as many diodes (or larger diodes) are used.

With heat-pipe cooling, however, both the amount of coolant to wick and the height that it has to wick grow with receiver width. But since custom mirrors would be used for scaling the width, these could be made enough off-axis to allow covering the receiver width with coolant and thus not need wicking for coolant transport, preventing the scaling bottleneck.

For SAT if the *area* of the interconnect wire is scale by L (diameter by $L^{0.5}$) then the wire mass scales by L and the resistance by 1/L. Thus, the voltage drop and the efficiency loss remain constant and the wire mass per Watt remains constant. There is a catch – the length of the wire used to connect the cells increases, so wire cost scales with L and its resistance remains constant and its efficiency loss grows with L^2, but since with the RP3-inner mirror that cost is ~35% of the total wire cost and is only ~14% of the resistance-based efficiency loss, doubling the wire cost and quadrupling the in-receiver resistance loss is a small price to pay for reducing other costs.

- ***<u>Receiver Length:</u>***

Scaling the receiver length is very limited because currents rise and voltage per length drops with longer receivers, and diode and placement costs go up with shorter receivers. With DAT the added costs of shorter receivers (scaling down) are very small (less efficient use of common diodes, and slightly higher handling costs), so if anything, receivers would get shorter.

With SAT the receiver length is constrained between being wide enough to micro-track and narrow enough to not have part too out of focus with RP3 mirrors. However longer focal-length troughs would allow longer receivers. An equatorial receiver could be roughly twice as long because the receiver's maximum tilt is 23º instead of 60º (***sin***(23º) = ~1/2 ***sin***(60º)), but even by the start of temperate latitudes little more could be gained (***sin*** (45) is ~85% of ***sin***(60)). Since the savings are not large, and these same savings can be gained for all latitudes by scaling the trough width, a special limited-latitude version is thought not to be worthwhile.

- ***<u>Receiver Length and Width:</u>***

If everything is scaled linearly by a factor **S**, the voltage of a receiver remains the same while the Watts, Current and wire area scale with $\mathbf{S^2}$. The wire length scales with S, so the wire's resistance scale with **1/S** (length/area = $\mathbf{S/S^2 = 1/S}$). The voltage drops, and thus the percentage of power lost,

scales with IR and thus $\mathbf{S^2/S}$ or **S**. The mass of the wire (metal cost) scales with S^3, so the wire-per-Watt scales with $\mathbf{S^3/S^2 = S}$. The two costs (efficiency loss and wire cost) that increase linearly with scale can be traded off – at the same wire cost efficiency loss scales with S^2, or vice versa.

The Receiver Length has reached diminishing returns already, so doubling both the receiver length and width appears to go beyond optimal. However, with SAT, doubling the receiver width appears to be near optimal. DAT TLC seems close to optimal already, but with higher-voltage cells the current might be below the optimum point for the most cost-effective diode size.

3.5.4 Cell and Lens Row Width

Having multiple lens rows decouples lens width and cell size from receiver and trough scaling. If the size of the individual lens rows is kept the same (changing the number of lens row with the lens length), then the lens cost is invariant to scale, as are the cell and cell-attach costs.

- ***Cells:***

Larger cells reduce cost due to lower kerf losses (but a 4x larger cell only saves ~5% of wafer area), however these are offset by higher optical losses and glass costs from the thicker lenses, the cells get slightly warmer, and cell yield per wafer decreases.

- ***Lens:***

The limits on lens row width are the glass and the cell width. Wider lenses rows would allow wider cells (good up to a point) but would (at the same concentration) require thicker lenses that would use more glass (more cost) and absorb more light (unless clearer glass is used, which costs even more). Narrower lenses would reduce lens cost and losses, but also reduce cell width.

The optimum lens width is dependent on the glass clarity and inversely dependent on the glass cost. With today's triple-junction tandem cells, the extra bottom-junction photo-current fortuitously allows using a low-cost low-iron solar glass lens; future higher-efficiency tandem cells may not be so wasteful of photons and so may require a more-expensive lens glass. However, in volume the next grade of glass becomes affordable at about the same lens width, so in spite of changing glass and cells the lens width is likely to remain constant.

- ***Air-gap Lens:***

An air-gap lens greatly reduces the lens scaling cost and optical loss, allowing wider lenses and cells, especially on larger troughs where other limits are postponed. However, air-gap lenses need better AR coatings to become cost-effective, and airgap lenses create alignment challenges between the lens tile on the cones.

3.5.5 Scaling Down to Smaller Installations

TLC can scale down to a single trough with no changes and no loss of efficiency, and with the main economic penalties being a higher installation cost due to not being able to install many megawatts on a single site, and higher inverter and permitting costs per Watt. However other PV faces similar penalties, so this is not a limitation. A minimum tracker is ~6m^2, or 3 mirrors (~3 kW.

A single half-width trough is also simple: the dual-off-axis design is two halves stuck together anyway, so while the cost per watt goes up slightly due to not sharing the supports, this is minor.

- ***Scaling to Slanted Rooftops:***

A half-meter focal length puts SAT TLC into the range of south-facing sloped roofs, especially when the heat by-product can be useful for space heating. And southerly slants (or northerly slants for southern latitudes) are similar to being at lower latitude in that trough end losses and micro-tracking range are reduced, so slanted rooftops are a possibility for SAT.

3.6 Conclusion

This chapter has walked through the major aspects of the TLC optical design and how it produces optics that have a minimum material cost. This chapter then detailed the three optical stages and showed that the TLC optics are plausibly mass-producible at low process cost. The chapter has also covered the main trade-offs in optimizing the TLC optics for high optical efficiency and low cost, and shows that although the focus has been on cost, the design should be very efficient due to the optically coupled 2nd and 3rd stages and high transparency in the key CPV wavelength range.

TLC's three-stage optics produce a dense array of small focal spots, a novel configuration which leaves the challenge of complementing the optics with high-efficiency, low-cost receivers, and of sealing and cooling the module. These challenges are addressed in the next chapter.

CHAPTER 4

4. ASSEMBLY PROCESS & MATERIAL COSTS

4.1 Introduction

Chapter 4 covers the design of the rest of a TLC module to match the TLC optics. Chapter 4 starts with the microcell arrays proposed for TLC module and details their electrical connections by way of a heat spreader that also serves as an electrical backplane for the module, placing the cells in a novel massively parallel microcell array configuration. This chapter then details several methods for making the heat-rejection housing that protects and cools the rest of the TLC module.

This chapter then walks through the sub-assemblies of a TLC module and the mechanical support proposed for a TLC module, and then outlines the dual axis tracker considerations for supporting TLC modules, and then the shipping and installation of the modules, showing a plausible path to TLC modules that not only have a low materials cost, but that can be easily assembled, shipped and installed as well. (The models and validations this is based on are presented in Chapter 6.)

4.2 Cells & Electrical Interconnections

The lens tile focuses partially concentrated light from the trough onto an array of cells, with each cone focusing onto its own tiny cell. This section covers the cells and their electrical connections.

4.2.1 Cell Type

While high concentration makes ultra-efficient tandem cells affordable, any high-efficiency cells could be used. A 36%-efficient cell would have to be free to be worth using given today's prices for 44% AZUR 3C44 cells; conversely a cell would have to have efficiency over 50% to be worth twice the price of the 44% cells. While the early prototyping has been done with ~40% QDEC cells, the design is currently optimized for commercially available AZUR 3C44 cells and the efficiency and cost calculations are based on those cells (both QDEC and AZUR 3C44 cells are germanium-based triple-junction tandem cells).

4.2.2 Cell Size

A receiver has multiple lens/cell rows, with many (currently 16) high-efficiency solar cells in each row. To keep the lens cost low, lenses are kept small enough to use the best commodity glass with acceptable (a few percent) absorption losses, rather than requiring expensive instrument-grade glass. Small lenses that concentrate ~15x on the second axis produce very narrow foci, which means very narrow cells. The optimum with single-part commodity-glass lenses is 0.6 mm to 1 mm cells (Norman, et al., 2019) (currently the target cell size is 0.648 mm wide). That smaller cells and higher concentration go together is confirmed by SunLab (SUNLAB, 2014).

The cell length is limited by the aspect ratio (handling), the TIR cone height (thermal), and keeping CTE-mismatch stress low, and the aspect ratio is limited by capturing light at 1500X without heroic manufacturing efforts. Cells 0.9 mm long allow avoiding light at grazing angles onto the cells while keeping the TIR cone height modest. The TIR cones for 16 such cells can span the focus of the trough, so the current design has sixteen 0.9 mm cells per row, and a cell diagonal of 1.11 m (smaller than a 1-mm LED's 1.41 mm).

This gives each cell an area of only 0.58 mm^2. While this seems tiny and likely to create far too many cells to handle, for a 38%/1500x system such a cell produces ~0.33W. This is only ~10% of what a top-of-the-line SunPower Maxeon Gen II silicon cell generates, but the TLC cells can be handled by high-speed pick-and-place machinery that can place tens of thousands of cells per hour, while the fragile, far-bigger silicon cells require special handling by far more expensive machines that operate at one tenth the speed, so even with such small cells TLC placement cost is *much* *less* *per watt* than the placement cost of silicon wafers in a flat panel.

To keep the size in perspective, these cells are in the middle of the common size range for passive components for electronic equipment, and **over 3 *trillion*** passive components are placed per year (Zogbi, 2016), on similar-sized substrates, by low-cost automated pick-and-place equipment.

- ***Central-pad Cells:***

The typical two wide bus bars consume too much of a small cell, but TLC could use a narrow cross-cell central bus bar with a single small bonding pad inset from the one cell edge (a second pad is added to remove an orientation constraint). The central bus bar cuts the distance travelled in gridlines, so electrically a 0.9 mm cell looks like a 0.45 mm cell. This is lets narrow gridlines be

used for less shadow loss, increasing cell efficiency (on a longer cell this could be extended if needed – two contacts at ¼ and ¾ of the length would cut the grid length in half again, etc.).

Reverse-ball-bonding puts the lower bond on the die, and this can achieve heights below 63 µm (Chylak, et al., 2006), and CPV optical coupling silicone is often 100 µm, so reverse bonds could fit under the lens.

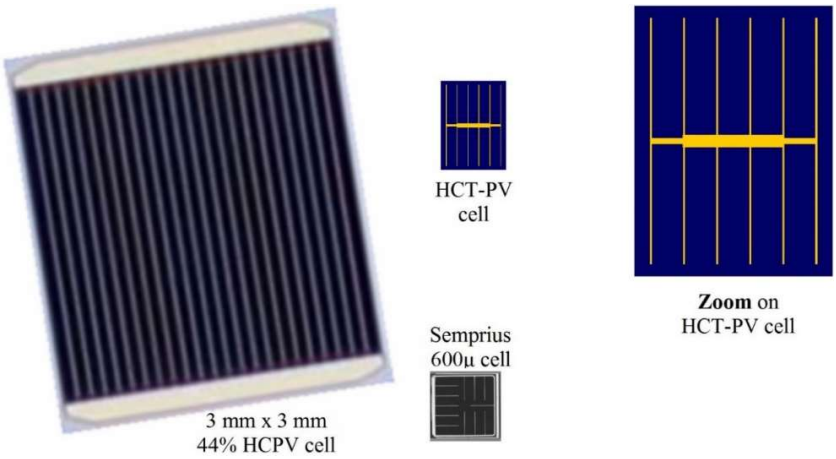

(a) (b)

Figure 4.1: *Cell Comparison: (a) Azur 3 mm, Semprius, TLC cells ~20x scale, (b) TLC cell detail at ~60x.*

- ***<u>Efficiency Comparison to Current Azur Cells:</u>***

  - *Grid Shading:*

From counting pixels on Azur's grid layout drawing, the Azur grid fingers appear to be 16.8 um wide shade and thus shade 12.34% of the cell, and the dual bus bars add to the die area but do not consume receptive area. A 3 mm x 3 mm cell thus has a receptive area of 7.89 mm^2, and with 22 bi-directional fingers, each Azur finger serves 1/46 of that, or 0.1715 mm^2 of receptive area.

With ~10% shading TLC's 0.58 mm^2 cell has an active area of ~0.53 mm^2. With 6 bi-directional fingers, each TLC finger serves roughly 0.044 mm^2 and at ~1.5x the concentration this will have a current roughly equal to 0.66 mm^2 of Azur's 3 mm cell at 1000X, or 38% of the area served, so a first-order optimization is to make them 38% as wide. (Keeping widths proportional to the Azur grid reduces uncertainty until the Azur grid is accurately determined).

TLC's bus bar end segments carry two fingers worth of current their whole length, and the optimum seems to be 3x the TLC finger width. The next 3 segments of the bus bar all appear to have an optimum width of 6 finger widths, so they are merged into one long pad that one can land a bond wire anywhere on. In total this combination shades only 8.88% of the cell, improving performance by **3.95%** from reduced shading loss.

- *Grid Resistance:*

An Azur grid finger serves 0.1715 mm^2 of cell area and is 1.5 mm long. Each TLC grid finger serves only 0.044 mm^2 so even at 139.5 W/cm^2 (the intensity at STC given the optical efficiency in the current-limiting wavelengths) vs 100 W/cm^2 (which the closest comparable Azur cell rating data is at) it has 36% of the current, so at 38% of the width its voltage drop per length is almost identical. It is 28% of the length and so has only 28% of the voltage drop.

Adding up the bus bar resistances that each finger's current sees raise the average resistance to only 59%, or a 41% reduction. The savings from this depends on the grid losses. Since the grid resistance is generally balanced to half the grid loss, this is ~41% of ~7% for a roughly **2.9%** increase in efficiency from reduced grid resistance.

- *Alternate Grid Optimization:*

Work on the cell grid in 2018 showed that although a second bond wire slightly reduce the cost per watt of a module it decreased the LCOE due to the higher efficiency from lower resistance. Several 2-bondwire optimizations were created. The lowest resistance of these is the asymmetrical menorah configuration, shown below (having the bond wires in opposite corners may be more convenient and offers almost as low resistance):

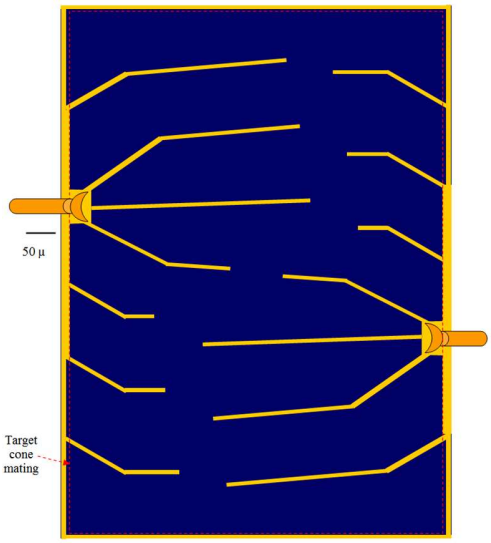

Figure 4.2: *Two bond-wire cell optimization.*

With the bond wires removed from the middle of the cell, the optical coupling no longer needs to be thick enough to cover the bondwire. This eliminates the issue of optical coupling control which had been identified as a potential 2% loss for uncontrolled optical coupling 60 µm thick. Instead, the cones can simply be pressed against the cells, either with a tiny amount (submicron thickness) of adhesive optical silicone, or simply left to couple on their own through pressure. The bondwire is so small compared to a cone that the soft silicone will simply conform around it, as shown below:

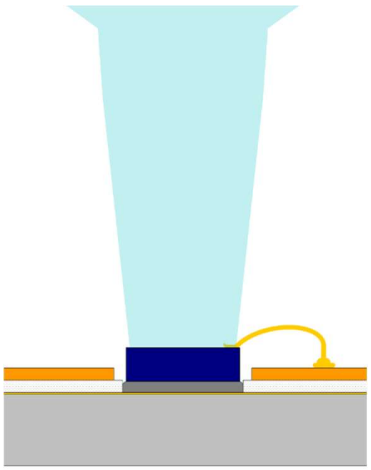

Figure 4.3: *Compressed Cone Coupling.*

- *Sheet Resistance:*

The TLC cell grid fingers are ~ SQRT (1000/1500) times as far apart as the fingers in the Azur cell, so the voltage drop from the sheet resistance in the TLC cell at 1500X is almost identical to the Azur cell at 1000X (1.068 time as great). From previous cells this is often ~1/4 of the grid resistance, so this slight change in a minor factor (highly pessimistic), this costs ~~ **-0.04%**.

- *Tunnel Junctions:*

Data for the tunnel junction resistance in the Azur cell has not been obtained, but Boeing Spectrolab's 5J cell work indicates that resistances below 5 mOhm*cm^2 are readily achievable, so that value is used for 10 mOhm*cm^2 for two tunnel junctions. At 1000x the Azur cell is 1.39A, or 15.4 A/cm^2, so the tunnel junction voltage drop is 154 mV, and at 139.5 W/cm^2 it will be 61 mV higher. 61 mV is 2.15% of the V_{MPP} of 2.82V, so this decreases the efficiency by **-2.1%.**

- *Perimeter Recombination:*

Recombination losses at the cell perimeter start to become significant on very small cells (Fidaner and Wiemer, 2013), but 1500x concentration delays this effect. The impact on V_{OC} has been determined for Solar Junction's 44% cell (Fidaner and Wiemer, 2013). The perimeter to area ratio of a TLC cell is the same as for a square 0.78 mm cell, so reference lines for a 3 mm square cell and a 0.84 mm cell have been added to Fidaner's graph (**Figure 4.4**).

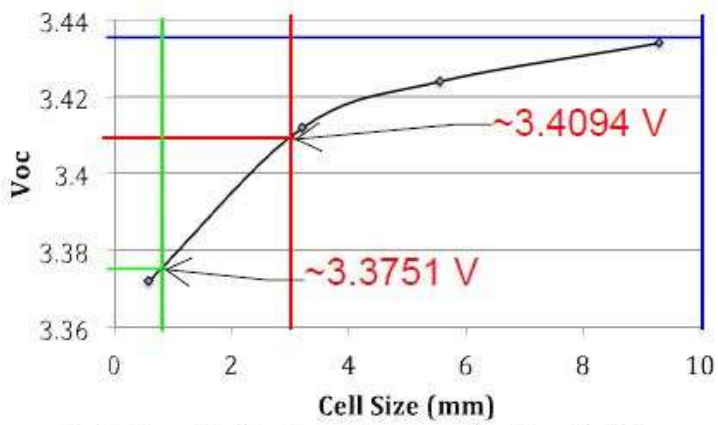

Fig. 5. Open circuit voltage change as a function of cell size.

Figure 4.4: *Cell Size Impact on V_{OC} (Fidaner).*

The percentage impact on V_{MPP} should be very similar to the impact on V_{OC}, and the difference in V_{OC} is 34 mV, or a 1.01% reduction, so increased perimeter recombination from the small cells should reduce the performance by ~1.01%. This is pessimistic because concentration decreases the perimeter impact, and while Fidaner does not mention the concentration that the V_{OC} graph is at, the highest concentration that any labelled data in the paper is at is 120 W/cm^2 (1200X) which is less than TLC's 139.5 W/cm^2.

- *Net Change in Efficiency:*

Taking these effects into consideration, under standard test conditions the small TLC cell at 1500x should gain 3.68% (relative) efficiency compared to the Azur cell at 1000 suns. The Azur cell is 42.9% at 1000 suns, so if made with the current Azur process, the TLC cell efficiency should be just slightly better than the Azur cell at 1000x, or 44.5% at a 25ºC cell temperature. The Azur cells' optical coupling loss of 0.4% reduces this to 44.3%.

- *Edge Fingers:*

Fingers could be put on the edges of the cell, just outside of the active area. If the finger spacing and width is kept the same, this would put only five fingers on the cell instead of six. Each finger shades 1% of the cell, and the tiny extra bus bar would be offset by the lower resistance for 1/6 of the cell, so output would be increased by ~1%. If the finger can be right on the edge, then this is ~13 extra microns of cell width, or 2.3% more cell; at current costs this is ~1% of the module cost. While the module cost per Watt would thus not change, the more efficient modules would reduce tracker, installation and maintenance costs by ~1% as well for slightly lower LCOE (and if the edge fingers can be narrower than normal fingers, the advantage slightly increases). However, if more than 10 µm of extra cell is needed due to dicing tolerance, then the extra finger would not be worth the extra cell area. Although etched streets should allow fingers right at the edge, until this is confirmed this efficiency boost is not included (it is shown in the 'menorah' cell grid).

- *Above-Rated-Efficiency Cells:*

How can a cell possible work better at higher concentration than the cell itself is rated for at its peak efficiency? Small cells raise a cell's efficiency markedly even to the point of beating the world record! This is because the 'World Record' for concentrator cells is restricted to cells of at least 5 mm^2 (Green, et al., 2016) to keep too-small-to-be-practical cells from beating out useful cells. This

has made sense because the few companies that use cells smaller than TLC cells spread them across the entire aperture area, which involves substrates too big to use low-cost pick-and-place equipment. However, the ~40X area initial area reduction afforded by the trough lets TLC use multi-cell substrates small enough to fit in high-speed pick-and-place and bonding equipment, which makes small cells very affordable, and thus allows using cells whose efficiency exceeds "the best cell's efficiency".

- ***Future Cell Contacts:***

An insulated Through-Cell via (TCV) could map the cell's top contact to the cell back, further reducing the contact area and eliminating the bond wire, and thus saving much of the grid shading and resistance losses. This currently would raise cell cost more than the cost of the cell area saved, but that might change at high volume with thin cell wafers (e.g., from substrate reuse).

- ***Variability:***

Small cells have higher variability because performance variations do not tend to average out within a cell. However, with many tens of small cells in parallel there will very low overall performance variation because cells will come from different wafer regions or even different wafers, reducing the receiver variation by an order of magnitude. Rather than creating a variability issue, small cells in TLC *reduce* variability (yet another advantage for TLC).

Other than removing 'dead' outliers with performance too low to be worth the space they would take, performance-binning could be skipped. However, binning can still be useful; cells on the very ends of receiver rows mostly catch rays from the edge of the sun hitting a misaligned mirror, and lower performance cells could be used here instead of throwing them out. This could increase the usable yield from a wafer by another few % without using cells too far out of spec.

- ***Cell temperature and DC Efficiency Estimates:***

Cell efficiency is key because efficiency increases reduce other costs (except inverters and AC connections) proportionately. Standard module test conditions (CSOC) are 20ºC ambient and 2 m/s wind. Wind is great on the large fin surface so cells should be around 58ºC under such conditions, costing ~4% of the output or ~1.6%$_{abs}$, so cell efficiency under test conditions should be right around 42.3%. With the 92.2% module optical efficiency estimate for the key wavelengths, that

gives cell+ optics module efficiency of 39.0%. Non-cell electrical losses are a tiny ~0.5%$_{rel}$, so the estimated DC efficiency is 38.8% with today's cells.

4.2.3 Future Cells

Advanced cells are ***not*** needed for TLC to beat silicon, so there is no chicken-and-egg problem. TLC beats silicon on LCOE even with ancient cells that have been available from multiple companies for many years. Cost and margin calculations for TLC introduction are therefore done with today's cells to provide an incontestable floor[1] until advanced cells are available. However, it is useful to know what is likely to be available to CPV, so cells expected at introduction in 2021 and in volume production in 2025 are discussed. Note that the cost model **conservatively assumes no CPV cell efficiency progress whatsoever.**

- ***Cell Expected at Introduction:***

Cells exceeding 46% efficiency were demonstrated by Soitec in 2014, and commercial cell efficiencies typically lag lab cells by about 3 years (but with silicon prices having far undercut traditional CPV, that has slowed). Boeing Spectrolab also has a 5J space cell that would be 46.7% if converted to a concentrator cell, so such cells in production in 2021 seemed reasonable and matched NREL/Fraunhofer forecasts (Wiesenfarth, 2017).

The Boeing Spectrolab 5J cell was expected to be the best cell available at introduction, so the available data on it has been looked at more deeply than other future cells. Spectrolab's analysis of a 1 cm^2 cell anticipates a peak efficiency of 46.7%, which is the base figure being corrected for size, concentration and temperature in the performance estimates.

With five junctions already reducing the current, small cells don't help as much due to the lower current creating a lower voltage drop and the higher voltage making that lower drop relatively even smaller, but higher concentration doesn't hurt as much either. Without a layout showing the

---

[1]While some CPV skeptics have tried to argue that even having CPV tandem cells at all can't be counted on, Cyrium showed that any LED fab can make 40% cells, so unless LEDs disappear the ability to make 40% cells can be relied on.

shading, top contact effects (IR drop, shading, and sheet resistance loss) can be approximated through scaling, assuming that same top conductor material and thickness is used. The tunnel junction loss is increased by the concentration and the number of junctions more than it is decreased by the higher voltage and resulting lower current per Watt. The perimeter loss of V_{OC} is slightly greater due to the larger reference cell, but this is balance by the higher V_{OC}.

The net effect is that under test conditions (cell at 25ºC) the TLC cell efficiency at 1500x (139.5 W/cm^2 after optical losses) is expected to be almost identical (~~ 0.14% lower) than a 1 cm^2 cell at 500 suns, so with the 5J cell the cell efficiency will decline from 46.7% to 46.6%. The temperature sensitivity is not known, but if it matches the Azur cell then the 5J cell would be 44.8% efficient under module test conditions, which with the optical efficiency for the three critical junctions and the even lower resistive losses in the wiring suggests a module efficiency of ~~41.0%. This is only 1.5% abs or 4% relative higher than with the 44% 3J cells.

- ***<u>Cells at 1 GW Production:</u>***

Tandem CPV cell efficiency has historically increased 0.8% (absolute) per year. While commercial cell progress will be slower until a viable CPV product is launched, once any HCPV system is commercially competitive with silicon then CPV cell progress will accelerate because the R&D cost of a percentage point of improvement in CPV is roughly two orders of magnitude lower than a comparable increase in silicon module efficiency. By 2025 this could make ~50% cells commercially available (that's *six years* after an NREL team working on them estimates that they will be in the lab), so 50% cells are likely by the time TLC reaches at 1 GW. The additional gain would not bring more tunnel junctions, so this would produce a ~~44% module.

- ***<u>Eventual Efficiency:</u>***

Once there is a major market for terrestrial tandem cells for CPV, money should flow to cell R&D. However, it will be hard to break the 55% cell barrier because adding junctions brings limited returns, especially at higher concentration. A 55% 5J or 6J cell is likely to be a practical limit, which would produce a ~~48% module with the current optics. An improved AR coating could add almost a percentage point, and clearer glass could also add, but it is going to be difficult to break the 50% module efficiency barrier in a commercial module (however with first-surface dielectric mirrors, moth-eye ARC coatings, instrument grade glass and active cooling, a 'hero' module should be able to exceed 50% if made with 55% cells).

4.2.4 Cell Cost

NREL shows a CPV *cell cost* for a 30% efficient 1000x module at 18¢/W_{peak} (Horowitz, et al., 2015), or \$5.40/cm^2. Almost all of that cost is the substrate and the metallizations, so it is largely independent of the number of junctions and the resulting cell efficiency.

- ***Cost Reduction from High Concentration:***

Sarah Kurtz of NREL once said "If you are concentrating more than 500X, ask yourself why". While she was absolutely correct when CPV was then competing with \$4/W silicon PV, silicon panels have fallen more than 10x since then so CPV cost must also drop dramatically to compete, and at 500X (and \$5.40/cm^2) *the cells alone would cost more than the entire 1500x TLC module, so 500X is no longer affordable unless CPV cell costs fall*!

Even single-axis-tracker TLC could reach very high concentration (~1200x) in a system optimized for 0º-37º latitude, in spite of multi-stage optics that incur Petzval effects, because it is based on linear technologies that are easy to make accurate. Thus, rather than optimizing the optics to *produce* a high acceptance angle, the mechanics and optics are optimized to *minimize the consumption* of acceptance angle.

TLC for dual axis trackers eliminates the 'tilt' broadening and 'smile' aberration, allowing roughly 3x higher concentration (up to ~5000x). Current cells generally lose efficiency significantly after 1500x (but small cell less so), so this extra capability is spent on using an off-the-shelf CSP mirror, on acceptance angle and on manufacturing tolerance, but if cells could take higher concentration while delivering the same efficiency, it could shave >10% off of TLC's total costs at 1 GW production.

- ***Azur Pricing:***

CPV-13 provided an opportunity to talk to Azur on pricing for their cells. At 1000 wafers, or 75,000 cm^2 or ~5 MW, diced wafers of tested Azur cells on blue tape cost €220 to €240, or \$245 to \$265 at the early-2020 exchange rate. \$260 would be \$3.50/cm^2, which yield and kerf losses would raise toward the NREL estimate (but Azur is still below the NREL cost estimates).

- *Kerf Losses and Yield:*

Smaller cells produce more good-cell area per wafer. Not only do the cells fit the wafer edges better, but yield goes up due to throwing out less cell area for each defect. For 5 mm x 5 mm cells

the NREL cost study lists an 80% yield; much of that would be random spot losses so 43x smaller cells would increase yield significantly. Papers on the defect distribution are hard to find, but if half of the loss is random defects, then the small cells would increase yield per wafer by roughly 10%. This matches information from an ex-cell supplier, so 90% yield is assumed.

Azure has a 70 µm kerf, and an on-line die calculator says 10305 cell-sized dies per wafer (assuming a 1.6 mm edge keep-out). Smaller cells have higher yield, and at a 90% yield that is 9275 cells per wafer, and thus 3.09 cents per cell, or roughly 9¢/W (depending on the Euro).

- *Cell Cost Reduction with Cell Volume:*

By 50K wafers the price is EU 165 to EU 170, or $185 to $190 at current exchange rates, or $2.70/cm^2, or ~6 ¢/W. This current pricing at 250 MW is reduced by 5% (adjustable) per doubling for future 1 GW cost estimates, which is pessimistic – if semiconductors in solar have a lesson to teach, it is that costs not only fall with volume and time but fall faster than expected.

NREL's bottom-up analysis shows that with current cell manufacturing technology, cell prices should be $2.10/cm^2 (after kerf and yield losses) at a production level equivalent to what silicon cells are made at today. Other costs will have also decreased, so the cells are still expected to be 30% of the system cost at that time, so TLC will still encourage higher-concentration cells.

- *Future Cell Cost Reduction:*

NREL's analysis also shows that with known technology such as 5x substrate reuse (which has been demonstrated but not yet commercialized), the price would fall to <60¢/cm^2. At that price, backing off to 1200x or even 1000x concentration *might* become attractive since *current* cells are more efficient at 1000x than at 1500x, and 1000x would only raise the cell cost by 0.5¢/W.

- *Testing Cost:*

Small cells mean more cells to test. The NREL CPV cost study shows 'Edge isolation, Test and Sort' as ~4.5 ¢/W at 1X, or 0.133% of the $33.86/W epi wafer cost. The study shows epi wafers at $11,200/m^2, and 0.133% of that is $14.88/m^2 for cell testing. The study uses 5 x 5 mm^2 cells; with the typical 250 µm bus bars that is 36,000 cells/m^2, or 36,000 cells for 1488¢, or ~0.04 ¢/cell. This has significant relative uncertainty from measuring such a short column on the NREL chart, but at the current cell size it is only 0.12¢/W so the absolute uncertainty is not significant.

This is pessimistic because testing 50 times as many cells should not cost 50 times as much at 100 MW, and even 10 MW is 5x more cells. At high volume it is even more pessimistic because a single contact point per cell should make testing a whole row at once with a multi-point probe practical. However, it is a referenceable cost, and it is a low enough cost that there is no point in refining it further, so it is used as is to be conservative.

- ***Dicing Cost:***

Small cells do have higher dicing cost when sawn or even when separated with TLS (thermal laser separation). However, TLS is very low cost. For a 150 mm SiC wafer of ~2 mm dies, the total cost was calculated to be only $2.70 (Zuehlke and Gebhardt, 2016), or 2.6¢/cm^2. The much smaller TLC cells have ~3 times the perimeter per area, so that would be 8.3¢/cm^2, or ~0.14¢/W, an increase of ~0.13 ¢/W above the cost for 5 mm square cells.

- ***Conclusion on Cell Size versus Cost:***

TLC's small cells pack better on a wafer and have better yield, and are more tolerant of variability, and so produce more Watts per wafer. On the other hand, dicing and testing costs do go up, but only by ~**0.25¢/W** compared to 5 mm x 5 mm cells. Thus, *small cells are highly affordable.*

4.2.5 Cell Attach Cost

- ***Cell Placement Cost:***

Even a small CPV cell costs an order of magnitude less per Watt for placement than placing silicon cells in a flat panel. A high-end $120K machine can places about 120,000 chips per hour, and even a mid-range machine places over 10,000 chips/hour with ~30 µm accuracy. Costing from the LED industry (Southern Machinery) says ~0.0233¢/LED; adjusting for U.S. labor rates brings this to ~0.03¢/cell. At 0.33 W/cell this is an insignificant ~0.09¢/W for cell placement. While high-speed machines have lower accuracy than slower machines, they can still achieve average error of ~5um, and worst case <20 µm accuracy (Farris, 2021), which is good enough to allow the use of a controlled volume of solder in a bounded area to make components self-centering during reflow.

- ***Micro-Transfer Printing:***

The small cells and the placement of numerous cells on a compact substrate are also ideally suited for micro-transfer printing. While Semprius has pioneered this for transferring small cells (2/3 the

size and lower concentration for ½ the power), the 40X sparser final placement forces Semprius to transfer to an intermediate substrate, and then dice that intermediate substrate and then use pick-and-place equipment to transfer each cell again to its final large substrate. In contrast, if TLC were to use micro-transfer-printing, ½ as many cells would be moved, and each cell would be moved only once. So, while TLC does *not* need micro-transfer-printing, it could make better use of it than Semprius does.

- ***Solder:***

A cell also needs solder underneath. SAC (lead-free SnAgCu) solder paste costs $1/cm^3, or 0.1¢/mm^3, so for a 100µ paste (about 60µ bond-line, pessimistic for flat cells) one gets 10 mm^2/mm^3 or 100 mm^2/¢. For a 40% module at 1500x that's 60W/¢, or 0.017¢/W.

The cell pattern is ideal for solder (and for conductive epoxy in prototyping). The cell back is all one terminal, so the solder pad is 12 times the minimum pad size for a *normal* PCB (based on Millennium Circuits specifications) and has 8 times the minimum clearance between pads (and advanced PCBs handle 008004 capacitors, in which the entire 2-teminal device is ~20 times smaller than a TLC cell). The pad size is six times bigger than a 12-mil (300 µm) pad size where Metcal's Roush already doesn't find any challenge (Las Marias, 2017). This relaxed pad size and spacing lets solder be stamped or stenciled onto the cell sites (or lets solder be plated onto the wafer back before dicing). Even when stencil-printed, the solder transfer efficiency will be very high – efficiency increases with pad size and is already ~95% for a 120µm stencil at only ~20% of the cell size (Harter, et al, 2017).

- ***Wire-bond:***

Wire-bonding is a reliable, high-volume process, with over *15 trillion* wire bonds made each year for semiconductor interconnections (Levine, 2016). Kulicke and Soffa says that wire-bonding is 48 cents for 500 bonds, or 0.1 cents per bond, for 15-mil gold wire in a 500-wire chip package (Qin, et al., 2015). A receiver is the size of a chip package so the cost should be comparable; while ~200 wires (two per small CPV cell) is less than 500 wires, in TLC all of the bond wires are the same short length and are all parallel, so the cost should if anything be lower. Slightly thicker wire uses so little extra gold that for 3-J cells (higher current and lower voltage going to a 25 µm wire is optimal at 0.11 ¢/wire for a 0.62 ¢/W wire-bonding cost.

This is probably pessimistic because the cost is from an article on reducing costs by ~40% by using copper wire rather than gold. Copper-wire-bonding is fairly straightforward, with the aluminum pad ideally being at least 1 µm thick and a thicker barrier metal under the pad, and a slightly larger pad to allow for metal splash (Notes from STATS ChipPAC on requirements for successful copper wire-bonding). Copper wire would further reduce resistance as well as the cost, but gold is currently in the models.

- ***Wire-bond Yield:***

Wire-bonding is a high yield process, with a failure rate on the order of 10 ppm (Harman, 1992). and roughly 200 wire-bonds per receiver, only about 2 receivers per thousand will have a defective bond, or one for every 15 modules. With one wire bond per cell this would remove one cell from a receiver's output; cells can trade current for voltage efficiently within a few percent of their maximum power point, and the loss of even the most illuminated cell is only ~1.5% of a receiver's power so the other cells on the receiver will decrease their voltage by a few times that 1.5% to compensate. This would cost very roughly 5% of the receiver's power, and one receiver in a thousand losing 5% of power is 0.005% power hit and thus not an issue. And with two wirebonds per cell, the second wirebond provides redundancy; the electrons travel farther on the average, reducing the cell's output by a few percent and thus the receiver's output by only a few hundredths of a percent, which will have no noticeable impact on a module's output.

- ***Total Cell-attach Cost and Small-cell Cost:***

The total cell-attach cost, even with very small cells, reverse bonds, and pessimistic assumptions, is only ~0.73¢/W, or < 1¢/W including dicing and testing, *so the small cell size is **not** an issue*.

4.2.6 Variance in Cell + Attachment Cost with Cell Size

The main cell size cost impact is the cell kerf loss, which adds ~11% of the cell area on a wafer with 40 µm streets (Masimo) or ~19.5% with Azur's 70 µm kerf. However even 19.5% is not much above the 11% with standard 5-mm dual-busbar square cells. Small cells have ~~10% higher yields and pack better on a wafer, so with 70 µm kerfs the Watts/wafer is roughly equal.

The optimum cell size thus at a cell width where kerf losses are offset by yield increases and lens width allows commodity glass with acceptable absorption, and a cell length limited by aspect ratio

for the cones to bring light to the cells. This appeared to be a cell width of ~0.65 mm, matching a lens width of ~9 mm, with a cell length pushed until higher aspect ratio would complicates the cones, or about 0.9 mm.

This optimum was studied in a paper presented at CPV-15 (Norman, et al., 2019), in was found to be a fairly broad optimum for cells that have excess lower-junction photocurrent, with little change in the total cost/W even for cells up to a millimeter wide. However, for five-junction cells without the fortuitous extra photocurrent, cells much wider than 800 µm would lose efficiency due to the thicker glass required (although antimony-doped glass should remove this constraint).

4.2.7 Higher-concentration Cells

Higher-concentration cells would be the highest-priority use for any extra acceptance angle. 2000x cells would save ~2.5¢/W, or ~10% of the entire system cost, while leaving the cells close to 0.5 mm wide. Beyond 2000x the tolerances get too tight to be manageable and the cell cost savings have rapidly diminishing returns. But cells that remain efficient at 2500x would be useful even in the current design because they would reduce the need to even out the intensity of the focus in the direction of the trough's focus.

4.3 Receiver backplane/Spreader

4.3.1 Replacing DBC/AlN with an IMS

A CPV receiver is typically built upon a heat spreader that serves as a mechanical support a cell and also serves as an electrical back contact for the cells. Most CPV systems use DBC/AlN (copper on aluminum nitride) under cells because copper provides surface electrical conductivity while AlN provides electrical isolation and a low CTE, and both have high thermal conductivity to spread heat well. Typical CPV receivers have DBC/AlN 5x to 10x the cell size, or 3 to 6 watts per cm^2 of AlN at 600x. But DBC/AlN costs ~\$3000/m^2 (Horowitz, et al., 2015), or 30 cents/cm^2, so that has been 5¢/W to 10¢/W. Even at 1000x, analyzing the NREL cost breakdown shows a substrate cost of 5 ¢/W, which is far too high when competing with low-cost silicon flat panels.

Microcells are highly resistant to CTE mismatches, so TLC can use a higher-CTE substrate. Low-iron solar glass has a high CTE (for glass), and the aluminum heatsink keeps compressive force on the lens tile, resulting in an overall module CTE of ~13 ppm/K. This is high enough that a low-cost IMS (insulated metal substrate) can be hard-bonded (e.g., epoxy) to the lens tile's thermal contacts.

While aluminum would be cheaper than copper, copper's CTE of 17 ppm/K is only 4 ppm/K higher than the expected overall module CTE, while aluminum's CTE would be 9.5 ppm/K higher, putting more than twice the stress on the receiver-to-lens-tile attachment. Copper's closer CTE match allows the receivers to be hard bonded to the thermal contacts of the lens tile strongly enough to withstand thermal cycling. This might also be true for aluminum, which would save one cent a watt with the standard IMS process, but this would have to be experimentally verified.

4.3.2 Standard IMS Process

A TLC receiver substrate was built with a standard IMS process for initial prototyping. The cells were on ~200 um (6-oz.) copper on thick ~350 um of 2 W/mK dielectric on ~500 um base copper.

Two standard-IMS optimizations have since been worked out, thermal optimization, and electrical optimization. Either provide sufficient heat spreading and also sufficient electrical conductivity, and both would fit on one minimum-sized panel so both could be tried for prototyping. The thermal optimization has the advantage of being more like the planned production receiver, while the electrical optimization would be easier to convert to an instrumented version that would measure the per-row current for profiling the trough's focus.

With either optimum the top copper is 4 oz (140 μm thick) for very low electrical resistance and good heat spreading. The isolation under the top copper should have a thermal conductivity of at least 2 W/mK, with higher conductivity better (some IMS isolations go as high as 7 W/mK). Compared to having cells directly on the base copper, the thermal optimum should add ~4K to the cell temperature and reduce the receiver efficiency by 1% through increasing electrical losses. This is acceptable for prototyping (and better than the ~8K hit in the first prototype).

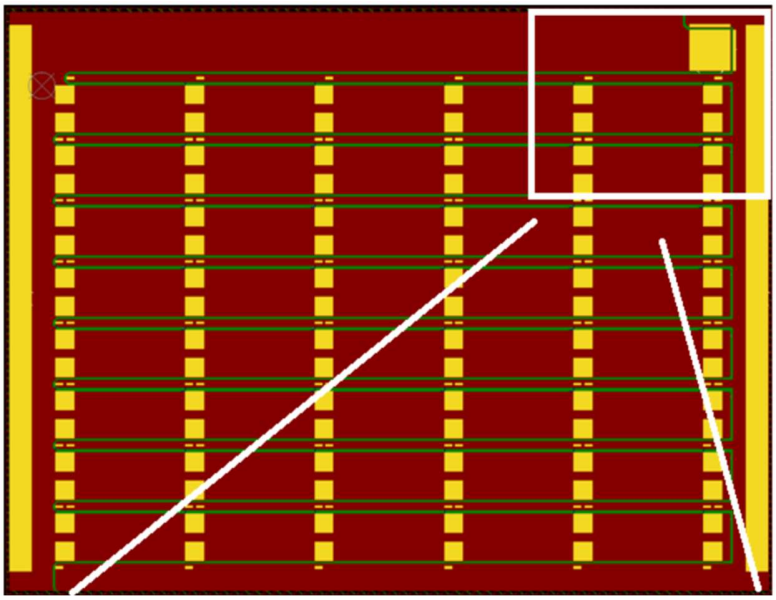

Figure 4.5: *Standard IMS layout for Thermal Optimum (2X scale, unpopulated).*

In figure 4.5, gold is backplane copper while the power plane (top) copper is brown.

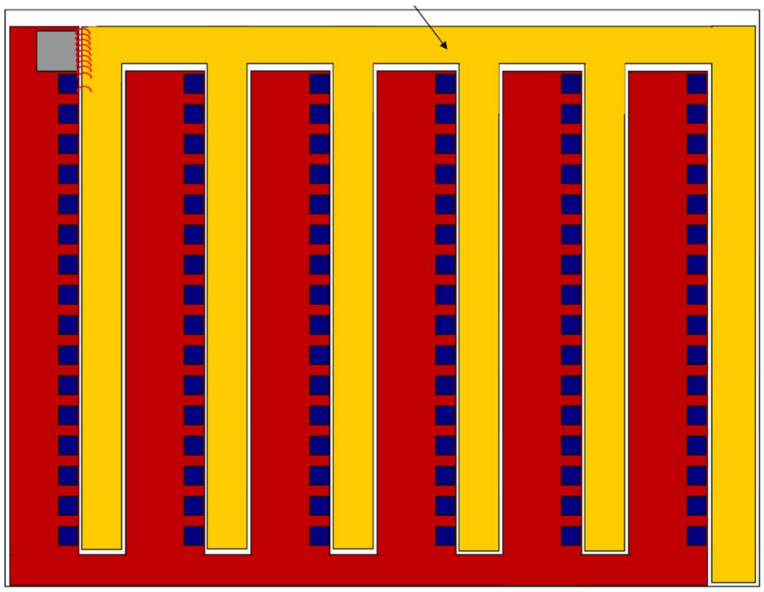

Figure 4.6: *Standard IMS layout for Electrical Optimum (2X scale, populated).*

In figure 4.6, gold is *power-plane* (top) copper while the backplane copper is brown; the CPV cells are blue, while the diode is grey.

4.3.3 Optimized IMS

Even a full plane of 4 oz. copper would not be as good a heat spreader or electrical conductor as the 0.5 mm to 0.8 mm copper base that the IMS is built on, so much better thermal and electrical performance can be achieved if the cells are directly on the thick copper base. This could be done with an IMS variation called selective dielectric removal (SDR), but SDR would add ~3¢/W, or ~20% of the module cost at a gigawatt volume, which makes it not cost-effective.

The TLC receiver pattern is much simpler than most circuit boards, so potentially very-low-cost IMS processes are proposed that would take advantage of the specific properties of the TLC receiver pattern. The ideal top copper pattern is a sheet with a hole for each cell site (and a matching hole in the prepreg); this leaves the top copper contiguous and structurally stable so the pattern could be punched before lamination to the base copper rather than etched after lamination.

The main challenge is that the 'prepreg' isolation used to bond the top copper to the base copper flows during lamination, with even 'no-flow' prepreg flowing several hundred microns. This can be tolerated only if the holes in the top copper are oversized enough that the prepreg still avoids the cell sites, or if the prepreg is removed from the cell sites afterward, and it still reduces electrical performance. In either case, the final positioning of the cell is provided by solder-surface-tension-alignment to holes in photo-imageable soldermask, which can provide roughly 10 µm resolution and hence cell position control.

Numerous variations have been proposed, all of which produce the same functional structure. The current top candidates are briefly described below and testing one or more of these is a moderately high prototyping priority because they should provide the high performance of SDR while reducing, rather than increasing the cost. All would be made as panels (typically 18" x 24") and then singulated to roughly 100 receiver substrates.

Many of the variations involve cutting or punching holes in the prepreg. While this is possible with glass-fiber prepreg such as FR-4, glass fibers are tough on a punch so it is better in these cases to use a room-temperature-punchable prepreg such as FR-1 (some of which is rated to 130°C). For mass production, paper laminates allow significant savings because all holes, slots, and even the whole profile can be punched (P&M Services comment on XXXPC & FR1/2/3). Even some thermally conductive prepreg such as CEM-3 are punchable (Hurley: Thermally Conductive Composites).

- ***<u>Punched, with Soldermask Last:</u>***

This variation is closest to the standard IMS process. The top copper sheet and the prepreg are punched with the hole pattern, with the holes oversized by several hundred microns on each side. This can be done with the sheets separate or with resin-coated-copper or with prepreg B-staged (pre-adhered) to the top copper. The top copper and prepreg are then laminated to the thick copper base sheet; the prepreg will flow toward the cell site, but should flow less than a few hundred microns since it will have a tall space (three times its own thickness) to flow into once it emerges from under the copper, but this will have to be verified experimentally since beyond a few hundred microns the holes would have to be oversized so much that the top copper would become fragile and higher-resistance. After lamination, photo-imagable soldermask is applied and imaged to define the cell sites precisely relative to one another (this key solder-mask-last step in this variation comes from Etienne Léveillé). (See **Figure 4.7 & 4.8**).

- *Slot Punched, with Soldermask Last:*

Punching a hole for each cell site quickly runs into limitations if the "no flow" prepreg flows too much. With just over a millimeter between cells, if the flow reaches about 0.5 mm, then there will be no copper left between the cells. However, the solution is simple punch a slot for 16 cells instead of 16 holes each for one cell. This can handle prepreg flow of several millimeters if needed, although the bond wire set a practical limit of about 1 mm of prepreg flow (but even some low flow prepregs, such as Arlon 49N or Isola FF406N, only flow this much). Punching a slot does increase electrical resistance by forcing the electrons to flow around the entire row of cells instead of between the cells of a row, but with 4 oz. copper this only costs ~0.3% of a receiver's output.

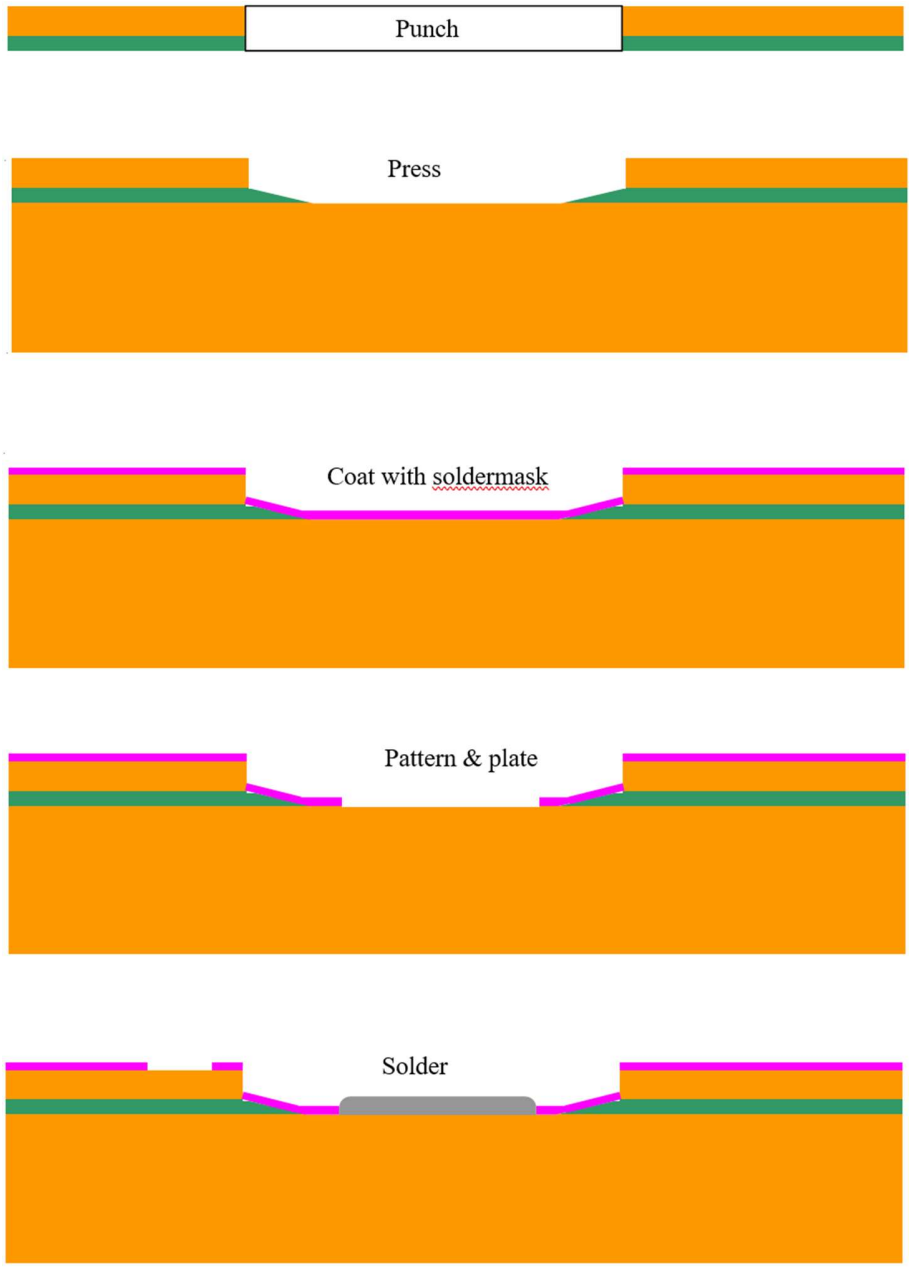

Figure 4.7: *Punch Press Pattern Plate substrate.*

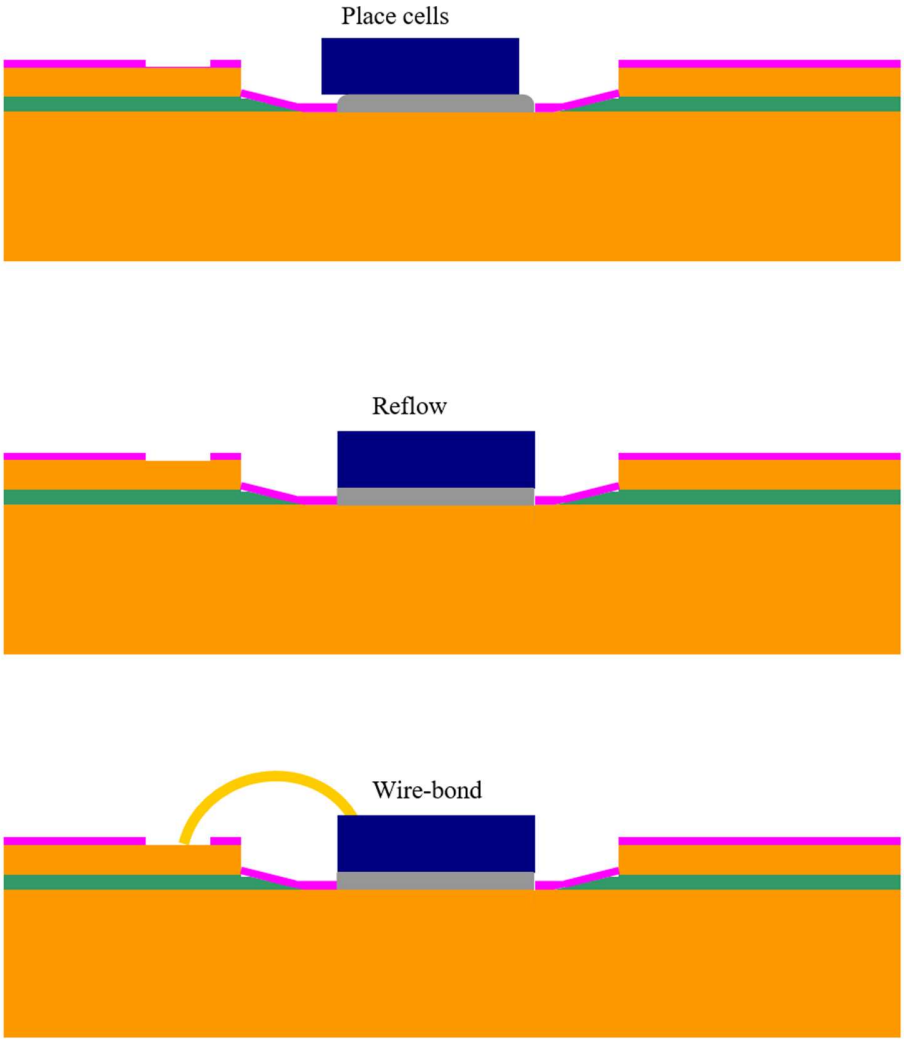

Figure 4.8: *Place, Reflow, Wire-bond.*

- *Solder Before Lamination:*

If prepreg flow would require excessive oversizing of the holes in the copper and drilling, routing and sawing all prove impractical, an option would be to soldermask the base copper first to define the cell sites, and then apply solder to the cell sites (dipping / HASL or dispense plus reflow). SAC 305 solder would not melt during lamination, and as long as the solder is thicker than ~120 µm it not only would physically block flowing prepreg but would be taller than the prepreg and could be mechanically cleaned of prepreg after lamination (if more height is needed than the solder can provide (e.g., to use thicker than 2 oz. copper), then cell-sized pieces of copper sheet could be soldered in place to act as posts to get above the height of the top copper).

- ***Laminate, Saw for SDR, Soldermask Last (Preferred):***

Rather than punch holes before lamination, intact sheets could be laminated first mechanically patterned next, and soldermask patterned last. This eliminates all concerns about the prepreg flowing too far and intruding on the cell sites. The geometry of TLC's highly parallel cell array should allow quickly removing the dielectric using ordinary sawing (or routing or drilling).

Sawing is an extremely low-cost process in high volume, and this can be done on large panels with no need for alignment between the substrate layers. A large panel of laminate comprising ~100-µm-thick copper on prepreg on ~ 0.5 mm thick copper (all standard materials in the circuit board industry) will be processed, and then be singulated into roughly 100 receiver substrates.

The panel is sawn with a shallow cut (top copper thickness plus a few tens of microns for tolerance) wider than a cell (cell width plus a few hundred microns) to remove the top copper for each row of cells. **Figure 4.9** shows a small section of a panel with the area that will become one receiver substrate outlined with a dotted-line rectangle (a full panel would probably be 10 substrates by 11 substrates). In figures 4.9 through 4.12, the power-plane (top) copper is shown in gold while the backplane copper is shown in brown. In 4.12 the CPV cells are blue, while the diode is grey.

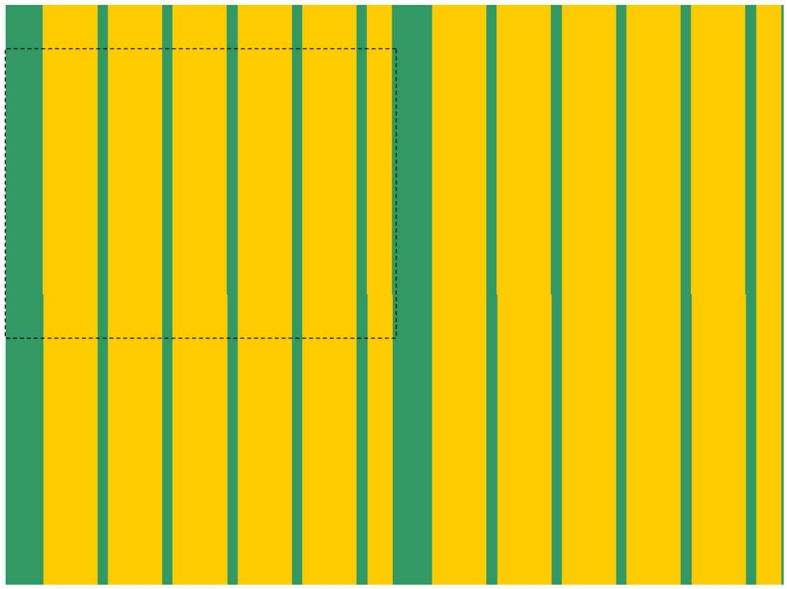

Figure 4.9: *Region of panel with strips of top copper removed to expose prepreg.*

As a second step, slightly narrower but deeper cuts and a slightly narrower cut (prepreg depth plus a few tens of microns for tolerance) are made to remove slightly narrower strips of prepreg, exposing the thick copper backplane. **Figure 4.10** shows only one substrate of the panel. The two cutting steps ensures that the saw does not smear metal from the top copper where it could contact the bottom copper and short out the receiver. While the most common prepreg has glass fibers embedded in it, using a "punchable" prepreg avoids glass fibers dulling the saw.

While both copper layers will be copper-colored, this and subsequent illustrations color the thick backplane orange and the thinner power plane gold to be easy to distinguish from each other.

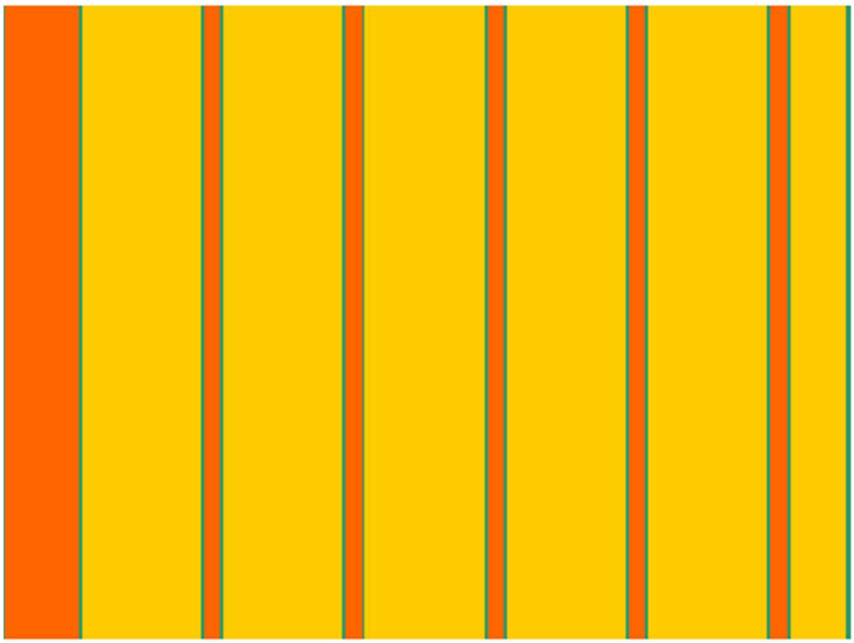

Figure 4.10: *Substrate-sized region with narrower cuts exposing backplane.*

Photo-imageable soldermask is then applied to mask off anywhere we do not want solder to flow. While the cut-widths are oversized enough (100 µm to 200 µm) to cover misalignment across the panel, within a receiver-sized region the photoimaging should provide ~10 µm relative accuracy for the cell sites (Limata, 2019). The depth of field during imaging can be up to 500 µm, allowing the tightest focus to be chosen at the depth of the backplane because the top-copper features do not need precision. (Not shown in the illustration below are the tiny (~200 µm diameter) bond-wire landing sites, which would show as tiny gold dots beside the cells).

LDI uses ordinary photoresist inks so the added cost is machine time, and combining the Limata datasheet with an EE Journal article (Maxfield, 2020) allows approximating this. A panel holds ~4kW of TLC receiver substrates, so the entry-level machine's 85 panels/hour is >300 kW/hr. The machine is ~$300K + $10K/year, so for a 10-year life the cost is $40 K per year, or $10/hr for two shifts. Allowing for an operator brings this to ~$60/hr, or 0.2 ¢/W. However, this is highly pessimistic because the area needing LDI exposure is ~1% of the substrate versus a typical 50%, so one can silk-screen large areas and just LDI around cells and bond pads to cut cost further.

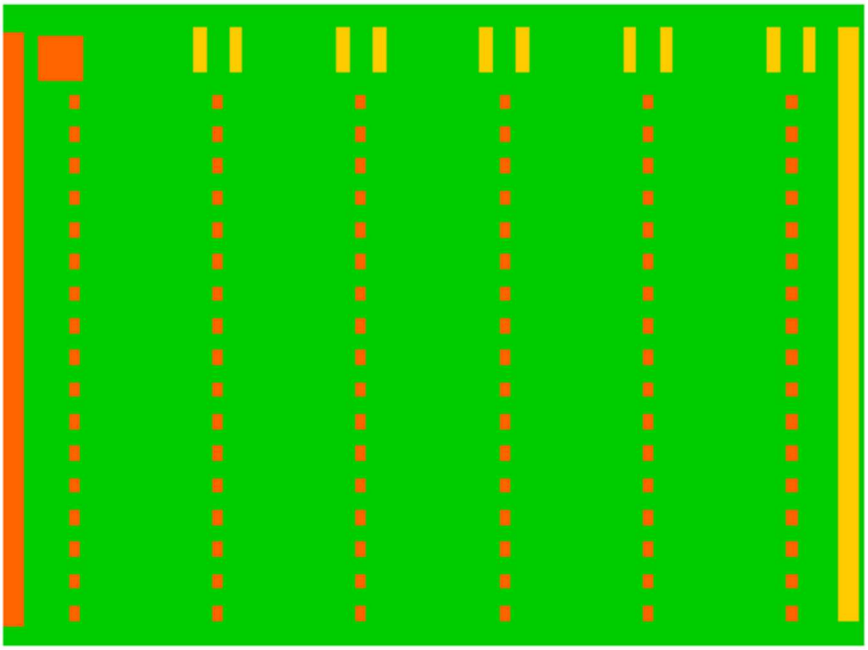

Figure 4.11: *Substrate-sized region with photoimagable soldermask.*

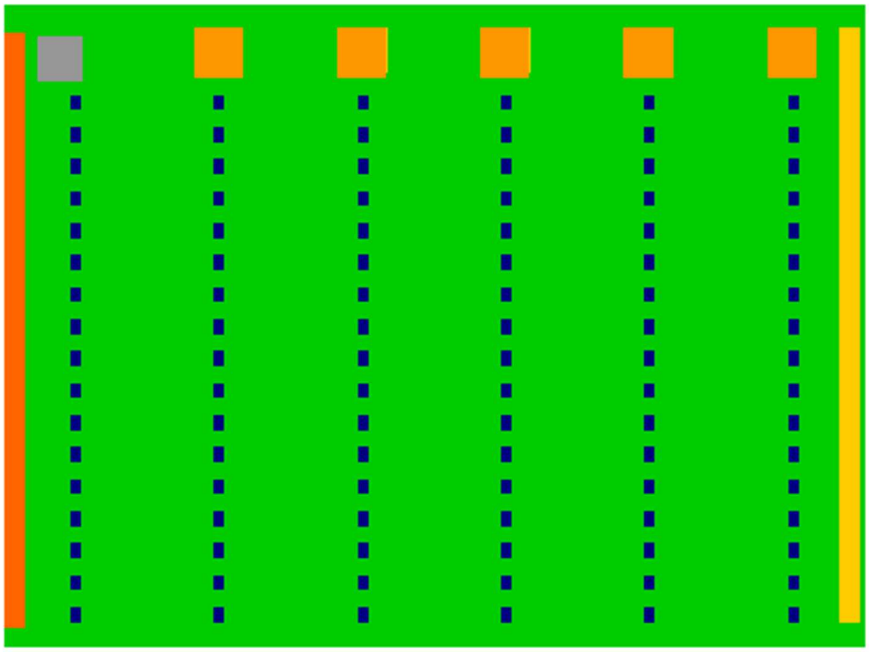

Figure 4.12: *Assembled receiver on sawn substrate.*

This divides the power plane into strips with a gap where each cell row goes, so in assembly these are rejoined by placing copper squares (roughly the size of the diode) to bridge each gap. While this loses the conduction shortcut between the cells, with thick copper the loss of power is only ~0.3%. The diode gets pushed to near the positive end of a receiver.

In addition to the thick copper backplane providing better thermal and electrical conductivity, this replaces a mask/expose/etch/strip cycle with two simple sawing steps, so in high volume it is expected to cost less as well (and be better for the environment).

- *Drilled or Routing:*

An alternative to sawing would be drilling blind holes through the top copper and the prepreg and a few tens of microns into the thick-copper base sheet (to ensure a clean surface). Although the cells are rectangular rather than round, the holes would be less oversized in the critical direction than they would have to be to accommodate prepreg flow. While drilling sounds expensive, drilling similar-sized holes in a circuit board is a highly automated process, so drilling a regular array of identical-sized holes in a panel of substrates is expected to be quite affordable (possibly less expensive than etching, although more than punching). After drilling, photo-imageable soldermask would be imaged to precisely define the cell sites within the holes.

An alternative to drilling 16 holes would be to route a single slot for each row. This loses the conduction shortcut between the cells, but with thick copper the loss of power is only ~0.3% (the diode gets pushed to the near the positive end of a receiver).

- ***Effects of Cells Directly on Base Copper:***

The most obvious effect of the optimized IMS processes is improved thermal performance. With the cells directly on the thick copper heat spreader, the heat does not have to spread in the top copper, and the receiver prepreg is not in the thermal path (thermally conductive prepreg is still used to line the metallic lid, but that is after the heat has spread from ~75 W/cm^2 to ~2 W/cm^2). The electrical performance is also improved because the entire plane of thick copper serves as the backplane for the cell array, and the array's power plane can use more of the top copper layer (except for the holes or slot for the cells). There should also be cost advantages because punching or sawing is less expensive than etching, and ordinary prepreg is less expensive than thermally conductive prepreg, but until a process is established and optimized this is just an expectation.

- *Other Isolation Thermal Considerations:*

In normal operation the top metal and its isolating adhesive are not in the thermal path, but if the module is sufficiently miss-tracked, light starts going beside the cones and several tens of suns can hit the insulation next to the cells. COMSOL ray tracing showed the this was much lower than anticipated because light leaks from the cone tops rather than the bottoms, but the insulation can receive ~~30 mW/mm^2 of light. A poorly thermally conductive plastic is generally 0.13 W/mmK, so jet-black adhesive could add as much as 240K/mm. But at ~50μ thick that's only a 12K temperature increase, and any adhesive that can take solder temperatures can easily take that, so normal prepreg is safe to use when the cells are directly on the base copper.

4.3.4 Receiver CTE Mismatches

- ***CTE Mismatch and Cells:***

The cell must handle cooling from soldering down to a cold winter night, which for SAC solder is a ~260K swing. The stress on the cells is compressive, and hard copper keeps it that way (soft copper can yield on cooling and then work-harden to produce tensile stress on re-warming), and the cells are very strong under compressive stress. Thermal cycling of a 1 mm LED (a 25% larger

diagonal than the target TLC microcell) on copper showed no deterioration after 1000 cycles from -40ºC to 100ºC (DeMilo et al., 2007). Since 200 cycles from -40ºC to 90ºC is said to be sufficient for ten years in flat panels (Kurtz, et al., 2013), 1000 cycles should be sufficient for 50 years.

- ***CTE Mismatch and Solder:***

However, between the stiff cells and strong copper is solder, and it bears the brunt of any thermal expansion mismatch. DeMilo used a hard AuSn solder, which puts even more stress on the cells, but AuSn is a stronger solder, so SAC solder needs separate confirmation.

Not only does the solder have to handle the maximum 260K drop from assembly to winter night, but it then has to handle repeated 30K-to-50K cloud swings, ~60K to 80K daily swings and up to 130K seasonal swings (cold winter night to an occasional grid-disconnect on a hot, sunny summer day). But TLC's 0.555 mm half-diagonal is 18% lower than a 1-mm LED, and these handle 165K swings on FR-4 substrates that have similar CTE mismatch. Thus, with the small cells the solder reliability should not be an issue at all, even for a >50-year product lifetime.

This is further confirmed (Nurmi, 2005) by accelerated aging of SAC solder joints between low-CTE resistors and an FR4 substrate (a ~30% higher CTE mismatch) with thermal shocks, from -40ºC to 125ºC, at >100ºC per minute (a 27% higher temperature change, incurred in 90 seconds rather than in a season). Even under this brutal accelerated aging, 0402 resistors (1.6 times bigger) survived for ~3000 thermal shock cycles (Nurmi, 2005). Since the thermo-mechanical aging of solder is exponential with size, with temperature change and with CTE mismatch (Lechivic et al., 2009), this is highly indicative that the solder under the smaller TLC cells with less of a CTE mismatch undergoing smaller and more gradual temperature changes will have a very long life.

The other form of solder aging is aging of interfacial intermetallics, which is exponential with temperature and linear with time at high temperature, so with typical temperatures below 80ºC and only brief excursions near 100ºC (e.g., a few seconds when grid goes down when there is high sun and little wind), intermetallic aging should not be a lifetime limiter.

4.3.5 Receiver Material Costs

Even for fairly thin copper, the cost of copper dominates the cost of copper sheet. With copper at $6.80/kg, PCB copper foil was ~$9.50/kg, or only $2.70/kg above copper-by-the-ton. Since a 1 kW

module uses 0.38 kg of copper sheet, the cost of the copper sheets is ~0.36 ¢/W. For the dielectric, thermally conductive prepreg is overkill but Ventec VT-4a2 is affordable at only 0.074 ¢/W.

4.4 Receiver Electrical Connections

4.4.1 Cells in Parallel on a Receiver

All cells on a receiver are electrically in parallel, so the cells on a receiver do not need to match photocurrents. This means that the cells do not need identical illumination, so TLC does not need even light on the cells and thus does not need even light across the trough's focus. This is important because a trough's focus is not even on intensity *across* the focus (even a perfect image of the sun itself is not even in intensity).

If the mirrors and tracker are properly aligned, then the cells in the end of each row *nominally* receive almost no light, and one "dark edge" can broaden when part of the width of a trough is shaded by another trough on another tracker (e.g., early morning). However, enough stray light bounces around within the lens tile that even these nominally dark cells should receive a few suns intensity, and thus have an open circuit voltage about 80% of that of the fully illuminated cells.

Even without the stray light, the QDEC cells used in early prototyping were tested and found not to suffer from leakage-induced thermal runaway, as will be detailed in the following section. This experiment and its results were presented at CPV-16 and published as a paper shortly thereafter (Norman, et al., 2020). This has also been experimentally verified by putting a prototype with QDEC cells on sun, although there the stray light comes into play.

However, AZUR 3C44 cells have been tested and found to have borderline susceptibility (less than a factor of 2 safety margin) to leakage-induced thermal runaway if no stray light is assumed. The testing will be repeated with stray light of at least a few percent of one sun to see if the beneficial impact of stray light is as significant as expected. A prototype with AZUR cells should also be tested on sun to experimentally confirm that leakage induced thermal runaway is not a problem with AZUR cells.

4.4.2 Massively-parallel Microcell Arrays and Thermal Runaway in Unilluminated Cells

- ***Foreword:***

This paper contains additional details of the tested cone molding process. This paper is kept separate from the rest of this thesis because in addition to the work of the present PhD candidate, the paper also includes work performed by other authors.

Authors: Richard Norman, Etienne Leveille, Luc G. Frechette and Vincent Aimez

All authors are associated with the University of Sherbrooke. Richard Norman is a PhD student, Etienne Léveillé is a research professional, and Vincent Aimez and Luc G. Frechette are professors.

Date of acceptance: August 19, 2020

Acceptance Status: Published

Reference: Richard Norman, Etienne Leveille, Luc Frechette, Vincent Aimez. Massively-parallel microcell arrays and thermal runaway in unilluminated cells. 16TH INTERNATIONAL CONFERENCE ON CONCENTRATOR PHOTOVOLTAIC SYSTEMS (CPV-16), May 2020, Denver, United States. pp.04000. AIP Conference Proceedings 2298, 040001 (2020); https://doi.org/10.1063/5.0032175

Contribution to the thesis: This article contributes to the thesis by detailing a potential showstopper for the TLC architecture. The article introduces a simple test and covers the testing and analysis performed to show that, at least for the cells used in prototyping, it is not a problem due to the use of microcells.

French title: Réseaux de microcellules massivement parallèles et emballement thermique dans des cellules non éclairées

French abstract:

Des travaux antérieurs ont montré que, dans certaines conditions, des cellules CPV multimillimétriques en parallèle peuvent déclencher un emballement thermique dans une cellule non éclairée lorsque des cellules entièrement éclairées pousser le courant à travers elle (M. Steiner G. S., 2013) (H. Lv, 2018). Un réseau de microcellules hautement parallèles a été proposé (Norman et al., 2018) pour s'adapter aux irrégularités de la focalisation d'un miroir à auge parabolique et

pour obtenir les avantages d'efficacité et de refroidissement des microcellules sans la pénalité financière liée à la gestion de nombreux micro-récepteurs individuels. Cet article évalue la vulnérabilité du système de microcellules massivement parallèle proposé à ce mode d'emballement thermique et constate qu'un test simple montre que, pour le type de cellule testé, il n'est pas vulnérable.

Note: No corrections have yet been requested by the members of the jury.

- ***Article:***

Abstract: Prior work has found that under some conditions, multi-millimeter CPV cells in parallel can trigger thermal runaway in an unilluminated cell when fully-illuminated cells drive current through it (M. Steiner G. S., 2013) (H. Lv, 2018). A highly-parallel microcell array has been proposed (Norman et al., 2018) to accommodate the unevenness of a parabolic trough mirror's focus and to gain the efficiency and cooling advantages of microcells without the cost penalty of handling numerous individual micro-receivers. This paper evaluates the vulnerability of the proposed massively-parallel microcell system to this thermal runaway mode and finds that a simple test shows that, for the type of cell tested, it is not vulnerable.

- *Background:*

TLC is a proposed three-stage-concentration CPV module designed to be producible at low cost (Norman et al., 2018). As illustrated in Fig. 1, an inexpensive parabolic trough (T) concentrates ~40X onto a long, narrow, glass lens tile that has a series of linear lenses on the front. Each lens (L) concentrates ~10X on the 2nd axis to form a short 400X focal line. Each 400X line is further concentrated by a row of silicone total-internal-reflection cones molded on the back of the lens tile; each cone (C) concentrates ~2.5X on the 1st axis and ~1.5X on the 2nd axis, delivering 1500X light to a tandem microcell (**Figure 4.13**). The initial ~40X concentration greatly reduces the rest of the module's size and cost.

The proposed TLC lens tile has 192 lenses along the focus of a trough mirror segment (an RP-3 inner mirror), and each lens feeds a row of 16 cones across the trough's focus, so a TLC module has over 3000 cones and thus over 3000 microcells. To avoid thousands of single-cell receivers, each TLC receiver has 6 rows of 16 microcells (each ~0.58 mm^2), with all 96 microcells electrically in parallel. 32 of these receivers are mounted on the back of the lens tile, and are electrically connected in series along the trough's focus, as seen in the energy-flow side of figure 1.

This arrangement has significant advantages. First, cells in parallel do not need to be current-matched, so this can efficiently handle the inherently uneven concentration across the trough's focus (Cooper, 2014), and having 6 rows in parallel along the focus averages-out the sub-centimeter mirror-imperfection inhomogeneities along the trough's focus (and also averages-out cell efficiency variations), minimizing current-matching issues between receivers as well.

Second, while small cells have cooling and efficiency advantages over multi-millimeter CPV cells (Fidaner and Wiemer, 2013), microcells have had higher assembly costs from handling vast numbers of tiny receivers. Small multi-cells receivers reduce this (N. Hayashi, 2017), and in TLC the trough's initial concentration will pack a receiver's 96-microcell array within a 50 mm x 50 mm area that fits high-speed placement and wire-bonding equipment for very low receiver-assembly cost. The sturdy, compact receiver will then be handled as a single part during the rest of the module assembly.

Each 96-microcell TLC receiver will produce ~30W, and its I_{SC} of 12A matches a low-cost bypass diode from the silicon PV industry. Copper tabs will auto-interconnect 32 receivers in series as

they are mounted on the back of the lens tile; the electrons flow from receiver to receiver along the trough's focus in a very short (low resistance) electrical path. Standard PV panel connectors are used at the ends of the 1720 mm long, 1 kW TLC module.

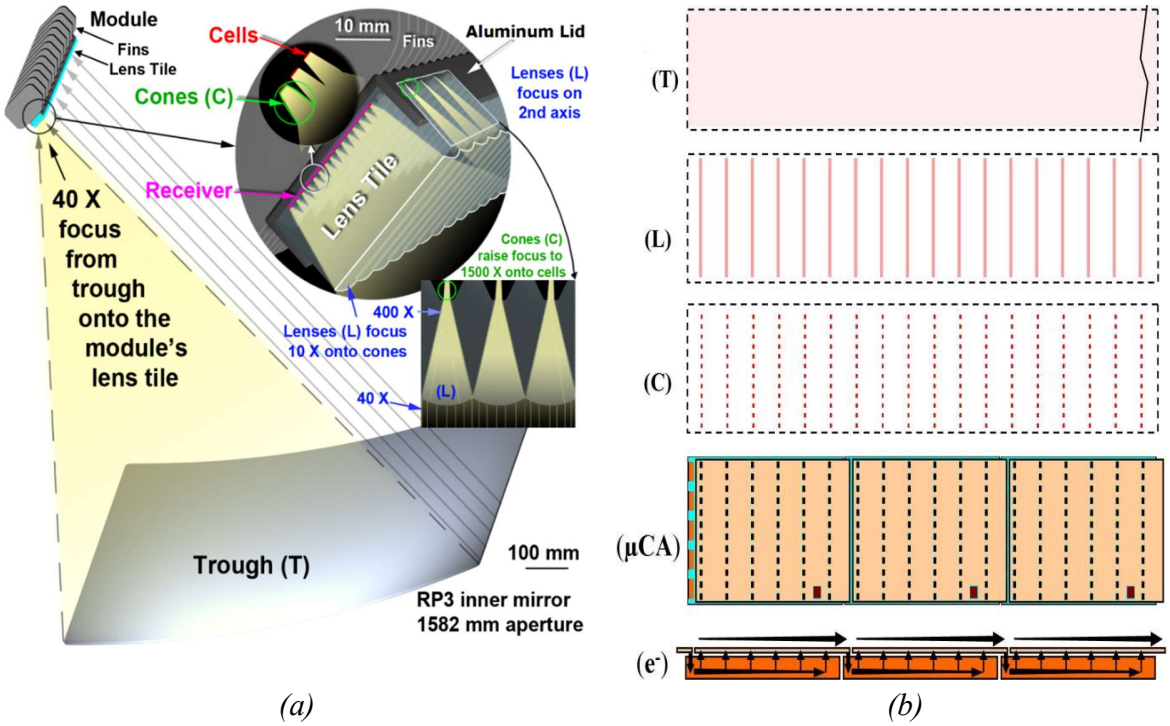

Figure 4.13: *(a) Visual diagram, (b) Energy flow.*

T: A **T**rough mirror focuses ~40X on the first axis onto long, narrow module (2-axis tracking).
L: Linear **L**enses on module front focus ~10X on the second axis, forming a series of short ~400X focal lines.
C: CPC **C**ones further concentrate the short 400X focal lines into arrays of small 1300X-1500X foci.
µCA: Microcell arrays on small receivers mate to the tiny foci. A receiver's cells are in parallel and share a diode.
e⁻: Short electron path is low resistance – a 1 kW module has 32 receivers (3 shown) in series along the trough focus.

- *The Parallel-Cells Concern:*

A TLC receiver with an array of 96 microcells in parallel is electrically similar to a standard single-large-cell CPV receiver. Instead of one large cell with many bond wires, TLC has many tiny cells each with a few bond wires, but the cell stack-up is the same as a large cell so the voltage is similar, and an array's 56 mm² cell area is between one 5 mm cell and one 10 mm cell, so the cell array's total current is similar to a standard CPV receiver. The illumination across a TLC array is uneven (a trough's focus is weaker near the edges, and there is centimeter-scale variation along the focus from mirror imperfections), but TLC's Peak-to-Average Ratio (PAR) is expected to be < 2, and many standard-cell CPV systems have more uneven illumination across each cell, with PAR > 2 (R. Herrero, 2015).

However, all diodes (including CPV cells) leak, and prior work on parallel cells (M. Steiner G. S., 2013) showed that when a receiver has multiple multi-millimeter cells in parallel, under some conditions the illuminated cells inject sufficient current through a dark (unilluminated) cell in the same receiver to heat the dark cell enough to increase its susceptibility to injection enough that increased injection further heats the cell in an escalating cycle of thermal runaway. While the point at which runaway is triggered depends on the details of the cell and on the cell cooling, Steiner et al.'s analysis of the Flatcon® system (M. Steiner G. S., 2013) found the most vulnerable case to be an unilluminated cell when the module was operating at the VOC of fully illuminated cells, and that the danger grows with concentration, and that beyond 800X there is sufficient current available for injection to be a concern with as few as three cells in parallel.

More recent work by Lv (H. Lv, 2018) evaluated the issue in detail for 30.25 mm^2 cells and showed that if the dark area was < 20 mm^2, the added heat from injection was dissipated rapidly enough to not trigger runaway even at 1500X. This suggests that microcells might be less susceptible, but in TLC many closely spaced microcells might be dark (e.g., while tracking a module from off-sun to on-sun, most of an array will be fully illuminated before the cells near one edge of the array are in the trough's focus), and their total dark area could exceed 20 mm^2.

The aggregate current from most of a TLC receiver's array of cells would be sufficient to overheat multiple microcells, so in TLC runaway is limited *by the cell's ability to resist injection* rather than by current available for injection. A dark cell's susceptibility to injection rises rapidly with temperature, so at some critical temperature the injected current would overwhelm a dark cell's cooling and trigger thermal runaway. The question thus becomes whether TLC's cooling is sufficient to keep cells below this critical temperature during operation of a TLC module.

An unilluminated cell doesn't have heat from concentrated sunlight, so in the worst case (M. Steiner G. S., 2013) of a totally dark cell, the Watts from injected current heating (W_I) replace a cell's normal thermal load from optical energy (W_O). If, at the maximum normal cell temperature, the power injected into a dark cell (W_I) is less than the optical heating (W_O) of a fully illuminated cell, then in normal operation a dark cell will be cooler than a fully illuminated cell, and thus cooler than the maximum normal cell temperature, and cooler than the cell being tested and thus *less* susceptible to injection, and thus below the critical temperature. $W_I < W_O$ defines a safe zone that is simple to test hold an unilluminated cell at the system's maximum normal operating cell

temperature, and inject current through the cell at a fully-illuminated cell's V_{OC}; if injected Watts (I*V) is less than normal optical heating, then $W_I < W_O$ and any dark cell will be below the critical temperature and not suffer runaway if voltage spikes to V_{OC} in normal operation.

While cell susceptibility to injection limiting W_I to $< W_O$ is *sufficient* to avoid injection-induced thermal runaway, it is not a necessary condition. Outside of this zone a receiver's cells might be safe, but more detailed analysis (e.g., multi-physics simulation) would be needed to determine safety, as was done in Steiner (M. Steiner G. S., 2013) and Lv (H. Lv, 2018).

- *Determining Whether the TLC Receiver Falls in the WI < WO Safe Zone:*

QDEC triple-junction CPV cells will be used in TLC prototyping, so QDEC cell V_{OC}-versus-concentration data was obtained from Wheeldon, et al. (J. Wheeldon, 2011). V_{OC} is very linear with the log of concentration, so extrapolating on the semi-log plot of V_{OC} versus concentration (figure 3 from Wheeldon) is straightforward and shows that the V_{OC} at TLC's 1500X would be ~3.16V when the QDEC cell is at 25° C. Under normal conditions the average cell in a TLC receiver is expected to produce a maximum of ~400 mW of heat (as well as ~ 340 mW converted to electricity rather than heat), so the injected-current limit for the $W_I < W_O$ safe zone is ~400 mW / 3.16 V = ~127 mA.

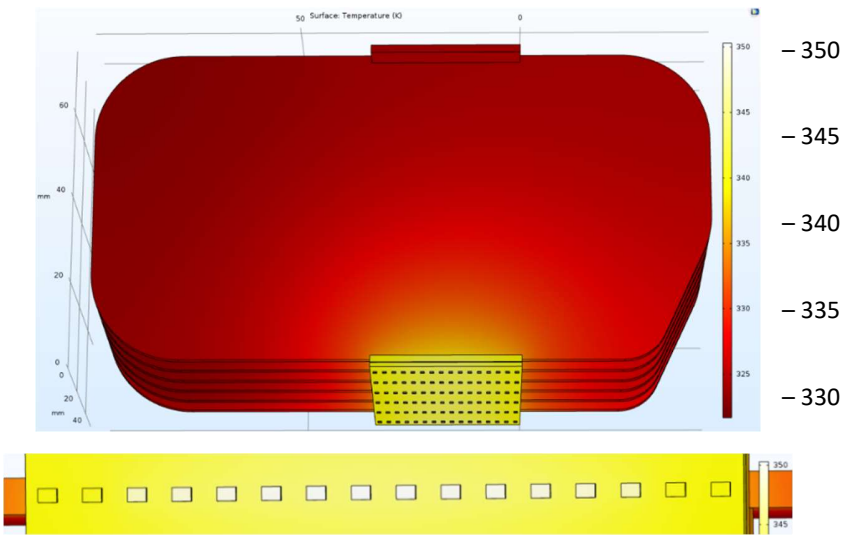

Figure 4.14: *With no wind, COMSOL calculates a maximum TLC cell temperature of 50°C above ambient (300K).*

Microcells shed heat well, and TLC's microcells share a finned aluminum heatsink whose fins are always vertical (good heat transfer even with no wind) and which have an aggregate surface area significantly larger than the trough's light gathering area. The TLC heatsink is sized so that COMSOL® thermal modelling of TLC module shows a worst-case normal-operation cell temperature of 50°C above ambient, so even at a very hot 50°C ambient the TLC cell temperature should not exceed 100°C (**Figure 4.14**). Thus if 3.16 V drives less than 127 mA through an unilluminated cell held at 100°C, the TLC receiver will be in the $W_I < W_O$ safe zone in normal operation (up to at least 50ºC ambient). The TLC cell size planned for production is 0.9 mm x 0.65 mm, so 127 mA is 217 mA/mm^2. Customizing cells would, however, be expensive for the small quantity needed for early prototyping, especially since the knowledge gained from prototyping might change the optimum cell size or other cell attributes. Early TLC prototyping will therefore use pre-existing cells of a larger size (with only part of the larger cell illuminated). When the current *available* to inject is large, a dark cell's injected current is proportional to the cell area (H. Lv, 2018), so larger cells are more vulnerable to this type of thermal runaway (the heatsink, and thus the main thermal resistance, does not change, so more injected current means a higher temperature). The injected current density through a cell, as a function of voltage and temperature, can thus determine not only whether the planned production module will be in the $W_I < W_O$ safe zone, but also how large the cells of a given type can be for prototyping before the cell size pushes the prototype out of that safe zone.

- *Experiment and Results:*

To determine its susceptibility to injected current as a function of voltage and temperature, a small QDEC cell was pressed against a strip of copper tape adhered to a temperature-controlled hotplate. A Keithley programmable voltage source had one probe pressed against a contact point on the cell's top metallization, and the other probe pressed against the copper tape (and thus in electrical contact with the back of the cell). The hotplate was heated to the initial test temperature and allowed to stabilize for a few minutes, and the attack voltage was then ramped from 0 to 3.3 V in 0.1 V steps, as shown in figure 3. The hotplate temperature was then increased and the voltage was again ramped from 0 to 3.3 V (or until the current source's limit was reached at ~ 60 mA/mm2), and this was continued until a hotplate temperature of 100°C was reached.

At the end of the 100°C run it was discovered that the microscope light used to align the probe to the cell had been left on, so the cell was not as dark during the tests as planned. To confirm that

the less-than-one-sun illumination makes little difference, the 100°C test was rerun with the light off (100 NL = no light). As figure 3 shows, this had little impact on the injected current at higher voltages.

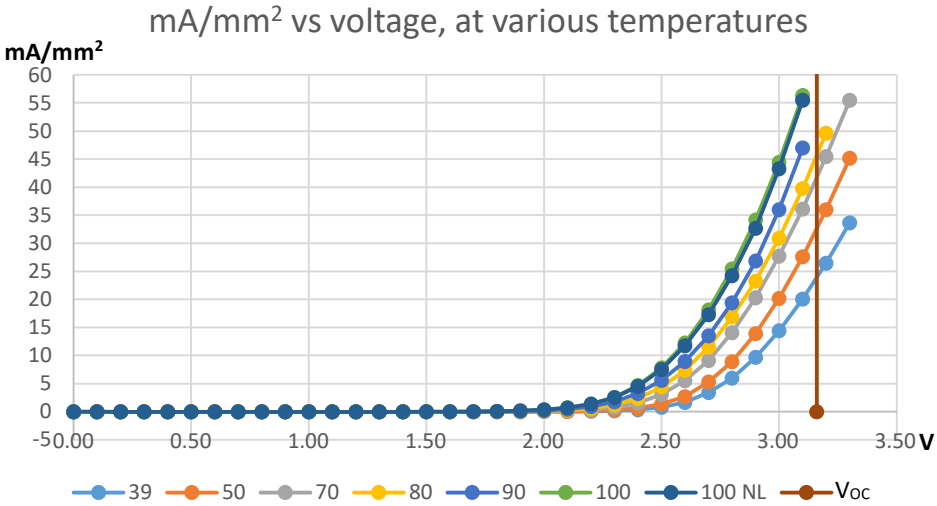

Figure 4.15: *Injected current in increases rapidly with both voltage and temperature.*

- *Analysis:*

As can be seen in **Figure 4.15**, injected current increased rapidly with both voltage and temperature, but even at 100°C and at the reference-temperature V_{OC} of 3.16V, the ~ 65 mA/mm^2 injected current was well below the 217 mA/mm^2 $W_I < W_O$ safe zone limit for production-sized cells. But cells more than ~3.3 times larger (217 mA/mm^2 / ~65 mA/mm^2 = ~3.3), or more than ~1.9 mm^2, would be outside the $W_I < W_O$ safe zone at V_{OC} = 3.16V.

However, that safe zone calculation is based on the V_{OC} of attacking cells at 25°C, and since the cells of a TLC receiver share a copper heat spreader and share an aluminum heatsink cooled by the same ambient air, if the illuminated cells each produce more heat than a dark cell, the attacking cells will be at least as hot as a dark cell. For QDEC cells, V_{OC} drops by ~4.3 mV/ºC (J. Wheeldon, 2011), so a 100ºC attacking cell (75° hotter than 25°C) will have a V_{OC} ~322 mV lower, or ~2.84V. Since the injected current rises rapidly with voltage, the attacking cells' temperature has a significant effect, as can be seen in the semi-log-scale graph of figure 4.

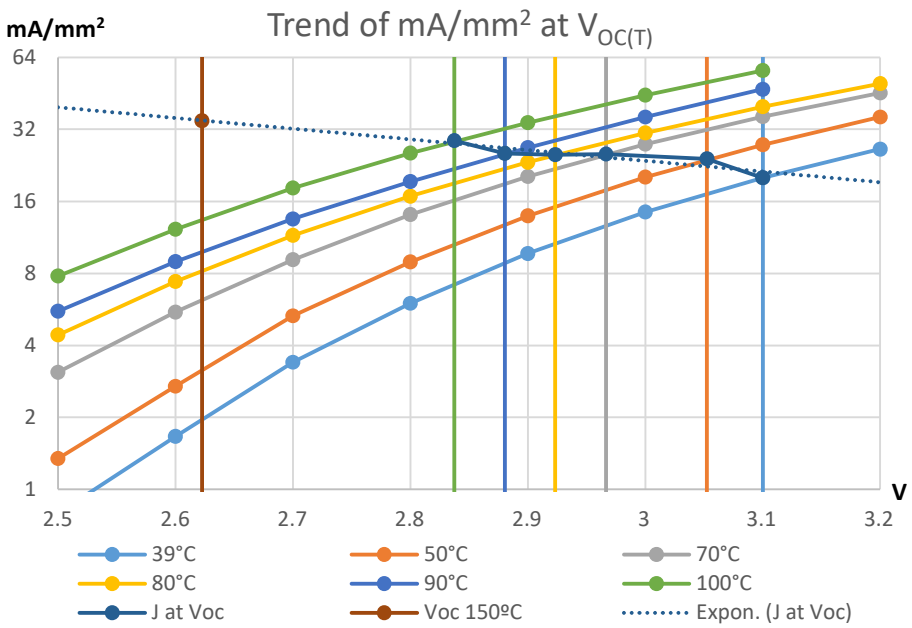

Figure 4.16*: Approximate trend of injected current versus temperature when attack-V_{OC} is temperature-corrected*

In **Figure 4.16**, VOC is calculated for each temperature at which the current was measured, and a vertical line is added for the VOC for each temperature. The approximate maximum injected current density for each temperature is at the intersection of the measured current density for that temperature and the VOC reference line with that temperature. These intersections are tagged with black dots connected by a black line for visibility, and a least-squares-fit trendline is overlaid to allow rough estimations of injected current at the VOC at other temperatures.

The optical heating used so far has been for normal operation near the maximum power point (the most important case), where a significant proportion of the incoming optical energy is converted to electrical energy, rather than heat, until a voltage spike to VOC (e.g., the grid goes down). However, if operation near VOC rather than near VMPP is sustained, the cells' full incoming optical energy (except reflected infrared) is converted to heat, almost doubling the thermal input (~740 mW versus ~400 mW) which raises the receiver temperature significantly over the course of a few minutes. Because the whole array of cells shares a heatsink, this raises the temperature of dark cells as well as of the illuminated cells. (All incoming optical energy becoming heat also covers the case of enough dark cells to convert the array's entire electrical output to heat, since electrical output comes from the optical input.)

An additional VOC reference line, labeled 'VOC patho', has been added at the VOC calculated for the ~150°C maximum cell temperature approached in pathological conditions (staying on sun with the grid down so that all optical energy becomes heat, at 1000 W/m2 DNI and 50ºC ambient with no wind). The intersection of the trendline and the pathological-case VOC of the illuminated cells approximates the pathological-case injected current density. Although measuring one cell and extrapolating an exponential effect out to a modeled performance leaves significant uncertainty, an injected-current density of ~35 mA/mm^2 is far enough below the 217 mA/mm^2 $W_I < W_O$ safe zone limit for the expected production cell size that this type of thermal runaway is unlikely to be a problem for this system with this type of cell. The extrapolation suggests that prototypes with QDEC cells up to roughly 6 times larger than the planned production cell size, or 1.8 x 1.8 mm^2, would survive even pathological-case testing. However, extrapolating an exponential leaves high enough uncertainty that smaller cells will be used for safety until the injected-current density at the pathological-case temperature and V_{OC} can be confirmed on multiple QDEC cells.

- *Discussion:*

The purpose of this approach is a quick test, that can be performed before a physical module prototype is available, to determine whether a given type of cell would be susceptible to injection-current-induced thermal runaway if used in a given parallel-microcell array. It does not evaluate how hot a cell would get if the safe-zone limit were exceeded; it evaluates whether a dark cell is in a safe zone in which it stays cooler than fully illuminated cells. The cell type tested here had ample safety margin, but in borderline cases some refinements may be useful.

The pathological case already covered the injection current heating's effect on the overall receiver temperature, and the normal-operation analysis can also include it. Although a dark cell will be cooler than an illuminated cell, the injection-current heating does increase the heatsink temperature, and thus the overall receiver temperature including the illuminated cells. From figure 4, ~28.7 mA/mm^2 is injected at 100°C. In a 0.58 mm^2 microcell this is ~16.6 mA, or ~47 mW at 2.84 V. In a TLC receiver's 6 x 16 array, any azimuth miss-tracking that leaves a row of 16 cells dark also removes at least a row's worth of incoming light, so the heatsink and thus the cells will be cooler. Altitude miss-tracking cannot leave more than 18 cells (3 columns of 6 cells) dark without missing significantly more optical energy than the ~850 mW that 18 dark cells would produce. 850 mW is ~2.2% of the normal thermal load, so it would add roughly 2.2% of the 50ºC above ambient, or

only ~1.1ºC, which still leaves ample margin. However, with more-susceptible cells, or in less-well-cooled systems, this would be more significant.

Cell size has some effect on V_{OC} [5], with smaller size reducing Voc. Over the size range discussed for prototyping, V_{OC} would be ~0.04 V to 0.06 V lower that the cells measured in Wheeldon (J. Wheeldon, 2011), so a V_{OC} determined from Wheeldon errs on the safe side, but if the test cell were much larger or smaller than the planned cells, the V_{OC} impact should be accounted for.

In a CPV system on sun, no cells likely to be truly dark. Steiner identified a totally dark cell as the worst case for current injection (M. Steiner G. S., 2013), and due to the logarithmic rise in V_{OC} with intensity, a cell receiving only a trivial amount of optical energy has considerable voltage opposing the injected current, and so is less susceptible to injected current. In the TLC design, where the entire cell array shares much of the optical system, this is likely to provide a significant safety margin – if even 0.1% of the light is scattered evenly to the cells within the shared optics, a "dark" cell will have a V_{OC} of ~ 2.5 V opposing the injected current.

Injected current does not heat the cell evenly; in a dark cell the area under the grid fingers will be hotter because it presents the lowest resistance path through the cell, so more of the current flows there. Before runaway is triggered, the impact is minor for 3J cells of typical sheet resistance. For example, at a grid finger pitch of 120 µm (AZUR SPACE, 2020) and a width of ~20 µm, the farthest distance from a finger is only ~50 µm. Each mm^2 of cell has an aggregate path 16 mm wide (1 mm of each of 8 grid fingers, times 2 directions) to flow that 50 µm, for a path resistance of 1/320 of a "square". A typical triple-junction-cell first layer sheet resistance is 200 to 220 Ohms/square (M. Steiner P. K., 2015), giving 1 mm^2 of cell a resistance of ~ 0.65 Ohms from grid fingers to farthest points, so even if all 35 mA/mm^2 of the pathological case flowed this far, the voltage drop would only be 23 mV, and since the current flow drops off slightly faster than linearly (due to the decreasing injection-current density), the voltage drop will be slightly less than half of this, or ~0.011V. It can be seen in figure 4 that < 0.011V makes very little difference in the current injected per area, and hence the area right under the fingers will receive very little extra current and will not be significantly hotter. This effect increases with current, and by the safe zone's limit of 217 mA/mm^2 it makes a ~70 mV difference, which corresponds to a ~25% difference in injected current per area, or a variation +/- ~50 mW/mm^2. The impact of this on temperature is still minor because this just changes where the heat is produced, and even if that local extra 50 mW/mm^2

flowed vertically through the entire ~200 µm thick cell before spreading laterally it would add less than 1°C to the cell temperature under the grid line. But if the critical temperature were to be exceeded, runaway would start under the gridlines, and, once started, would escalate there fastest.

With TLC's shared optics it is hard for a middle cell (in a row of 16) to be dark, and the edge cells have slightly better cooling, increasing margin against runaway. If edge cells had worse cooling, this should be accounted for.

W_I greater than W_O would not rule out a given cell for a given parallel-cell CPV receiver, but more detailed analysis (e.g., multi-physics modeling) would be needed to determine at what temperature the added heat from higher injected current would overcome the receiver cooling's ability to extract the added heat and lead to an escalating cycle of further temperature and injected current increases. When a hot cell can borrow cooling from other cells through sharing a heatsink, thermal runaway could be delayed to well beyond the limits of the $W_I < W_O$ safe zone. However, even if it is established that a given cell in a given system is well within the $W_I < W_O$ safe zone and thus that injected-current-induced thermal runaway would not occur, it is still good to track off sun when the grid goes down to avoid the higher temperatures rapidly aging the cells and modules.

- *Future Work:*

A variety of CPV cells will be obtained and will be tested for susceptibility to injected current versus temperature to determine which types of cells are most suitable for highly-parallel microcell arrays. At least two types of CPV cells will then be assembled into TLC mini-modules and characterized on sun under a range of DNI, ambient, and alignment conditions. Some CPV cells are unlikely to be suitable; for example, a technical note (Bensch, 2011) for at least some AZUR SPACE cells recommends that to avoid overheating "The inhomogeneity of the cell illumination shall not exceed factor 2" (a much more stringent requirement than a PAR of two), which suggests that such cells would be unlikely to be suitable for a highly-parallel microcell array with uneven illumination. However, due to the electrical similarity of a TLC receiver to a single-large-cell receiver, it is expected that most CPV cell designs that can tolerate a PAR >2 on a single large cell would be suitable for TLC's microcell arrays.

- *Conclusion:*

Measuring the injected current density versus voltage and temperature and calculating W_I versus W_O provides a simple way to confirm a given CPV cell's suitability for use in massively parallel

receivers once the module's cooling performance has been either modeled and analyzed or built and characterized. The cell' injected-current versus voltage and temperature can also be used to determine a safe cooling performance for a given cell type and size such that $W_I < W_O$ ensures that injected current will not trigger thermal runaway.

With the cell's injected current density versus temperature and voltage measured, and the module cooling characteristics modeled, this process was used to estimate the maximum cell size that can be safely used in TLC prototyping without needing detailed multi-physics analysis. Well-cooled QDEC microcells are shown to be resistant to injection-induced thermal runaway, and the results suggest that QDEC cells up to roughly 1.8 mm x 1.8 mm could be used for TLC prototyping without triggering thermal runaway even in worst-case test conditions.

- *Acknowledgments:*

The authors thank Programme Innovation, IRAP and TFI for financial support for this work, and thank the University of Sherbrooke for access to expertise and test equipment and CMC Microsystems for the use of COMSOL®. The authors thank TFI for access to its detailed TLC design and use of overview illustrations and descriptions for TLC; TLC is a trademark of, and the TLC design is the intellectual property (pat. pend.) of Terra Firma Innovations Inc. (TFI). The authors thank LN2 - a joint International Research Laboratory (Unité Mixte Internationale UMI 3463) funded and co-operated in Canada by Université de Sherbrooke (UdeS) and in France by CNRS as well as Université de Lyon (UdL, especially including ECL, INSA Lyon, CPE) and Université Grenoble Alpes (UGA). LN2 is also associated to the French national nanofabrication network RENATECH and is supported by the Fonds de Recherche du Québec Nature et Technologie (FRQNT).

4.4.3 Back Contact to Substrate as Conductor

Cells in parallel can all be soldered to the 0.5-mm copper heat-spreader. With ~20 mm^2 of cross-section, its 0.04 mOhms costs just 0.02% of the power with 3J cells (and half that with 5J cells).

4.4.4 Top-Contact Conductor

All cells on a substrate in parallel also lets all cells be wire-bonded to a common front conductor, often referred to as the top metal, or the top-copper. A copper sheet (ENIG plated) allows the bond

wires from the cells to terminate wherever it is convenient on the flat surface, and 2-oz top copper has a resistance of 0.32 mOhm for the whole length of the receiver, costing 0.16% of the power (and half that with 5J cells). The top of the bypass diode also gets connected to this top metal.

4.4.5 Bond-wire Resistance

While the bond wires could land anywhere on the top metal, it is expected that they will simply bond from each cell's pad to the nearest spot on the top copper in the direction that the electrons have to flow anyway, slightly reducing resistance while keeping all bonds short and parallel. Although the bond wires are small diameter, they only carry the current from one cell and they are only ~0.5 mm long. With two 25-µm gold-wire bonds per cell, the resistance costs ~0.07% of the power (~2 mV) with current 3-Jct cells (and (~0.04% with 5J cells).

4.4.6 Bypass Diodes

Having all cells in parallel also lets a single bypass diode be used per receiver; 12A diodes cost ~16¢ each in 10 MW quantity (Vishay quote, 2016), or **0.5 ¢/W** for the target initial receiver size. Prices per Ampere will come down only slightly at GW volume (since diodes are already made in large volume anyway for flat panels, this is purchasing negotiation rather than manufacturing cost), but this diode is ~2x oversized for the upcoming 5J cells and smaller diodes will cost almost proportionately less, so 0.25¢/W is estimated for GW production in 2025.

The diode is soldered similarly to the cells, taking 0.06 ¢/W of solder. It is also wire-bonded, at 0.6 A/bond this uses 20 bond-wires with 3J cells, at a cost of 2.2 ¢ per receiver or 0.067 ¢/W.

- ***<u>Smaller Diodes:</u>***

In several of the substrate layouts the size of the diode forces the substrate to be wider (this is not the case when a per row slot is used because the width provides copper to reduce resistance). Smaller diodes are available, which would allow a narrower receiver substrate, but multiple smaller diodes would currently be needed. Currently these would cost more than one 12A diode, but if the substrates cost more than expected this might be worthwhile (especially with higher voltage/lower current cells). 1A diodes are currently 0.9 mm on a side and 2A diodes are 1.15 mm on a side.

4.4.7 Interconnecting Receivers

The broad areas between the cell rows will mate to the lens tile's thermal contacts when the receiver is bonded to a lens tile. For the five thermal contacts that land in the middle of the receiver, structural epoxy will be dispensed, either on the thermal contacts or on the receiver, to secure the receiver in place (**Figure 4.17**).

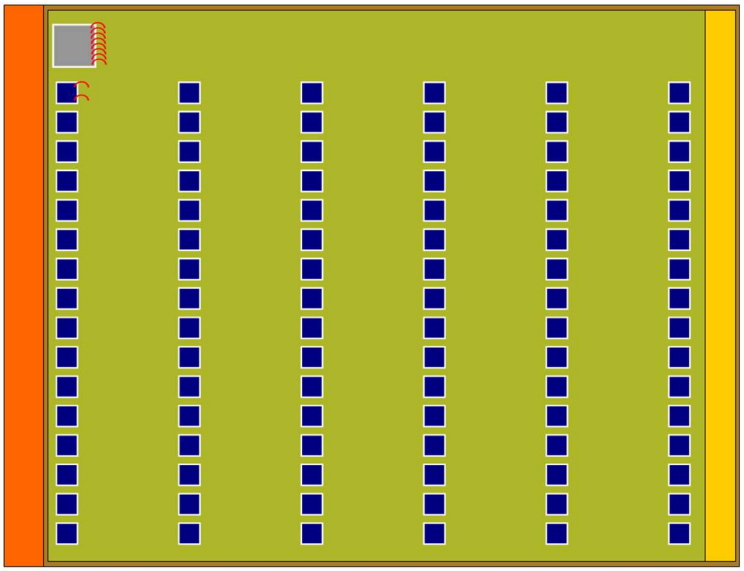

Figure 4.17: *Exposed Conductors ~2.5x scale (backplane is orange, power-plane is gold).*

The thermal contacts where two receivers meet will also serve to electrically interconnect adjacent receivers in series. A strip of copper foil will be applied to each of these thermal contacts before any receivers are placed (in early prototyping this will simply be adhesive backed copper foil tape). During assembly, electrically conductive epoxy will be dispensed on these thermal contacts. Regardless of which processes is used to make the receiver substrates, one end of a receiver will have the receiver's backplane (the positive contact for all cells) exposed and the other end will have the receiver's power plane (the negative contact for all cells) exposed (**Figure 4.17**). These exposed ends will contact the electrically conductive epoxy, allowing the electrical conductor on the thermal contact to connect the positive end of one receiver to the negative end of its neighbor, thus placing the receivers electrically in series (**Figure 4.18**). This avoids all connectors cages and all interconnect wires *within* the module, providing lower resistance and much lower cost.

All receivers in a module are nominally identical and all receive nominally identical illumination, and the module support does not shade any receivers. With dozens of cells per receiver averaging out hero cells and poorly performing cells, receivers will be very well photocurrent matched. The receivers can thus be electrically in series with essentially no mismatch losses (the diodes will handle any defective receivers, or receivers in modules shaded by dirt, clouds, or other trackers).

Figure 4.18: *Electron Flow through Receivers (thickness exaggerated).*

- ***Stress in the Interconnect Strips:***

SunPower uses copper tabs to connect silicon cells due to the high reliability this provides for the modules (Campeau, et al., 2013). SunPower adds stress-relief cut-outs to the tab copper; this should not be necessary in TLC since TLC substrates are 1/3 the length and also have 1/3 the CTE mismatch. But if needed, such cut-outs could easily be added in TLC at trivial cost since only ~10% as many tabs are needed per Watt as with SunPower cells and each is only 25% as long.

4.4.8 Module Wires

Insulated wires with connectors are *only needed at each end of the whole module.* The module terminal wires will be hermetically sealed in place with compressed elastomer (like a NEMA-rated enclosure) or epoxy. These module wires are not part of the normal receiver; instead, a small one-row "receiver extender" is used at each end of the module (the extra row prevents photocurrent mismatch issues from the less intense light from the very edge of the module's mirror), and this receiver extender has a poke-in connector cage *(see also 'Module Ends')*.

The cages from the NREL CPV costing study are used; these are rather large, but they will accommodate overkill 14 ga terminal wires (18 ga would be sufficient). The lens tile will have pockets rolled in to accommodate the cages. While the NREL CPV cost study has two cages per 7.5W cell assembly for ~2.7¢/W, TLC only uses two cages per ~1 kW module, for 0.02 ¢/W.

4.5 Heat Rejection Housing

4.5.1 Overview of Heat-rejection Methods

With the receivers lined up along the focus of the trough, the heat must be efficiently carried away to prevent overheating. The best way to do this depends on whether the heat is a useful byproduct to be collected or is waste to be rejected, on whether the trough is on a single-axis tracker with the receivers micro-tracked or is on a dual-axis tracker with a solid module, and on the width of the trough and hence the amount of heat to reject.

- ***Finned Heatsink:***

The simplest heat rejection is to have fins (or a finned heat sink) directly on the back of the module. However, a simple heat-sink doesn't scale well because taller fins need to be further apart to prevent the viscosity of the air from slowing down the airflow, and the wider gap decreases the heat transfer coefficient for conduction into the air.

While a solid purely conductive heatsink was the first type looked at, the design at that time comprised a wide mirror on a single-axis tracker, with micro-tracked receivers in an enclosed housing, which created challenges for this simple type of heatsink. Once an idea has been rightly dismissed under one set of circumstances it is hard to rethink it when circumstances change, but when the newer design was described to Mahmoodreza Salimshirazy (A cooling expert at the University of Sherbrooke), he asked why a simple solid conduction heatsink was not used, and the old reasons to reject it were found no longer be valid. An RP-3 mirror, even on a dual-axis tracker or a polar-axle single-axis tracker, is near the mirror width-limit for cooling with a solid-material finned heat sink. However, the simplicity of only the air moving is a significant advantage, and, initial calculations show that even at an RP-3 mirror width this probably wins for lowest total cost.

- ***Gathered Heat:***

Heat can also be gathered rather than rejected. This is basically a variant of pumped coolant, but with the coolant distributed from a large pipe running the length of the system rather than being cooled locally through fin tube and reused. For DAT this would require gathering coolant from numerous two-axis trackers, so heat byproduct is much easier on a long trough in SAT TLC.

- ***Finned Heatsinks Look Best:***

Only DAT TLC is considered in this section because it currently offers the lowest overall LCOE (SAT cooling is described in detail in Appendix B of this document). Also, only fins-on-lid heatsinks are considered (heat-pipe/fin-tube heat sinks are described in appendix A.5 and A.6).

In DAT TLC the module lid seals hermetically to the glass lens tile to keep out moisture, conducts heat from the receivers to the heat-rejection fins, and mechanically supports the fins.

4.5.2 Fin Heat-Rejection Scaling Background

Except where heat is a useful byproduct, the heat needs to be rejected to the surroundings, typically to the ambient air. Air has low heat capacity and low thermal conductivity, so a large surface area is needed to transfer the heat to the air. Although the heat-spreader receiver substrate spreads the heat out within the module, the heat is still concentrated by the primary concentration of the trough so a heat-rejection surface area much larger than the module is needed. A standard way to get a large surface area near a heat source is with arrays of thin metal fins. Aluminum is the lowest-cost formable, stable solid per unit of thermal conductivity, with magnesium at 60% of the conductivity per dollar and steel at 40% (sodium is actually better at almost twice the conductivity per dollar, but bursts into flame on contact with water, and some natural graphites are better, but are not as formable as a metal). All fins currently considered on any embodiment of TLC are thus aluminum.

- ***Natural Convection:***

Most of the potential heat rejection solutions use natural convection to move air across the fins. As the fins warm the air, the air's density decreases and the air rises, and fresh air flows in at the bottom of the fins. Natural convection is boosted by whatever wind there is, so when the fin surface area is sufficient to avoid overheating the cells when there is no wind, this provides quite good cooling under more typical conditions.

- ***Fin Height:***

In natural convection the velocities are low enough that velocity pressure is not significant, and the air between the fins is in laminar flow, so the friction and the buoyant forces between fins both scale with the path length that the air travels. If the fin spacing remained constant, taller fins (longer paths) would thus have more area to transfer heat from but no more air to transfer it to. To rebalance

the heat transfer, the fin spacing scales with the fourth root of the path length (Bar-Cohen, et al., 2003).

With wider fin spacing, the heat must be conducted farther in air, so this decreases the heat transfer coefficient with the fourth root of the path length. With a wider spacing also decreasing the *number* of fins roughly with the fourth root of the path length (since the fin thickness is small compared to the gap), the heat transfer at a given temperature only scales with the square root of the fin height.

Taller fins also have farther to conduct the heat within the fin, and fewer fins also have more heat per fin, so the fin thickness grows significantly for taller fins. Since more fin area needed due to the lower heat-transfer coefficient, the total fin mass grows rapidly. Thus, taller fins quickly reach both diminishing returns on cooling performance and escalating costs of fin material.

- ***Fin Width:***

Fin width does not normally increase the path length of the vertical air flow between the fins, and so does not normally affect the heat transfer coefficient or the fin spacing. *But the orientation of the fins changes as the trough tracks the sun, so what starts out as fin width becomes fin height during tracking*. With simple fins this becomes significant when the sun is still high and the heat is near it maximum, so wider fins can only be pushed slightly farther than taller fins.

On one side of the module, wide fins would also shade the trough. This happens at the mirror's off-axis offset, which for RP-3 mirrors is ~40 mm, minus the spread from the half-angle of the sun, which for an RP-3's focal length is about 8.5 mm. With an RP-3 mirror, rectangular fins can thus only project ~30 mm on one side of the module center before they start shading the mirrors. But the fins do not have to be symmetric on the module; cantilevering the fins on the other side just pushes the mirror rows slightly further apart. Nor do the fins have to be rectangular; folded fins, for example, can be folded from sheet patterned to eliminate the trough-shading corner, or even to produce trapezoidal fins that can be wider at the top without shading the trough.

4.5.3 Monolithic Aluminum Heatsink Lid

An all-aluminum heatsink has a higher CTE than the glass lens tile, but if the heatsink is tempered to prevent plastic deformation during thermal cycling and is bonded to the lens tile at higher than the maximum operational temperature, the CTE-mismatch stress will remain compressive. A 1.5

mm to 2.5 mm thick fin base is targeted to match the overall module CTE to galvanized steel supports, which will also be high enough to allow a copper receiver substrate (**Figure 4.19**).

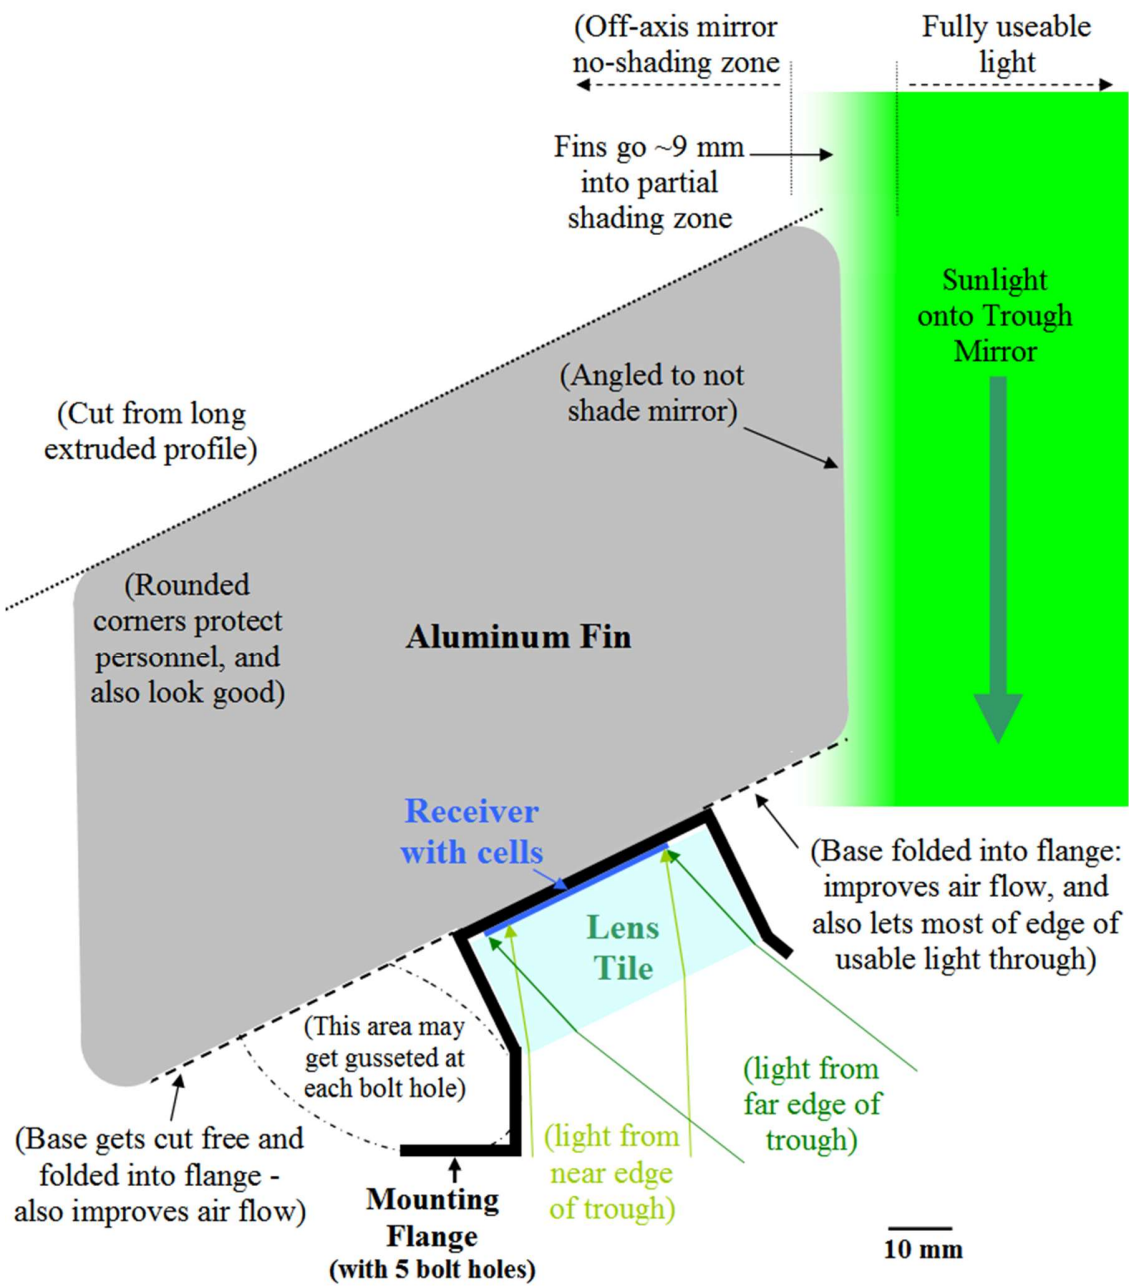

Figure 4.19: *Sketch of Monolithic Extruded Heatsink Lid.*

The target heat-sink lid is a one-piece aluminum part, which could be either extruded or skived. Extrusion is most common at the TLC fin thickness (which is near the upper limit for skiving), and extrusion is less expensive than normal skiving, so extrusion is considered first.

- ***Extrusion:***

Practical extrusion width is limited to <<1720 mm, so multiple widths will be joined together. While special extruders can go as wide as 600 mm, most extruders are limited to ~350 mm and there is much more competition at 200 mm and below, so 9 widths, and thus 8 welds, is costed. Each heat-sink requires 140 mm per weld, for 1120 mm of welding (or brazing) per heat-sink.

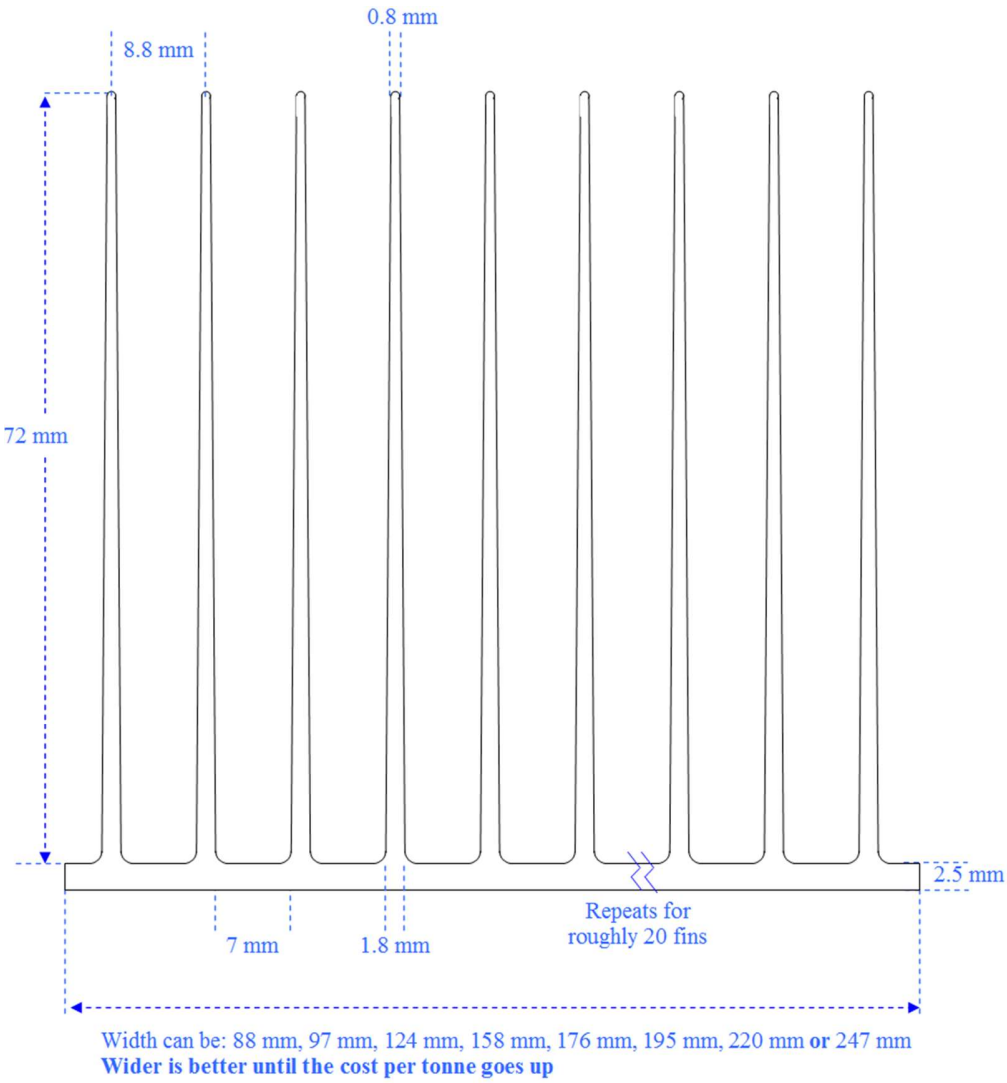

Figure 4.20: *Extrusion Profile for Monolithic Heat-sink Lid.*

- ***Post-Extrusion Processing:***

  - *Welding or Brazing:*

Aluminum takes more power to weld to a given depth than stainless steel does, but even so a 3 kW multi-mode laser can weld 3-mm aluminum at 2540 mm (100 inches) per minute (Amada Miyachi America, Inc., 2016), or about half the speed attainable in welding stainless steel. At this speed 1120 mm of welding on long straight lines is <½ minute of beam time.

A cost study on TRUMPF's TruLaser Robot 5000 series automated robotic welder provides a reasonable approximate of the cost. The module length uses the longest (2000 mm) dimension, leaving only 1000 mm for the weld length; this fits 6 heat-sink widths with spare space for fixturing (over 100 mm extra after accounting for the slanted cuts ion the fins). Even 10 MW is ~10,000 identical parts, so some automation is useful immediately. A conveyor can carry fixtured assemblies to the welding chamber (and carry finished parts out) to minimize non-welding time, so 0.5 minutes is allocated for the assembly exchange. 6*1120 = 6720 mm of weld is 2.65 minutes of beam time, so the total laser welder time is 3.15 minutes for 6 heatsinks.

A TRUMPF cost study shows €149 as the average hourly cost in Germany (considerably higher than laser machining in the U.S., which is typically $110 to $120 per hour (Olny, 2014), but then the robotic welder may be more expensive). The welding is automatic, so the operator can fixture the next assembly during the welding (placing six straight 1-meter-long parts in a jig, and pulling two clamping levers, can easily be done in 3 minutes). The welding cost is therefore €7.82 or $9.15 per 6 modules, or $1.53 per module at current exchange rates (which is 0.145 ¢/W).

***Segmented heatsink*:** Rather than join the extrusion widths, CTE-mismatch effects can be minimized by using a one-piece alloy-410 stainless lid and attaching multiple extruded-fin segments. This can be done with a flexible epoxy with an acceptable thermal penalty.

  - *Cutting:*

A heat sink then needs to be cut free from the welded heat-sink block, or its segment cut from the raw extrusion. This is a simple cut that can be automated (it is cut on a 25° slant to maximize fin area without shading the mirrors, but that doesn't affect a circular saw much), and the depth is only

~75 mm. A minute per heatsink is still allowed for parts handling in low volume. In high volume, gang sawing would reduce the handling to one minute for six heatsinks.

The extrusion base is then cut free from the fins except for the part that will go right over the receivers on the lens tile. Cutting two straight cuts (which can be made at the same time), totaling 100 mm deep with only 20% metal fill is ~350 cm^2 of total kerf). With no requirement for high accuracy or smooth edges, and with proper tooling, this should take < 1 minute per heatsink. Saws are relatively low cost, and at $60/hour machine time this total $1 per heat-sink, and it is highly automatable in high volume to further reduce the cost.

In low volume one or two identical extrusions could be used, with a fin or two cuts from each end to leave material for an end flange (9*22 fins is 198, or four more than the number needed to match the rows of cells). In large volume separate 'left end' and 'right end' extrusions would be used; at ~$10K per extrusion die this is affordable even at 10 MW, saving time and materials.

- *Folding:*

The freed sections of the base are then folded to form the two flanges. This will use a custom folding die that tightly clamps the key area of the lid so that it is not distorted, so folding is reduced to placing a lid in a jig and pulling a lever, and a hydraulic press then makes all folds at one; this should also be < 1 minute. The sharpest corners might also be trimmed or otherwise rounded (for installer safety, and it looks better). Add a minute for corner trimming done in two passes with a shear that trims all fins in parallel (and can also punched the mounting holes and module terminal-wire holes), and it total 2 minutes of low-cost machine time. At $60/hour that is $2.

- *Material:*

Each fin is currently 72 mm tall, 140 mm long, and tapers (currently linearly) from 1.4 to 0.8 mm thick, for 11.1 cm^3 of aluminum, so 194 fins is 2151 cm^3. The base is 140 mm x 1750 mm (including end flanges), and 1.5 mm thick for 368 cm^3, so the total is ~2550 cm^3 (allowing for corner radiuses). At 6063's density of 2.69 g/cm^3, that's at most 6.86 kg of aluminum. Extrusions are typically made in lengths of up to 5.8 meters, so the ~30 mm lost to the slanted end cuts is 0.5%, so 6.9 kg of extrusion is needed per heat-sink lid. Similarly simple extrusions are available at ~$0.70/kg above raw aluminum (currently $1780/tonne), so that's $2.50/kg or ~$17.15 of aluminum extrusion, or ~1.6 ¢/W.

- *Tempering:*

Tempering is desirable to prevent the aluminum from plastically deforming on a cold winter night and then putting tensile stress on the glass when on-sun in the summer. As-extruded "O" temper 6063 is almost acceptable as is, and would be easy to form, but aluminum is usually quenched as it is extruded (partially tempering it) and is then usually tempered to T5. T5 is overkill for a TLC lid's thermal expansion, but at T5 temper the heat-sink is more robust and helps the lid serve as a structural element, and T5 also slightly improves thermal conductivity (to at least 201 W/mK, and some references say 209 W/mK), so T5 is an excellent final temper.

However, T5 only has "good" formability instead of the "excellent" formability of O temper, so while "good" formability might be good enough, it also might be necessary to untemper (from either quenched or T5), and then re-temper. If un-tempering is needed, heating to 345 ºC is enough to partially remove tempering, a few hours at 415ºC and then slow (30ºC/hour) cooling to 260ºC can return it to "O" temper (Atlas Steels, 2013). Re-tempering to T5 is a simple aging process (60 to 90 minutes at 205ºC). >30 kW of lids fit per cubic meter without nesting the fins, and >50 kW/m^3 if pairs have fins interdigitated, so even a 1 m^3 oven could support 100 kW per shift, and even a 1 GW/year plant running two shifts would only need ~20 m^3 of tempering oven.

So re-tempering may not be needed, and if it is it is roughly at most a minute of handling time to rack or hang a pair of lids for transport to the oven, and oven time is essentially free. Handling time is low-cost time, so ½ minute per lid at $30/hour is $0.25 per heatsink.

- *Total Cost:*

The total cost with a 1.5 mm base thickness is thus ~$(17.10+1+1.17+1.45+0.25) = $26.50 per module, or ~2.5 ¢/W.

- *Accepting Slight Shading:*

For the DAT version of TLC, the fins always point at the sun, so if the base is cut free to form the flange, then the fins cast only a partial shadow. The taper of the fins means that only the 1.4 mm width of the fin base counts toward shading, and that is only ~16% of the 8.85 mm fin pitch. RP-3 mirrors are 40 mm off-axis, of which the module's half-width consume 18 mm; and the sun's image widens 8.5 mm in each direction from focus to mirrors, leaving 13.5 mm of extra shading-free width, or 15 mm of flange. The desired flange is ~25 mm, so the fins for the extra 10 mm of flange project 9 mm into the edge of the usable light.

The useable fraction of the light starts at 0% and end up at almost exactly 50% (9 mm being almost exactly half the sun's image's width from an RP3 mirror). The light increases slower at the leading edge, so the average is considerably less than half this, so using a 25% average useable light intensity in this shading zone sets a pessimistic upper bound. So, the fins shade 16% of the light in a 9 mm zone in which less than 25% of the light is usable anyway, and the total light lost is <4% of 9 mm or the equivalent of <0.4 mm of full light. On a mirror with >1500 mm of light gathering width, this is < 0.03%, which is an acceptable violation of the no-shading goal.

(Even if the fins were slanted enough to cast full shadows (due possibly to a shearing process to free the base from the fins), the loss of <0.15% of the total light would be acceptable, although then a slightly shorter flange might be an optimal compromise.)

- *Impact of Future Cells on Finned-Cooling Costs:*

While the general impact of higher efficiency cells is to proportionately reduce all pre-inverter costs (module, mirror, tracker, labor, land and field O&M, but *not* wiring runs to the inverter, inverter, or grid interconnection), higher efficiency has an additional beneficial impact on cooling costs due to having less heat to reject even on a per-module-area basis. For many cooling methods this is linear, but due to the poor scaling of the fins it is quadratic for the amount of aluminum needed for cooling with fins directly on the module lid. Thus a 50% cell being 13.6% more efficient than a 44% cell should reduce cooling aluminum ~25% as well as increasing output 13.6%. With fin aluminum being ~15% of the total module cost at 1 GW, this would be a nice boost.

Higher-efficiency cells will have more junctions, which competes for some of the same benefits as microcells, so a '50% cell' would probably boost TLC efficiency by only ~~8%, so the actual savings would be ~15% of the cooling aluminum, or almost half a cent per Watt in addition to the overall 8% cost/W reduction from the higher module output.

4.6 Receiver Module

A Receiver Module (typically referred to simply as a 'module') comprises two major sub-assemblies: a populated lens tile, and its heat-rejection housing. These two subassemblies are mechanically bonded together and hermetically sealed with an adhesive such as structural epoxy.

These are bonded at roughly the maximum in-service temperature (bonding at ~90°C) so that the higher CTE of the cooling subassembly and the strength of the adhesive ensure that the module expands and contracts as a whole and that the seal remains hermetic.

With no moving parts in the module, all electronics in a glass/metal sandwich, and only glass and metal exposed, a separate cover glass is not needed. This gives DAT TLC a ~5% efficiency boost over SAT TLC, which along with removing the micro-tracking and pumped cooling cost makes up for the higher cost of dual-axis tracking.

4.6.1 Lens Subassembly

The lens subassembly contains the optical and electronic heart of a TLC module. It is produced by taking a lens tile (19-mm roll-formed low-iron glass sheet cut to module-size, with Winston cones then over-molded on back), and then placing populated receivers on it. Placing the receivers on the lens not only minimizes X and Y uncertainty in the lens tile placement, but also the Z-uncertainty relative to the lens.

The receivers will be placed cells-down, and pick-and-place alignment passes the part being placed over a camera that checks the bottom of the part, so the cell array itself can be aligned rather than the receiver exterior. The Z-height uncertainty is moved to the back of the receiver, where even at a mediocre thermal conductivity of 1.5 W/mK, an extra 50 μm after the thick copper has spread the heat to ~2 W/cm^2 would cost only ~0.7ºC.

- ***Dispensing:***

Three different coupling agents are currently specified: a refraction-matched optical coupling for coupling the cells to the cones, an electrically conductive adhesive for electrically interconnecting the receivers, and a mechanical adhesive for securely attaching the receivers to the lens tile.

- *Optical Coupling:*

Optical coupling with a refraction-matched coupling agent (typically optical silicone) prevents an air gap which would reflect light due to s sharp refractive index change. The narrow (cell width) area minimizes voids. The cones will be pressed against the cells so the optical coupling can be extremely thin. Silicone will optically couple to glass under even slight pressure, so a separate operant optical coupling agent may not even be needed.

The optical coupling could be stamped onto the cones, but if it is stamped onto the cells then over-stamping is harmless, especially if allowed to partly cure to avoid low-viscosity issues. In high volume it can be transferred to all cells in parallel with an array stamp. In low volume the optical coupling can be jetted onto the cells using readily available nanoliter dispense equipment. For example, the Hernon Sureshot 3500 (Hernon Manufacturing Inc., 2018) jet dispensing valve can dispense adhesive volumes ranging from 0.1 nL to 10 nL per shot at 500 shots per minute. Each 1 µm thickness over the cell area is ~0.6 nL, so this allows coupling thickness from ~0.2 to ~20 µm to be dispense in a single shot (and if thicker is desired, other dispense equipment delivers larger droplets).

- *Electrical Coupling:*

Electrical coupling connects the power plane of one receiver to the backplane of the next. In low volume this is probably easier to dispense onto the receivers since these are the size of things that typically have conductive epoxy dispensed on them, but in high volume this could be dispensed on all inter-receiver thermal contacts of a lens tile in parallel. Placing the receiver will squish this into a thin layer – the area is large so it does not have to be a high-end conductive adhesive, and it will work even if full of voids.

- *Mechanical and Thermal Coupling:*

A structural adhesive is used to mechanically bond the receiver to the lens tile and to conduct heat from the lens to the receiver. This adhesive can be dispensed on either the lens or the receivers, whichever is more convenient. Again, in high volume this could be dispensed on all of a lens tile's intra-receiver thermal contacts in parallel.

This adhesive does not have to produce a rigid bond – even with a flexible coupling the CTE mismatch of 3 ppm/K for 60K is only a maximum shift of less than 7 µm, which the silicone cones could easily flex to accommodate. Nor does it have to be flexible – even a rigid bond would put only a few percent of the tensile strength of the glass on the lens during temperature cycling. The rigidity is expected to match that of the thermal gap filler attaching the lens assembly to the housing; this was expected to be a fairly rigid bond because it will help tie the lens to the module housing so that it all contracts and expands as a unit.

Although this adhesive in in the thermal path for cooling the lens, it does not have to be very thermally conductive. The lens only absorbs a few percent of the trough's 4 W/cm^2 focus and has

roughly half of the focal area to conduct through; at ~0.2 W/cm^2 even a 100 µm layer of household epoxy at 0.13 W/mK would add less than 2K to the lens temperature, and a filled mechanical adhesive like JB Weld would add only ~0.5K.

The mechanical adhesive does not have demanding requirements, so it can be quick curing to minimize storing of uncured modules. Ordinary few-minute structural epoxy is suitable.

- ***Dispense and Flip:***

If the dispensing is onto receivers (likely in low volume), a module's worth of receivers can be placed in a tray for the dispensing. A cover with an opening for each receiver is then placed on the tray to hold the receivers securely in place. After the dispensing is completed, the tray is flipped over and its back removed, allowing a vacuum tip to pick up the receivers by their backs.

- ***Receiver Placement:***

The lens tile will already have the interconnect strips in place and any dispensing will have been done. Each receiver will already be populated with wire-bonded cells and a diode and will have been pre-tested under a flash tester, so each receiver will be ready to incorporate into a module.

A receiver is the size of a large packaged chip like a CPU (~40% of the area of a business card), and it will have its flat back exposed so standard pick-and-place equipment can use a vacuum tip to pick it from a tray a placed it on the lens tile. This is a fast and accurate process already used to handle billions of similar-sized chips per year at low cost (much lower cost than placing a much larger and far more fragile silicon wafer in a flat panel, and with ten times more Watts per placement than a wafer, too).

- *Z-Stop Stand-off:*

The thermal contact area between the cone rows will serve as a stand-off to set the lens height relative to the receiver. For highest accuracy stand-off ridges can be rolled in between the rows and those areas left epoxy-free to provide hard stops. With a receiver pressed against several stops, the uncertainty in height will be very small, and excess optical, electrical and mechanical coupling agent can be avoided, preventing adhesive from going where it is not wanted.

- ***Placement and Curing Temperatures:***

Ideally the modules are receivers are placed and cured at an elevated temperature to minimize and cell stress on the lens during high-temperature operation. However, heating after placement is problematic because a receiver has a CTE ~ 8 ppm/K higher than the lens glass (before the heat sink constrains the glass), so temperature change would produce a ~300 nm/K shift of the corner cells relative to the lens. Even a 30º change in temperature is thus only a few microns, but huge changes should be avoided. A reasonable compromise is a room temperature or slightly warm placement temperature (e.g., 30ºC to 40°C), and a cure temperature of 50ºC to 80ºC.

- ***Finishing the Lens-tile Subassembly:***

Once the mechanical adhesive cures to handling strength, the lens-tile subassembly can be handled, the other adhesives do not need to be fully cured, as surface tension and viscosity will hold them in place if the tile is moved. Minimal sub-assembly finishing is needed:

- Adhesive squeezing between receivers can be avoided with proper dispensing, but this can be checked and any adhesive squeeze-up removed.
- The lens subassembly can also be flash tested if desired.
- Standoffs could be used between the lens and the housing to ensure enough isolating adhesive thickness to provide the desired level of electrical isolation, *currently this is not needed because a layer of thermally conductive prepreg supplies the isolation*. However, if a fairly thermally conductive epoxy is found that is lower cost than the prepreg, then a stand-off could be as simple as a narrow strip of Kapton tape run along each side of the back of the assembly with many epoxies providing ~30V/micron of isolation, 4-mil tape would ensure 3000V and 6-mil tape would ensure 4500V of isolation. Thicker stand-offs could be used for lower-dielectric-strength adhesives.

- ***Module Connectors and Wires:***

  - *Connector Wires:*

The wires to the module connectors pass into the lens where the side of the metal back seals to the lens. The wires have weatherproof elastomer seals (like NEMA enclosure cord seals) compressed by aluminum tabs to protect them from off-focus sun; these conduct the heat to the heat sink in a path only a few millimeters long).

A wire is inserted through a groove in the lens and into the connector cage on an end substrate. To keep the size of the groove in the lens small, 14-gauge Exane wire is used. Exane wire is "ideally suited for rugged transportation industry applications"; its coating is highly abrasion resistant and is weatherproof enough to be used in marine installations, and its higher-quality isolation is thinner (4 mm diameter) than normal solar wire, and its fine strands (19/27) bend easily. It also has a high temperature rating, which is reassuring near a trough's focus even if it is always in the shade of the cooling system. The Exane wire is only 27 cents/foot, so the 400 mm (200 mm on each end) costs only 0.03 ¢/W.

- *Connectors:*

The connectors are ordinary weatherproof solar connectors, keyed to ensure correct connection. For 1000V systems these can be MC4 at 42¢ per pair (Quantity 500, Shanghai Jiukai Wire & Cable Co.). For 1500V, the 2017 Mouser catalog has TUV-certified Helios H4 connectors at $2.74 per pair for quantity 10 (a drop of 15% in price from the quantity 1000 price earlier in the year); for a 1.05 kW/module that is only is 0.26 ¢/W (at that is highly pessimistic since price drops with volume and even 10 MW uses a thousand times that quantity).

4.6.2 Bonding the Sub-assemblies

The module lid has already been discussed under 'Heat Rejection'. Regardless of which style of finned lid is used, the lid is shaped to accommodate the lens sub-assembly, including having notches for the module's terminal wires.

Bonding the lens subassembly to the lid involves two bonds: the lens glass to the metal lid, and the receiver backs to lid (or to an isolating layer thereon).

- ***Lens to Lid:***

The lid bond not only hermetically seals the lens to the metal lid, but mechanically couples them so that they expand and shrink together with temperature. Curing the adhesive at near the maximum operating temperature ensures that the only significant force that the lid puts on the lens is compressive, and glass is very strong under compression. Thus, the lens should not need strengthening (although chemical strengthening would be affordable at the tiny per-kW lens area).

- ***Receivers to Lid:***

The receiver backs need thermal conductivity to, but electrically insolation from, the metal lid. With heat spread to ~2 W/cm^2, moderate thermally conductivity is acceptable. The isolation needed is 1000V plus twice the string voltage, so the isolation needed depends on whether the module is for 1000V or 1500V strings.

A sheet of thermally conductive prepreg (used for insulated-metal PCBs) looks ideal. Ventec's VT-4a2 has a DC withstanding strength of 4500V for a 75 µm thickness, a thermal conductivity of 2.2 W/mK, and costs only \$13.02/m^2, or 0.074 ¢/W. This could be applied to either the back of the substrates (before cell placement) or to the inside of the module back cover (expected).

A thermally conductive compliant layer is still needed on the other face of the prepreg, but it does not have to have high dielectric strength. If the prepreg is applied to the back of the substrates, and any gaps are sealed (can be with ordinary epoxy or silicone), then even an electrically conductive gap filler could be used. Prepreg takes longer to cure than a thermoplastic adhesive takes to melt and cool, so prepreg might be more expensive until bonding is automated.

Any warps over 50 µm (none expected) would be filled with grout, so only an average of 25 µm is needed or ~1.5 cm^3 of epoxy, or about 3.3 grams per module at a density of 2.2 g/cm^3. Epoxy at 1.6 W/mK costs ~\$120/kg in North America (Master Bond), or 0.034 ¢/W. This is slightly pessimistic because 1.5 W/mK is advertised at ~\$80/kg in 10 MW quantity (several suppliers in China) or 8¢/gram, which would save 0.01 ¢/W if the risk of epoxy from China is accepted.

4.6.3 Module Rail

The mounting flange is designed to bolt to an angle iron that supplies sufficient stiffness to support the module under the maximum wind or snow load. This allows a location-dependent angle iron size to be used rather than designing all modules to support a massive ice load. It is expected that one long rail will support multiple modules.

4.6.4 Module Vertical Supports

The off-axis trough allows simple vertical supports since a vertical support will not block any sunlight. A focal length plus bolting length support per mirror is allocated, plus one for the far end

of the last mirror in the row. Supports can also be guyed at the top or even in the middle if needed to prevent buckling. This allows standard profiles like 1" x 1" x 1/8" angle iron to be used; at one per module, this is 0.23 ¢/W. Wiring to the inverter can run in the shadow of a support for protection from the focused sunlight when the tracker is tracking the modules into the sun, and the connectors sit on the module back shielded from the focus.

4.6.5 Module Length

RP-3 mirrors are 1.7 meters long, and 1.7 meters is a good module length. Longer modules would be awkward for one person to handle, and shorter modules would require more connectors and supports, raising costs. A 1.7-meter module is just over 1 kW, or roughly three times the power of a top-of-the-line silicon panel and four times the power of a typical silicon panel. At ~~20 kg, the 1700 mm x 110 mm x 150 mm (mostly fins) module can be handled by one person if needed, although two people is preferred.

4.6.6 Module/Trough Ends

While 2-axis tracking eliminates the large seasonal spill zone at any interruption in the continuity of the mirrors along the trough (including the ends of a trough), unless the mirrors are tightly packed end-to-end in a given trough there will still be a small 'spill zone' above the gap between mirrors produced by the sun's non-zero diameter and any tracking errors. If mirrors can be placed just a few millimeters apart this can be small enough to be ignored except at the trough ends, and the trough-end strategies can be applied any for bigger inter-mirror gaps as well.

At a 1710 mm focal length the ~17 mrads allocated for the sun's image along the trough is ~30 mm, so in the spill zone between individual trough mirrors the intensity of the trough's focus tapers off over this distance. Since 17 mrads requires is all errors at maximum and aligned the same way, almost all the tapering occurs within 10 mm to either side of the mirror length. There are several ways to deal with these small spill zones:

- ***Identical Receivers for Tight Mirrors:***

Tightly pack mirrors only have spill-zones at the mirrors-row ends. Since the intensity tapering in a spill-zones all falls within one receiver, one way to deal with tapering of the focal intensity is to

keep all receivers identical and to accept the loss of one receiver's output on each end of a row of modules. For a 9-meter tracker with five modules per row this is only a ~1.25% loss.

- ***Identical Receivers for Mirrors with Gaps:***

If there are gaps of more than a few millimeters between mirrors, the spill zones would be significant enough to miss-match receiver photocurrents and losing two receivers at each such spill zone would be a ~6% loss. Centering a receiver over the gap would cut this to a 3% loss, but why put any receiver where its power is almost always lost? Shortening the module ~15 mm on each end brings the end receivers out of the spill zone, while reducing power by just ~1.8% regardless of mirror gaps or tracker size.

- ***Boosted-ends Module:***

However, the spill zone still has enough light to be worth scavenging. The tapering off of the light almost all occurs within +/- 10 mm (from above the mirror edge), so one extra row of cells is almost enough to cover for the lost light. Shifting the receivers inward adds full-strength light while foregoing the weakest edge, so a shift of less than a millimeter makes up the difference.

This only requires shrinking the receiver width by ~50 um, and an extra cell row on each module end is ~1% more module, which costs much less than a corresponding reduction of module power (and thus 1.25% of module, mirror, tracker and installation cost). The extra rows could even have only the 14 or even the 12 central cell sites populated to reduce the added cell cost. The remaining bits of un-captured spill-zone are the least-illuminated parts, totaling only ~~0.1% of a mirror's light. The boosting could be an extra row on each end receiver, but a separate single row 'receiver extender' is expected to be more convenient.

- ***Complementary Options:***

These module end approaches can be complementary. On a tracker where mirrors are close and aligned anyway, identical-receiver modules could be used for all intermediate modules and boosted-end-receiver modules used only at the row ends. Even using modules with both ends boosted, this cuts most of the extra cost, and in large volume left and right ends would reduce the end-loss compensation cost to two rows row of cells per tracker row, or about 0.2% of the receiver-module cost (but not the mirror, tracker or installation cost) on a mid-sized tracker, and a mere ~0.1% of the receiver-module cost on a large-heliostat-sized tracker.

- ***Module-end Option Costed:***

The boosted-ends option is used in the costing because it has more flexibility and is expected to be used initially. Its receiver size can be used in identical-receiver modules while only foregoing ~0.1% of the potential power, it is independent of the tracker size, and it does not require closely packed mirrors on the tracker, and it is the more efficient of the two options.

4.7 Dual-Axis Tracker

TLC is designed to be tracker-agnostic so that TLC can pick the most cost-effective tracker. While the DAT version will thus fit on almost any CPV-grade two-axis tracker, the tracker is such a large part of the costs that it merits discussion (it dominates the pie charts, obscuring information that is often more useful; therefore, it can be included in or excluded from the percentage summary and pie chart with a switch: 1 for include, and zero for exclude). In addition to some discussion of current trackers, potential directions for customized future trackers are also briefly discussed.

4.7.1 Mechanical Fit

While TLC modules and mirrors can be adapted to any reasonable length, using a standard RP3 1.7 m mirror makes multiples of 1.7 m ideal for a tracker. Many trackers have a 9 or 10 m length, matching five or six full-mirror TLC module, and a 5-to-7 m width, matching three or four row of modules. Mechanical fit becomes even easier on larger trackers, and heliostats appear to be headed to a roughly 10 m by 15m size, matching six rows of eight or nine modules each.

Smaller trackers are rare, but a potential Canadian cooperator (Morgan Solar) has suitable CPV trackers that could efficiently accommodate two rows of three modules each. Other small trackers might require custom mirror and module lengths (both minor annoyances that could be accommodated at 10 MW) or mirror widths (a moderate annoyance but still feasible at 10 MW), or mirror focal lengths (affordable at 100 MW with commercial glass mirrors).

4.7.2 Electrical Match

- ***String Inverters:***

String inverters still cost slightly more than central inverters, but the gap is closing rapidly; Huawei (now the world's largest inverter manufacturer, and making both central and string inverters) reports that their string inverters save 2.5 ¢/W over central inverters in non-inverter costs (Huawei, 2017), so the total costs are already close to even.

A string inverter such as the SUN2000-15KTL has three separate MPP-tracking inputs, which would match the reference Trackers Feina tracker's three rows of mirrors/modules, thus making the TLC module rows relatively insensitive to partial shading, especially with a dense N/S, sparse E/W field layout.

- ***Partial Shading Tolerance with Central Inverters:***

While the SAT version of TLC is almost completely immune to partial shading (since long troughs shade all modules in a string equally), the DAT version would normally have similar shading sensitivity to flat panels or Fresnel CPV. However, if the bottom rows on trackers are in series (for small trackers) with the bottom rows on other trackers, and the second rows with second rows, etc., then shading sensitivity can be greatly reduced. For a 5-module-per row tracker such as the one that the costing is for, feeding a 600V inverter matches 1 tracker, a 1000V inverter matches rows of two trackers, and a 1500V inverter matches rows of three trackers in series. For 5J cells one tracker matches a 1000V inverter and two match a 1500V inverter.

- ***Module Electrical Matching:***

The simplicity advantage of one module row feeding an inverter is worth customization. In the DAT version of TLC, changing the number of cell rows per receiver does not *require* any other changes because neither the lens nor the cooling is even aware that there are multiple receivers. Even full optimization is trivial, affecting just the diode size and bond-wire diameter.

The number of rows of cells per receiver can thus be easily optimized for a given combination of tracker and cell type. For the initial lens with 8.8 mm lens rows, the tables for 3J (Azur 44%) and 5J (Boeing Spectrolab) cells are shown below. Appropriate inverter voltages (at the high end of the MPP range for a typical inverter) are highlighted in light olive for 600V inverters (becoming rare at utility scale), dark olive for the 1000V inverters, and bright green for 1500V inverters:

Cell Vmpp at temp	Rows per Reciveret	1	2	3	4	5	6	7	8	9	10
2.70	Volts/Module	519	260	173	130	104	87	74	65	58	52
	Amps	`	4.0	6.0	8.0	10.0	12.0	14.0	16.0	18.0	20.0
Modules per row	2	1039	519	346	260	208	173	148	130	115	104
	3	1558	779	519	389	312	260	223	195	173	156
	4	2077	1039	692	519	415	346	297	260	231	208
	5	2596	1298	865	649	519	433	371	325	288	260
	6	3116	1558	1039	779	623	519	445	389	346	312
	7	3635	1817	1212	909	727	606	519	454	404	363
	8	4154	2077	1385	1039	831	692	593	519	462	415
	9	4673	2337	1558	1168	935	779	668	584	519	467
	10	5193	2596	1731	1298	1039	865	742	649	577	519

Table 4.1: *Voltage Matching to Inverter for 3J cells.*

Cell Vmpp at temp	Rows per Reciveret	1	2	3	4	5	6	7	8	9	10
4.78	Volts/Module	917	459	306	229	183	153	131	115	102	92
	Amps	1.2	2.4	3.6	4.8	6.0	7.2	8.4	9.6	10.8	12.0
Modules per row	2	1835	917	612	459	367	306	262	229	204	183
	3	2752	1376	917	688	550	459	393	344	306	275
	4	3670	1835	1223	917	734	612	524	459	408	367
	5	4587	2294	1529	1147	917	765	655	573	510	459
	6	5505	2752	1835	1376	1101	917	786	688	612	550
	7	6422	3211	2141	1605	1284	1070	917	803	714	642
	8	7339	3670	2446	1835	1468	1223	1048	917	815	734
	9	8257	4128	2752	2064	1651	1376	1180	1032	917	826
	10	9174	4587	3058	2294	1835	1529	1311	1147	1019	917

Table 4.2: *Voltage Matching to Inverter for 5J cells.*

It is NOT necessary to match a module to a highlighted cell, it is just a nice-to-have that becomes worthwhile for a large utility-scale installation. Especially at small quantities, factors like matching a common 12A diode from silicon panels are more important, and the impact of mirror imperfections becomes increasingly significant with fewer rows per receiver.

4.7.3 Dual-Axis Tracker Cost

Dual-axis trackers were so expensive that TLC was originally designed for single-axis trackers to avoid chicken-and-the egg dilemma of needing large volumes to bring the cost down before one can afford large volumes. However, as power towers have taken CSP market share from troughs, the two-axis trackers needed for heliostats have come down ~3x in cost in 8 years, which is fast enough to remove the dilemma.

Current CPV-grade single-axis tracker and dual-axis tracker costs are from 10 MW quotes from a supplier (Trackers Feina) with ~100 MW of trackers in field. Dual-axis tracker costs have already fallen enough that the cost-per-Watt-including-tracker are roughly equal; the higher capacity factor gives DAT TLC the edge on LCOE, and less of the cost being the module gives DAT TLC a significantly higher profit margin.

- ***Future Heliostat Dual-Axis Tracker Costs:***

The DOE targets heliostat (tracker plus mirror) cost at $75/m^2 in 2020 (Mehos, et al., 2016), so this is used for 1 GW. Based on a 4:1 ratio of tracker cost to mirror cost (from current quotes as well as from the DOE), this would be a $60/m^2 tracker. Since little learning is left and the basic glass is already produced by the 100 GW equivalent, the non-learning scaling curve is used thereafter.

This is quite pessimistic for several reasons:

- The DOE solar field targets of $75/m^2 in 2020 and $50/m^2 in 2030 also *include the installation of the heliostats*, whereas silicon-module tracker costs used for comparison do not include the racking or the tracker installation costs.
- TLC tracking is easier than for heliostats, so cost should be significantly lower. The drive mechanism is the largest heliostat cost, and improved calibration is the key to reducing the drive cost (Shemer, 2017).
    - Because heliostat tracking cannot be calibrated easily, heliostats must maintain calibration without feedback. In contrast, TLC tracking can get instant feedback from the sun's position at any time, eliminating the calibration challenges.
    - TLC's optical error target is also 1.7 times more relaxed in altitude and 2.5 times more relaxed in azimuth than the DOE's heliostat targets, allowing lighter structures and drives.
- Even at 100 GW the manufacturing scaling is still more than the DOE's 2030 target.

The DOE costs are thus pessimistic for TLC, but they are referenceable and affordable, and hence are good enough to use until a purchasing agent can beat them.

4.7.4 Space Frame

The vertical supports hold the module rails about 1750 mm from the rails that support the mirrors. With two angle-iron structures 1750 mm apart, it is very easy to tie these together into form a space frame. Although at the scale of a typical two axis tracker space frames are normally slightly more expensive than a central torque tube and tapered beams, the verticals, the diagonals, and the horizontal members in one direction are needed anyway so using the module supports to form a space frame is expected to be the lowest-cost way to form a structure to hold the mirrors and modules in place. The basic design details have been worked out far enough to assess effects such as the thermal expansion of the module rails relative to the mirror rails, which have been calculated and are minor for space frames up to the size of a typical tracker.

4.7.5 Floating Azimuth Two-axis Tracking

Pyron Solar (2004) pioneered a CPV system where arrays of one-axis (altitude) tracked CPV modules are linked together in a circular structure that floats, allowing the second axis (azimuth) tracking to be provided through simple rotation of the floating structure. Novaton extended this "floating" principle to a structure that floats on an air cushion, which is more applicable in deserts where water is scarce. In either case, the floating principle can potentially provide a carousel two-axis tracker for little more than the cost of the traditional altitude single-axis tracking.

For floating on water this can potentially work synergistically with using the module supports (rails and verticals) as part of a space frame because this provides very high strength and stiffness, allowing a large structure to ride out significant waves while maintaining accurate pointing. Because the module rail be warmed by the module while the mirror rails may be cooled by proximity to the water, there is a limit to how long an individual rigid block can be. At a 50° C difference the ~12 ppm/K expansion of steel will expand the top rail by 600 ppm, warping the structure enough to slant the ends by ~ 1 mrad per three times the focal length, or ~ 1 mrad per 5 meters, so individual blocks would start consuming is significant part of the tolerance budget if they are more than ~ 15 meters on a side (~ 75 kW). However multiple such rigid blocks can be linked together to turn as a unit on the azimuth axis.

For land-based floating solar a space frame is less important because the land generally does not move, but the module supports can still provide a convenient lever arm that is long enough to allow driving multiple rows of mirrors without thermal expansion of the drive mechanism being problematic.

- ***Synergy with Desalination:***

By supporting sets of long rows of mirrors that have only single axis tracking within the set, floating azimuth tracking (or any large carousel tracker) makes it much easier to gather waste heat instead of rejecting it to the atmosphere. Thus, floating solar is a natural fit for hybrid systems that use the waste heat such as for biofuels distillation, releasing CO_2 from an absorbent, or desalination. Since areas with good sun for CPV are often short of water, desalination is a logical use of the heat. While desalination can achieve a higher gained output ratio (water output for heat input) from higher temperatures, CPV cells become less efficient at higher temperatures so there is a balance between the value of electricity and the value of water that determines the optimum operating temperature.

4.8 Module Shipping and Installation

Receiver-Modules and mirrors are shipped separately because the mirrors can be shipped directly from the supplier. Receiver-Modules are mounted 'beside' their off-axis mirrors in the field.

4.8.1 Shipping

- ***Mirrors:***

With standard RP-3 mirrors, the mirror plus ceramic attachment pads is less than 2 cm thick so 50 mirrors fit per meter, and even allowing for curvature and a frame, roughly 110 standing mirrors fit across a shipping container. A megawatt of mirrors thus takes 9 rows of mirrors; at 1.6 meters per mirror plus 100 mm for framing this it 1.7 meters per row, so 7 rows fit in a 40' container (or 3 in a 20' container). ~800 kW fits per 40' shipping container, or 1.25 containers/MW (or three 20' containers per MW).

- ***Modules:***

While the active part of the module is very small, the fins increase the size markedly. Including a protective cardboard wrapper, a module takes ~120 mm x ~150 mm x 1750 mm (in cardboard sleeve). 1 MW fits in a standard 20' container, or 2 MW per 40' container. If shipped with the mirrors a container has room for the receiver modules above the mirrors, so no extra volume is needed and is still takes just three 20' containers per megawatt.

- ***Supports:***

Angle-iron supports nest together and hence take up almost no space. Even square tube would be only 0.1 containers/MW.

- ***Total and Comparison to Flat Panels:***

For TLC the total is thus ~1.85 standard 40' shipping containers per MW. For comparison, the best flat panel (SunPower's X21-345) takes 46 mm (unpackaged), for 48 modules per row in a container. At 1046 mm wide and the same 100 mm for framing, a row takes 1146 mm, so ten rows fit a container, for 480 panels of 345W each. That's 165 kW/container, or ~six 40' containers/MW. Even with the commercial RP3 mirrors, TLC thus has a more three-fold shipping advantage over the best flat panels (and a roughly 5x shipping advantage over low-cost 16% flat panels).

4.8.2 On-site Installation

- ***Mirror Installation:***

The tracker is leveled at a comfortable working height, with the mirror rails spaced to where the ceramic pads of the mirrors will go; 20-foot rails will conveniently fit the pads of four mirrors, but longer or shorter rails can be used as appropriate for a given tracker width. The rails need to be parallel to each other and the right distance apart (tape-measure accuracy), but their orientation to the tracker can be off by a few degrees with no harm.

A row of mirrors is then bolted to its rails. The lower pads can be bolted to their mirror rail with only a washer between the pad and the rail, while the higher pads either have a 250 mm support between the rail and the pad, or the rail itself is raised on 250-mm supports.

While all rows of mirrors could be installed before any receiver modules, without fancy staging it would be hard to avoid stepping on the mirrors to install the module supports and modules (the mirrors can take stepping, but it is not good practice). So once a row of mirrors is installed, a row of modules is installed.

- ***Module Installation:***

After the row of mirrors, vertical supports and their diagonal braces are installed. With the tracker level, the alignment of the verticals can be checked with a level. A good carpenter's level is accurate to 5 milliradians, which is borderline accuracy; a machinist's level is accurate to ~0.1 milliradian, which is overkill but well worth the few hundred dollars (on my machinist's level, the bubble being off by an entire 'major line' is only off by 0.4 mrads). A long angle-iron rail (23 feet for four modules) is then bolted to the verticals; predrilled bolt-hole accuracy is sufficient

After a row of module supports and the rail, that row's modules are installed. While a module is roughly the same length as a standard silicon panel, it is roughly eight times narrower and ~35% lighter, so it is easy to handle. The module flanges and the rail both have pre-punched (or pre-drilled) holes, and the module is simply bolted to the module supports – bolt-hole accuracy is fine (a few millimeters in any direction is within the tolerance budget), so as long as the supports are vertical (perpendicular to the tracker), no alignment check is even needed.

Each module is connected to the previous module in the row. The connectors are standard solar connectors and are keyed so that they cannot be connected wrong. The module orientation is known (the module only has bolt holes on one side), so the connector wires are short and thus low resistance. The wires with the connectors are clipped to the top of the module so that they are in the shadow of the cooling as the module is tracked on and off sun. At each end of a row a connector with a longer wire gets connected – this runs down the end-of-the-row module support and thence below the mirrors to carry the power to the tracker post.

Depending on the inverter voltage and the tracker size (see the discussion on trackers), a single row of modules may not reach the inverter's ideal voltage range. One of the few drawbacks of dual-axis-tracked TLC is that it lacks some of the single-axis trough's immunity to partial shading, so to compensate the lowest row of one tracker is connected in series with the lowest row of the next tracker so that their photocurrents match. While it would be more convenient to connect the rows on a tracker to each other, that would be less efficient when the sun is low enough that trackers

shade their neighbors because the lower rows will be shaded long before the highest rows. Even when a single row feeds an inverter, lowest rows are combined in parallel so that their voltages match, covering the case where just a corner of an array is shaded.

- ***Installation Costs:***

Although installing mirrors and modules separately at first sounds like more work than installing flat panels, a mirror is only about 50% heavier than a flat panel and it supplies three times the power, and the modules are trivial to install. TLC takes roughly the same number of bolts per kilowatt as flat panels, and only one third as many connectors clicks. The labor per Watt (or per kWh) is expected to be similar to silicon flat panels, and with the appropriate staging and high-volume TLC is expected to have lower installation cost (outside of the tracker) than flat panels, and a lower total installed cost even including the tracker.

This is cross-checked by heliostat costs, which NREL targets a \$75/m^2 in 2020 (Mehos, et al., 2016); Mark Mehos has confirmed that this is *installed* cost for heliostats, and that as of Nov. 2017 he thinks that it is 'realistic and achievable'. Since RP3 mirrors were only ~\$10/m^2 more than flat solar mirrors in May 2017, and since heliostats have higher stiffness and accuracy requirements and have alignment challenges that TLC doesn't, as well as more complex mirror canting, this sets a realistic upper bound of \$85/m^2 for the installed cost of RP-3 mirrors on two-axis trackers. For ~40% modules that is ~21.25 ¢/W.

At even 10 MW the non-mirror module costs are only 17 ¢/W, and this falls to <11.75¢/W at 1 GW. SunShot is targeting direct installation costs of 10 ¢/W (versus 12.3¢/W today), and with mirrors already accounted for TLC has ~70% fewer modules to install and connections to make than with flat panels (and lighter modules as well), so TLC module installation should be only 30% of that, bringing the total modules-on-tracker cost to ~36 ¢/W for TLC at 1 GW. Costs for wiring to reach the inverter bring this to ~37 ¢/W, and including the inverter brings it to 44 ¢/W.

In contrast, for silicon flat panels at utility scale SunShot is targeting a 2020 BOS equipment (not including modules, tracker or inverter) cost of 25¢/W, along with a tracking adder of 10¢/W and direct installation labor of 10¢/W, for 45¢/W to get the modules installed, and 52 ¢/W including the inverter – *on top of whatever the silicon flat panels themselves cost.*

The lower capacity factor from TLC not using diffuse light is site dependent; in the reference site of KC Mo., it makes the equivalent cost for TLC 58 ¢/W. This is only 6¢/W more that the non-

module costs for silicon flat panels, so TLC at even 1 GW would match the LCOE of silicon flat panels on single-axis tracker if the flat panels were sold for 6¢/W. And that is in mediocre-sun Kansas City, Missouri, and ignoring TLC's lower shipping cost and more efficient use of land and the lower site preparation costs for two-axis trackers than single axis and ignoring the lower degradation of tandem cells and the likely very-long module life of TLC.

4.9 Conclusion

This chapter has walked through the major aspects of the proposed TLC rest-of-module design and has outlined potential solutions to each major design challenge encountered. This chapter has then walked through the proposed assembly process and shown a plausible path to TLC modules that can be easily assembled, shipped and installed. The basis for cost estimations for each component was also presented, which will be used in the following chapter.

CHAPTER 5

5. ECONOMICS

5.1 Introduction

This chapter covers overall economics of TLC. The material costs of TLC were addressed (along with some process costs) in discussion of the design in chapters 3 and 4, so this chapter fills in the remainder of the anticipated costs by walking through the assembly and estimating the cost thereof.

Reliability is also key to long-term economics in photovoltaics, so this chapter then looks at how TLC can resist common failure modes and compares the potential reliability of TLC to the known reliability of silicon PV.

This chapter then summarizes the overall costs and value of TLC. Since the value depends upon location-dependent conditions such as the ratio of direct to overall sunlight, and temperature and wind, this chapter then compares the overall value of TLC to silicon PV in target markets. Since costs are highly dependent upon manufacturing scale, costs are evaluated as a function scale, and this chapter concludes with an analysis of potential limits to production scaling.

5.2 Assembly Walk-throughs and Costs

5.2.1 Components Preparation Walk-through

a) **Receiver Substrates**: Made as panel (~457 mm x ~610 mm) = ~3 kW of substrates
    - Thick backplane copper sheet, power-plane copper and prepreg are cut to panel size (these sheets may be pre-bonded, or bonded after cutting to size)
    - Power plane and prepreg are sawn to depth with oversized kerfs
        - Includes cell rows and inter-receiver contacts
    - Diode area is drilled to backplane
    - The unmasked areas are plated (probably ENEPIG) for soldering and wire-bonding
    - The panel is cut into roughly a hundred receiver backplanes

        - Probably V-grooved and snapped (slightly rough edges OK)
        - The edges might be wiped with ferric chloride to etch away any shorts from the sawing process leaving power-plane kerf copper along the edge
        - Routing and laser-cutting are also options

b) Integral Lens Sheet

- Low-iron glass is roll-formed into lens sheet with lenses and thermal contacts
- Thermal contacts are accurately over molded with hard plastic
- Lens sheet is cut to module-sized tiles (long edges can be rough)
- Lens tile is over-molded with silicone TIR cones
    - Over-molded thermal contacts align cone-molds to lenses

c) Heatsink Lid

- Lid profile is folded it 410 stainless
- Mounting holes (~5) holes for connector wires (2) are punched into lid
- Thermal interface surface is smoothed if needed
- Thermally conductive prepreg is bonded to thermal interface surface
- Fin profiles are extruded and cut every ~ 140 mm (on ~25°) slant into fin segments
- Fin segments have part of base cut free from the fins and folded to form flanges
    - Edges may be deburred, and sharp corners trimmed for looks and safety
- Lid has thermal adhesive dispensed and fin segments are clicked to lid
    - Flanges mechanically hold while adhesive sets
- Base may be gusseted to fins near mounting holes

d) Cell Fabrication

- Normal solar cell wafer has custom top-metal-layer (1 mask) for small cells
- Cell wafer might have solder plated onto back before dicing
- Cells are tested and bad cells mapped or marked
- Wafer is diced
- Bad cells are removed
- Good cells are put into tape-on-reel or another convenient format

5.2.2 Receiver Assembly Walk-through

➢ ***<u>Receiver Electrical Assembly (All methods):</u>***

- Cells and diode are pick-and-placed on each site on receiver substrate
- Solder is reflowed to center and bond cells and diode
- Cells and diode are wire-bonded
- Receiver can be flash tested if desired
- Receiver-extenders for module ends assembled (connector-wire cages instead of diodes)

➢ ***Receiver (and Receiver Extender) Placement on Lens Tile:***

- If adhesive optical coupling silicone is used:
  - Receivers are placed cells-up in tray
  - Very thin layer of optical silicone adhesive is stamped on to cells
  - Tray lattice-lid is closed (lattice only touches edges of receivers to avoid cells)
  - Tray is flipped, and tray back is opened
- Lens tile is placed in holder, cones up
- Viscous mechanical epoxy is dispensed onto intra-receiver thermal contacts
- Electrically conductive epoxy is dispensed onto inter-receiver thermal contacts
  - Each gets two narrow beads to avoid conductive epoxy between receivers
- Small drops of instant adhesive are dispensed
- Receivers are picked and placed (pressed) onto lens tile
  - Light through the lens tile can be used to turn the cones into alignment fiducials
- Standoffs on thermal contacts control cone compression, which instant adhesive holds
- Lens subassembly is allowed to cure
- Lens subassembly can be flash tested if desired

➢ ***Sub-assemblies mated:***

- Lid is placed in jig, fins-down
- Thermally conductive adhesive is dispensed down middle of lid
- Structural epoxy is dispensed at edges of lid
- Lens assembly is pressed into housing
- Epoxy is allowed to cure

➢ ***Module Assembly Finishing:***

- Terminal wires have connectors attached and seals slid on

- Remove protective silicone plugs from connector cages
- Dispense conductive epoxy into cages
- Slide wire ends into cages; epoxy is allowed to cure
- Wires are clipped to module back
- Protective tabs are folded/bonded over wires where they exit lens, compressing seals
- Module is flash tested
- Module is packed for shipping

5.2.3 Variations on the Receivers-to-Lens Assembly Step

- **Flexible adhesives:** With the 410 stainless CTE only 1.2 ppm/K higher than the lens glass, flexible adhesives can accommodate the +/- ~70 um shift between glass and steel, and the ~15 um receiver-to-glass shift for a low-stress module. Assembly is at room temperature.
- **Sealing lens-tile subassembly:** To ensure that epoxy only goes where it is wanted, it may be useful to seal the lens tile subassembly before placing the lid. If the high voltage prepreg is on the receiver backs instead of the inside of the lid (which would allow electrically conductive thermal interface epoxy to be used), then a bead of sealant between the backs of the receiver substrates would be needed. Also, if a low-viscosity epoxy is used to bond the lid to the lens tile, then the sides of the lens tile might need to be sealed first to prevent epoxy from flowing into the areas between the thermal contacts.
- **Dispense onto Receiver:** Depending on the dispense equipment, it may be preferable to dispense the structural adhesives onto the individual receivers rather than onto the lens tile. This it is probably simpler in low-volume production because a lens tile is much longer and might need custom dispense equipment.
- **Placing Receivers in High Volume**: In high-volume assembly cost can be further reduced with modified placement equipment that picks a receiver once and uses the placement positioning to the move the receiver to multiple stations. Unlike normal chip placement, the part being placed is held (via normal a vacuum grip) from the bottom rather than from the top. First the receiver is pressed against a stamp that transfers the optical silicone adhesive to the cells. The receiver is then pressed against a stamp that transfers electrical

epoxy dots for the previous receiver's tab to connect to, and finally against a stamp that transfers the pattern of structural epoxy dots. The receiver is then placed **up** against the lens, where the tack of the structural epoxy will hold it securely. The lens tile can then be stepped over one receiver-length (to minimized placement head movement for high speed).

- At each stage a camera checks the precise alignment of the receiver to the stamp or to the lens's cones, and actuators shift the receiver a fraction of a millimeter this way or that to ensure precision alignment. Standard pick-and-place and dispense equipment already uses similar camera-based final position checking, so this rearrangement to optimize for a single type of assembly would not involve new processes. 200μ accuracy is sufficient for everything except the alignment of the cells to the cones (where tighter alignment reduces the excess cell size needed).

Since these variations are minor, no work needs to be done on them now, and these variations are documented here as food for thought. It is likely that the simplest to implement process will be used initially and the most efficient process will be used in high volume.

5.2.4 Module Assembly Costs

While receiver micro-assembly cost is already accounted for, the cost of placing the receivers, plus placing the lens in the lid and placing the module terminal wires, need to be covered.

- ***Receiver Placement:***

The NREL Bottom-up Costing has a category for the pick-and-place and bonding of receivers onto a back-plate, which is comparable to TLC's placement of receivers onto the lens tile but must *also* include interconnecting the receivers since that could not be done beforehand. The category also includes assembling the back plate and lens plate to the module frame, which are comparable to placing the lens in the lid in TLC (and must include placing the module terminal wires as well). The NREL costing backup slide shows the non-material costs for this category as 4% of 11.2 ¢/W, or 0.45¢/W. With 5 mm x 5 mm cells in a 30% efficient 1000X module and thus 7.5W per receiver, that is 3.4¢ per receiver.

That cost is significantly pessimistic for TLC in comparable volume because while the TLC receivers are ~10x larger in area and require three dispenses instead of one, they are thin and flat-

topped for easy handling, they are placed on a 30x smaller 'backplate' and so require less placement head movement, the receivers do not require wires inserted into cages to interconnect them, and the large parts in the final steps are 40x smaller per Watt (or 10x smaller per receiver).

As a cross-check on the larger size, silicon wafer handling is automated and hence CapEx-heavy, and it only accounts for ~3% of total silicon PV production CapEx (Powel, et al., 2015). So, the handling cost should be *at most* ~3% of the ~$0.40/W production cost, which is 1.2¢/W or 4¢ wafer. A TLC receiver is 8x smaller and much less fragile than a wafer, so even with the added steps 3.4¢ per TLC receiver looks highly pessimistic at high volume.

While this 3.4¢ per TLC receiver cost is pessimistic at 100 MW or more, it is a referenceable cost and is reasonable for the introductory 10 MW volume, and 3.4¢ per 33-Watt receiver is only ~0.1 ¢/W in any case so it is not worth tightening the bounds on this cost.

- ***Assembly Materials:***

In addition to the 0.034 ¢/W of thermal epoxy for the receiver-backs-to-lid attachment, material costs in the assembly are electrically conductive epoxy for receiver-to-receiver connections (0.004 ¢/W), and structural epoxy for receiver-to-lens attachment (0.006 ¢/W) and to the flanges of the lid housing (0.017 ¢/W).

- ***Total Assembly Costs:***

The costs thus total ~0.16 ¢/W. It will take several runs with increasing automation to get the parts handling right, so hitting these costs at a 2nd 10 MW, will require multiple few-thousand-module runs, but even if production is rushed, they are probably correct for 100 MW and should be pessimistic by 1 GW. These costs will be revisited after a first prototype is completed.

5.3 Reliability

In addition to higher efficiency, TLC's higher value per Watt is also based on high durability. The value of durability can be seen in both the "Value of Modules" chart, where the value curves for more-durable modules are much higher than for low-durability modules, and in the "Pathway toward an LCOE of 3¢/kWh" chart, where it is a close second to module efficiency as the biggest driver of lower cost.

It is thus important to both design for high durability and to be able to predict the durability. The key steps to ensuring low degradation and long product life before the first prototype is built are:

- Understanding standard PV failure modes and their applicability to TLC
- Understanding any HCPV-specific failure modes and their applicability to TLC
- Understanding any TLC-specific failure modes
- Designing to eliminate or dramatically reduce applicable failure modes where practicable
- Assessing the risk and effect of any remaining failure modes

5.3.1 Temperature Cycling

The reliability of tandem cells under temperature cycling and UV radiation can be appreciated by considering the harsh environment that they experience in space. Nasa's rule of thumb on temperature cycling in space is -120ºC to + 120ºC, and space panels with tandem cells are often tested to even more-extreme temperatures (Fatemi, et al., 2013). A detailed reliability analysis of CPV cells also predicts a 70-year lifetime under harsher conditions than a TLC module (Gonzalez, et al., 2009). The only significant temperature effect is thus thermo-mechanical stress, and the cells are lower CTE than the copper and are thus under compressive stress after cooling from soldering temperatures; cells are very strong under compressive stress, so the temperature cycling seen in TLC is much gentler on the cells than that seen in space.

The temperature cycling puts alternating tensile and compressive stress on the solder, creating fatigue that could eventually propagate cracks to solder-joint failure. However, in accelerated aging testing of SAC solder joints, solder between low-CTE resistors bigger than cells, survived ~3000 rapid thermal shock cycles between -40ºC and 125ºC (Nurmi, 2005). This is even clearer in CREE's solder-joint testing for LEDs on AlN (CREE, 2017):

Cree solder-joint thermal shock test, -40C to 125C @ 1C/sec, LEDs on AlN with lead-free solder

Cell length	Cell Width	Diagonal	Cycles to 1st fail
2.45	2.45	3.46	2600
3.45	3.45	4.88	1800
5	5	7.07	1200
9	7	11.40	800
9	9	12.73	400

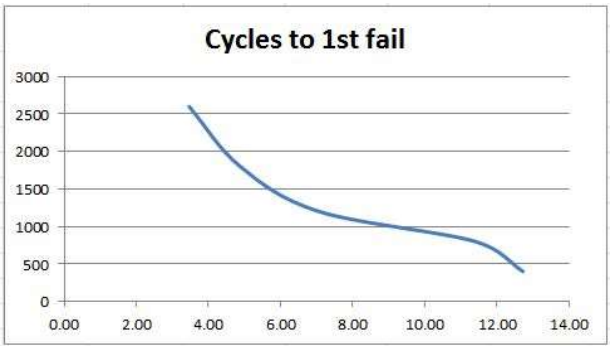

Figure 5.1: *Effect of Die Size on Solder Joint Reliability.*

The speed at which solder joints fail depends strongly on cell size, and small cells are almost immune. Not only are TLC cells on copper are gentler on solder than LEDs on AlN, but TLC cells have diagonals 1/3 the size of the smallest LED solder-joint tested on that graph so they should survive for many thousands of such shocking cycles. Since 200 milder cycles (from -40ºC to 90ºC) is sufficient for ten years in flat panels (Kurtz, et al., 2013), survivability for many thousands of shock cycles should ensure no solder-joint fatigue problems even in a 50-year product life.

5.3.2 UV

Space cells also experience harsh UV light, with only a thin coating to protect them. For CPV cells on earth, the ozone layer filters out all of the UVC and filters out ~80% of the UVB (UVA is useful in contributing to the first-junction photocurrent). The mirror backing also filters out over 90% of the UVB. Each 3 mm of low-iron glass filters from ~50% to ~100%, depending on the glass, and even at 50% a ~25 mm total pass through low-iron glass (mirror glass twice plus lens glass) reduces the UVB 250X.

To a first order the UVB is reduced by four orders of magnitude relative to AM0 in space, so even at 1500X concentration a TLC cell sees less UVB than a space cell sees without concentration. It isn't quite that simple, since the gentler UVB near 315 nm is filtered less, but even at 315 nm the ozone layer reduces the UVB ~5X, the silver reduces it ~20X, and the glass reduces it >10X, so even the gentlest of the UVB is reduced ~1000X and is thus only slightly stronger at 1500X geometric concentration when the sun is straight overhead; by AM~1.2 even this gentlest UVB is

weaker. And at 310 nm and below, the UVB reaching the cell is reduced relative to AM0 UVB even when the sun is straight overhead.

So even at 1500x concentration the cells see a milder UV environment than cells in space do. Relative to other CPV, other CPV usually only has ~3 mm of low-iron glass, and the instrument-grade glass secondary optical elements absorb little UV, so even at 315 nm the extra glass and the silver reduce the UVB by ~200X compared to Fresnel-lens box CPV, and the ratio gets even larger at harsher wavelengths. Even at 1500X, a cell in TLC thus receives less than 1% of the UVB that a cell in traditional Fresnel/Box CPV receives, so UV is should not be an issue in TLC.

5.3.3 Degradation

The expectation of low degradation is based on CPV cells being far less sensitive than silicon to light-induced degradation, potential-induced degradation, and even thermal degradation; for example, the Boeing Spectrolab sheet shows <0.5% degradation after high-temperature soaks equivalent to 25 years of normal use (Spectrolab, 2011), or <0.02%/year, while even the best silicon panels degrade at 10 times that rate (Jordan, et al., 2016) and cheap silicon panels degrade at ~25 times that rate.

5.3.4 Comparison to Silicon PV

Since silicon PV reliability has received the most study, the risks and impacts are best understood there so a comparison to silicon PV forms the bulk of the reliability assessment; this is based on an International Energy Agency study (Kontges, et al., 2014), supplemented by an NREL study (Kurtz, et al., 2013). The HCPV sections and especially the **TLC section are smaller but more specific.**

TLC was found to be less susceptible to the main causes of silicon module failure than the silicon modules are (largely due to the compact size of the sealed module in TLC).

Table A.1 in Appendix A.1 details the assessment.

5.4 Summary of Costs and Value

The detailed analytical models calculate costs versus production volume for a TLC design, given pricing information for components and processes, and calculations of the efficiency if built as designed. The example below shows this for the introductory volume of a second 10 MW.

5.4.1 Costs in ¢/W

***Table 5.1**: Table and Chart of Cost Estimates for TLC.*

Volume	Cell	Rating DNI	Efficiency	Module Optical Length	Module physical Length	Watts per module	Captial to Scale $M
10 MW	Azur 3C44	1000	**39.49%**	1698	1721	**1061**	0.59
Item	%	Cents/Wa	Includes				
Trough Mirror	24.1%	5.18	COTS RP3	1	1= INCLUDI	0=No Miror	
Lens tile	4.2%	0.91	Glass	Roll Formin	Coating	Cutting	Silicone
Cell	44.3%	9.51	Azur 3C44	Dicing adde	Testing adder		
Rcvr Assembly	5.9%	1.27	Coupling	SAC Solder	Wire Bond	Placement	Diode
Substrate	5.2%	1.12	Base Copper	Top Copper	Subst Proce	Bonding	Electroless
Connectors	1.4%	0.31	Terminal Wire	Terminals	Cages		
All-aluminum heatsink lid	11.6%	2.50	Extruded+welded				
Mechanical	1.4%	0.31	Wire-shield tal	Brace	Vertical Support		
Module Assembl	1.7%	0.37	VT 4a2 prepre	Thermal epo	Elec. Cond.	Epoxy Lens:	Epoxy lens
Tracker	0.0%	0.00	Trackers Feina	0	1= INCLUDI	0=No tracker	
Total	**100.0%**	21.48	10 MW				

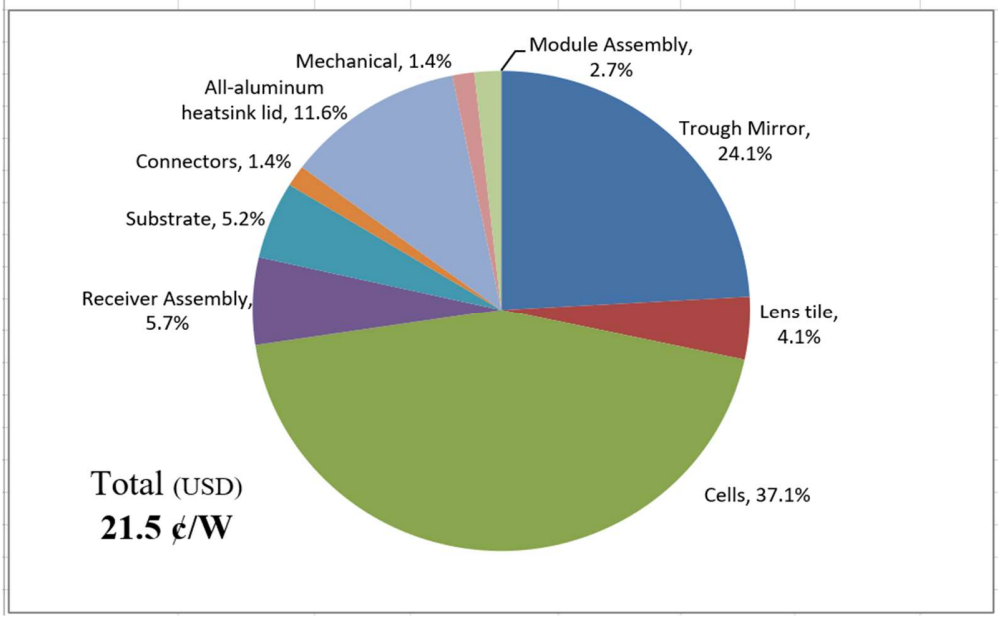

Costs so far have come down with time, with early estimates (cells, lens, heat spreader, manifold, placement, wire-bonding, mirrors, diodes, solder, dicing) dropping an average of >25% as quotes were received, individual steps were costed, and processes and materials were refined.

- ***Detailed Costing:***

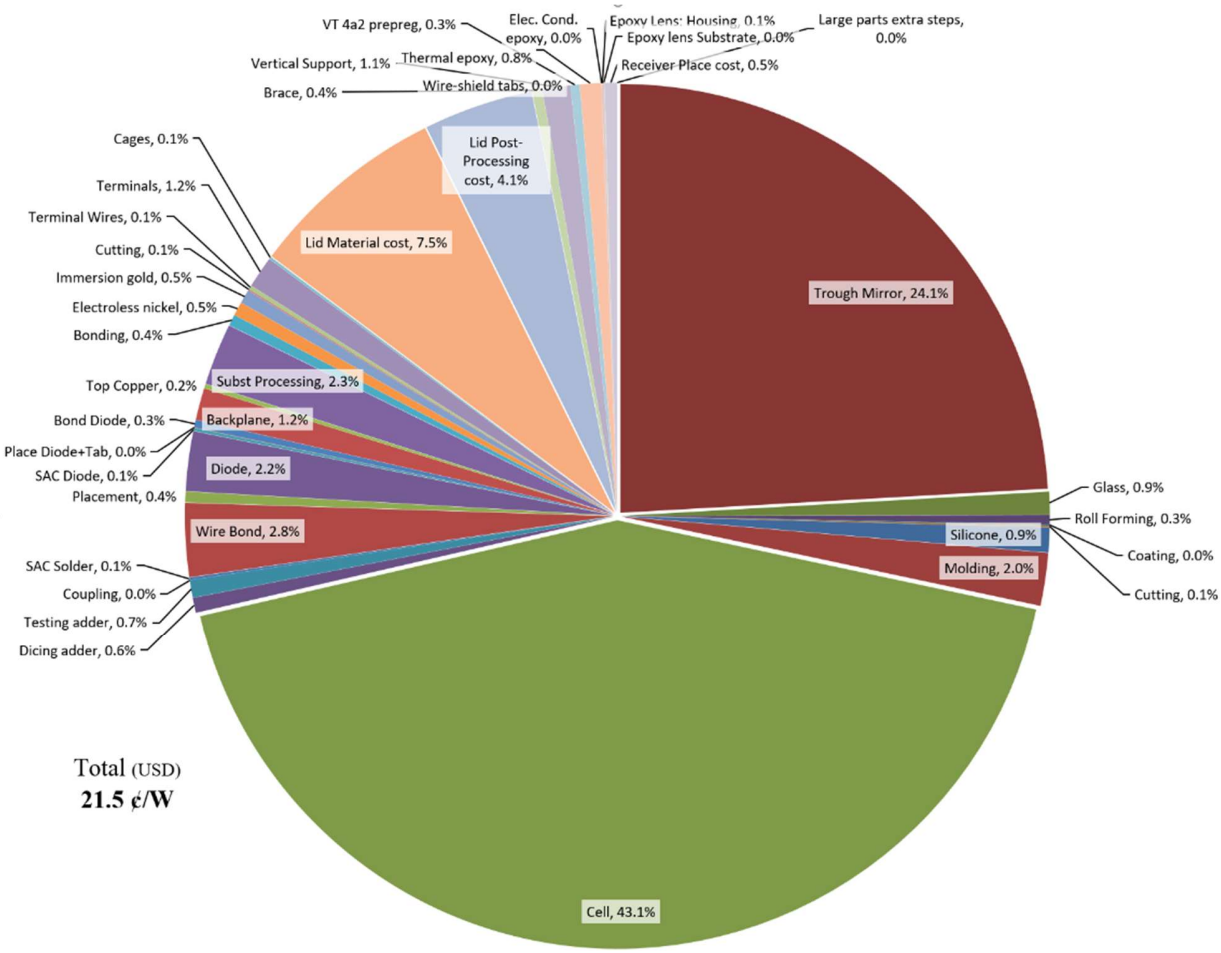

Figure 5.2*: Detailed Costing Pie Chart.*

A detailed costing shows where to focus efforts on improving costs. ~70% of the cost is from quotes from suppliers (all for quantities of 10 MW or less for initial costs); this rises to ~85% if the tracker is included! These are typically costs from the first quote that I got from the first supplier who responded, so a good purchasing agent with a purchasing budget should be able to improve these significantly through competitive bidding.

Another ~8% of the costs are from cost studies (such as the NREL CPV costing) or articles on costs from major companies in their fields (e.g., Kulicke and Soffa). Another ~12% of the cost comes from material costs, typically for small volume, found on various supplier web sites, but where quotes have not been obtained. A good purchasing agent should be able to beat these costs.

The final ~10% of the costs are deduced from dissecting costs of related but not identical materials or processes. These are mostly tiny costs; but the heat-sink post-processing, the substrate processing, and the silicone molding are significant, so proving out these processes are high on the prototyping priority list. More details can be found in the accompanying spreadsheet for the analytical model. The low materials cost should keep total cost low (Norman, et al. PVSC, 2019).

5.4.2 Cost Reductions

- ***Sources of Cost Reductions:***

- Multiple Suppliers:

Since the focus has been on technology and introductory cost, the price listed for most components is simply the first supplier who responded, or, for minor items, the prices found in a quick web search. With the bulk of costs now from off-the-shelf items, contacting more suppliers for each component has more potential to reduce costs than further design enhancements.

- Cost Changes with Production Volume and over Time:

Only a few of the suppliers contacted have been asked for volume pricing, and all of the web-search prices have been for <10 MW volume, so the only volume discount included is cells. Future cost reduction analyses from DOE/NREL and Fraunhofer are then used wherever applicable.

When costs are from cost studies on existing technologies, the opposite is common; these are typically for higher volume than needed for 10 MW of TLC, so reverse scaling is used to increase the cost for small volume. When no supplier or DOE/ Fraunhofer forecasts are available, the analytical model uses general scaling laws:

Automotive Scaling Law: A custom version of something that the industry already does in high volume typically has a 5% reduction in cost for every doubling of volume. This is an adjustable parameter on the Financial tab. Setting this to 1 removes all cost reductions from scaling, unless confirmed by quotes from suppliers or from DOE / Fraunhofer cost studies.

Swanson's Law: Dick Swanson of SunPower observed that silicon photovoltaics undergoes a 20% cost reduction per doubling of total volume produced (as the industry as a whole develops lower-cost techniques). This will also be an adjustable parameter on the Financial tab. Setting this to 1 will remove all cost reductions from industry learning scaling, unless confirmed by quotes from suppliers or from DOE / Fraunhofer cost studies. Setting it to a non-zero value applies that scaling once any component's volume exceeds the industries current production of similar components. (This feature is not yet implemented).

- ***Cost Reductions by Component:***

Component	Cost Reductions Included
Trackers	• 10 MW: Supplier Quote • 10 MW to 1 GW: Non-learning manufacturing scaling rule • 1 GW: Derived from DOE 2020 Heliostat cost target • >1 GW Non-learning scaling
Mirrors	• 10 MW: Supplier Quote • 10 MW to 1 GW: Non-learning manufacturing scaling rule • 1 GW: Derived from DOE Heliostat 2020 cost target • >1 GW: Non-learning manufacturing scaling rule
Lens Glass	Currently based on quote from Aohong for 10 MW of better ultra-clear low-iron glass sheets. Since not a specialty item, and current low-iron glass manufacturing is equivalent to terawatts of TLC per year, the only cost reductions would be sales negotiation. • No cost reductions are assumed at any volume.
Lens Roll-Forming	Derived from difference between patterned and un-patterned glass flat-panel module covers; scaled by the lens thickness (heating cost). Manufacturers texture in-line at high volume (eliminating the extra heating), so the thickness scaling is removed at 1 GW.

	• Thickness penalty removed at 1 GW for in-line rolling.
Lens AR Coating	Derived from difference between coated and un-coated flat-panel module covers. Since the volume would be much lower, this difference is multiplied by 10. This sanity-checks with the NREL CPV costing – 1/6 the glass plus the coating plus the cones at 100 MW is 0.51¢/W and 48 times less area is needed so this suggests 25¢/W, which is 50% higher than the 16.7¢/W shown by NREL (the difference is attributable to higher volume in Fresnel CPV). • No cost reductions are assumed at any volume.
Lens Tile Cutting	Derived from silicon wafering cost, not adjusted for much-easier-to-cut glass. Cost increased 10-fold at 10 MW volume to account for a carbide saw rather than a wire saw. • Cost reduction to wire sawing (pessimistically based on - harder and more-fragile silicon) assumed at 1 GW.
Silicone for Cones	Derived from the NREL Fresnel CPV study. Cost is increased due to low volume (reverse scaling). • Cost increase based on no-learning scaling at < 5 GW • Cost decrease based on no-learning scaling above 5 GW
Cone Molding	Derived from NREL Fresnel CPV study at volume equivalent to 5 GW. • Cost increase based on no-learning scaling at < 5 GW • Cost decrease based on no-learning scaling above 5 GW
Substrate Materials	• 10 MW cost based on supplier websites and quotes • > 1 GW Cost reduction based on non-learning scaling
Other Substrate cost	• Punched or sawn process assumed for > 10 MW

Cells	For 100 GW, the NREL cell cost estimate, including **only technologies already demonstrated in the lab**, is used. • From supplier quote for ~3 MW to 250 MW. • Non-learning scaling used for 1 GW to 10 GW • NREL study used for 100 GW (with or w/o reuse)
Cell Testing	• 100 MW Cost derived from NREL CPV cell cost study • Non-learning scaling used for other volumes
Cell Attach	• From cost studies for LEDs and chip packages. • No cost reductions since already high-volume processes
Diode	• 10 MW from supplier quote for ~10 MW. • No scaling except for volume purchase
Module connectors	• 10 kW cost from supplier catalog used at 10 MW • No-learning scaling used beyond 10 MW
Other	• No cost reduction assumed

Table 5.2: *Table of source and cost reductions used in cost model.*

- ***Effects of Scaling Production:***

The approximate effects of scaling production can be seen in the charts below. Please note that these are pessimistic, and that cost reductions will almost certainly be faster (Swanson's law for silicon has been over 4x faster!). Furthermore, these costs are based on today's cells, and assume **no cell progress** other than cost reductions from production volume and from NREL studies.

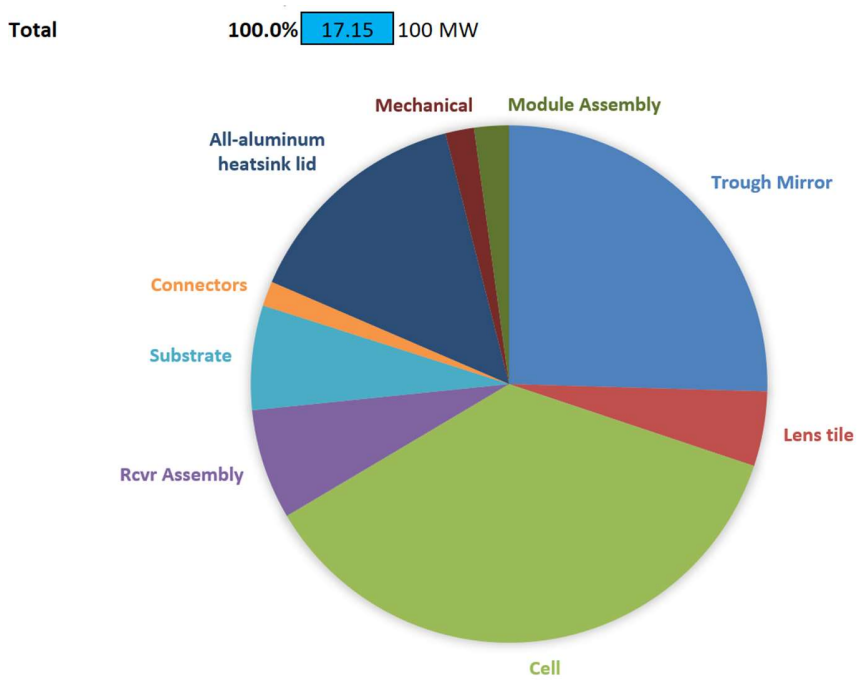

Figure 5.3: *17.1 ¢/W Cost Pie Chart at 100 MW.*

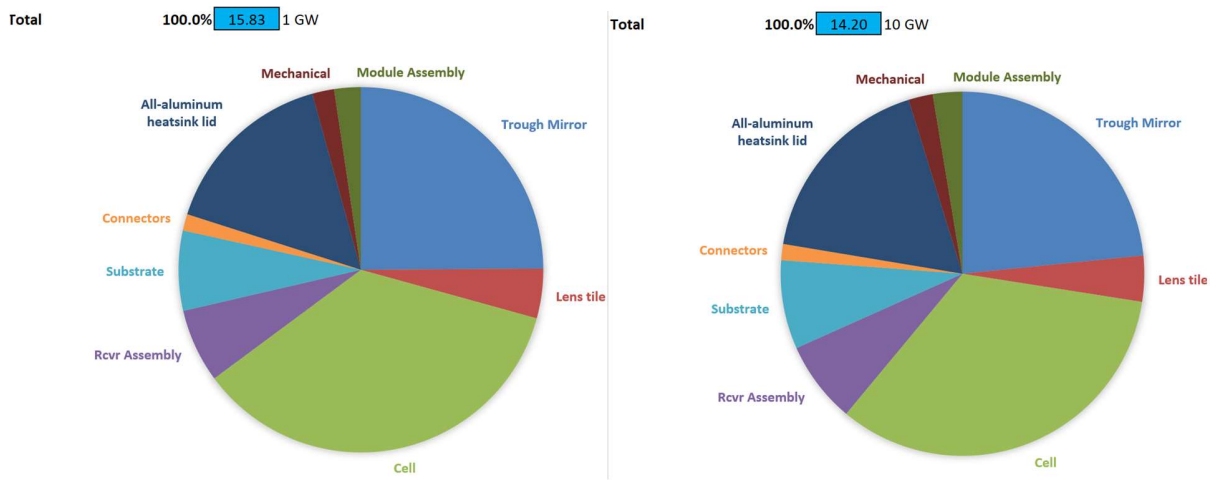

Figure 5.4: *Cost Charts 15.8 ¢/W at 1 GW (left) and 14.2 ¢/W at 10GW (right).*

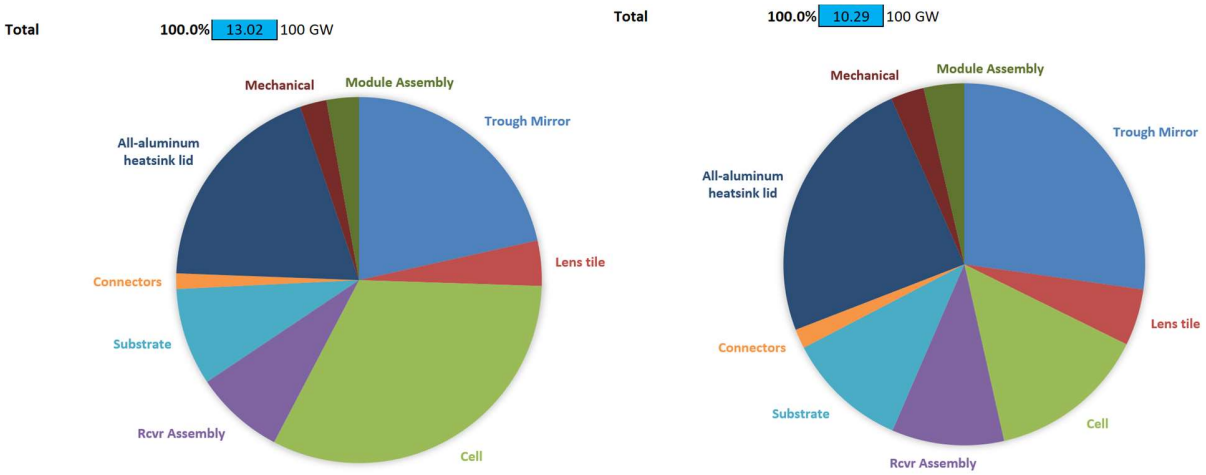

Figure 5.5: *Costs at 100 GW (Left: without substrate reuse, right: with reuse).*

The 100 GW cost is key because silicon panels are now made at roughly this volume. This shows that TLC can not only meet silicon on costs when made at 10 MW (10000 times lower volume) but can beat silicon dramatically (~60% lower) on cost per Watt when made at the same volume.

5.4.3 Value

- ***Capacity Factor Adjustment:***

At utility scale ¢/kWh is more important than ¢/W, and cost per kWh depends on capacity factor.

For comparing DAT TLC to other HCPV no adjustment is needed, but for a fair comparison to silicon on single axis trackers the different capacity factor needs to be accounted for.

Picking a typical Tier-1 module (Trina mono Si) on a flat single-axis tracker in SAM (a dominant case at utility scale) and adjusting the reference DNI for TLC until the capacity factors are equal gives a de-rating for TLC of 0.766 in the DOE's reference case of Kansas City.

This means that to produce the same annual kWh in KC, TLC needs ~1.3 times the nameplate watts installed as silicon PV on single-axis trackers would. This shows up both in equivalent module cost, which must be increased 1.3 times to cover the modules themselves, and the equivalent efficiency, which is lowered to 0.766 times the actual DNI efficiency, or to ~30%.

(While of less interest at utility scale, fixed-tilt is still used for smaller flat-panel installations. The capacity factor derating of TLC relative to silicon at a fixed 20º slant is 0.91, so equivalent costs are increased by only 10% and the efficiency decreased by only 9% to bring TLC's equivalent efficiency to ~35%).

- ***Efficiency and Lifetime Adjustments:***

Two other key value adjustments are from the much higher efficiency and durability of TLC relative to flat panels. A DOE chart showed that even in low volume TLC could easily meet the DOE's SunShot 2030 goal of a 3¢/kWh LCOE in Kansas City (KC).

As can be seen from the overlaid DOE chart, for a 3¢/kWh LCOE in KC, a 50-year TLC module could be sold at 22 ¢/W, and cover the 2-axis tracker cost while still meeting the DOE 2030 goal.

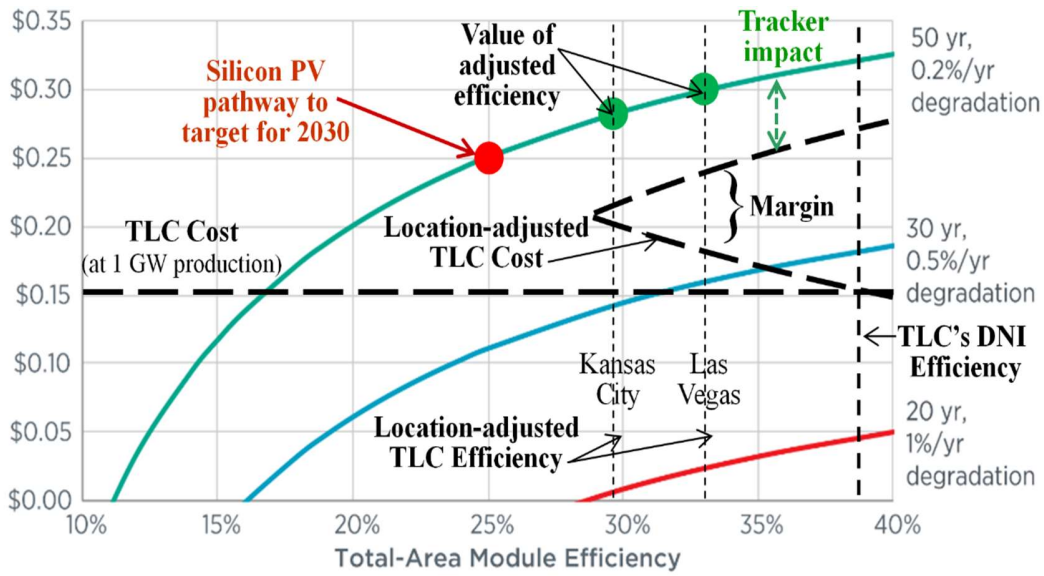

Figure 5.6: *Value of Modules* - ***Maximum selling price that would meet the DOE 2030 goal***.

In contrast, today's 20%-efficient, 30-year mid-cost silicon on single-axis trackers would have to sell for ~6 ¢/W to meet the same LCOE.

As this repeated chart from the beginning of this document shows, TLC can not only beat the DOE's SunShot 2030 goal of 3¢/kWh unsubsidized in Kansas City **(KC),** but can beat the cheapest silicon flat panels on cost while providing higher value due to higher efficiency.

This is a bit simplistic when comparing TLC to flat panels because it doesn't include the capacity factor effect on inverter loading ratio or separate out various maintenance costs.

Value Item	*SIPS 16% / 1 GW*	*¢/W 10 MW*	*¢/W 1 GW*	*~100 GW*
Higher inverter loading	-2	-1	-0.6	-0.4
Higher Maintenance	Just capital	-1.6	-1	-0.6
Total Value Adjustment	**-2**	**-2.6**	**-1.6**	**-1**

Table 5.3: *Table of Value Adjustments for TLC.*

- ***Details of Value Adjustments:***

  - Inverter loading Ratio:

By adjusting the TLC efficiency from down to 38% to account for not utilizing diffuse light in Kansas City, most of the capacity factor difference is accounted for. However, one significant factor remains: a higher capacity factor also allows a lower inverter loading ratio and thus fewer inverters for the same annual energy production. For the same annual kWh output, TLC thus takes either more inverters to capture its higher peak output or a higher inverter loading ratio (ILR) to provide more consistent output comparable to tracked silicon.

This impact of this is significant; at the predicted 5.85 ¢/W for utility-scale inverters in 2020 (Schneider, 2016), the extra inverter capacity to capture enough extra output on the brightest days is 1¢/W in a typical U.S. location! Thus a 1¢/Watt value penalty is given to TLC at introduction; continuing the rate of decrease, this falls to 0.6 ¢/W in 2025 and 0.4¢/W in 2030.

Brute-force adding of inverters to capture the higher peak output of TLC is an easy to calculate way to equalize the LCOE, and thus provides an upper bound on the cost: some blend of slightly more inverters and a slightly larger array would probably cost less, but it would require a lot of site-specific work to uncover a fraction of a cent per Watt of savings.

  - Maintenance:

Inverter replacement is also the biggest maintenance item so more inverters mean higher maintenance. Other maintenance is also affected by the module type: most per-area costs are

reduced proportionately by the higher efficiency, including tracker maintenance and vegetation management. Cleaning is more complex to compare – while TLC has only 2/3 the area to clean it is roughly 2.5 times as sensitive to dust so flat panels have the edge on cleaning cost, but in most dusty areas TLC gains from the higher direct-to-diffuse light, offsetting the cleaning penalty.

In the reference case of Kansas City as a 'typical' U.S. location, maintenance costs per kWh are expected to be ~18% higher, which at the rate expected for 2021 launch is an extra 0.136 ¢/W per year, which at a 5% discount rate has a value of about 1.6 ¢/W. Thus, a reduction of value of 1.6 ¢/W is assigned to TLC for product introduction.

Maintenance costs continue to fall, so the penalty is reduced to 1¢/W for 2025 and 0.6¢/W in 2030. This is highly pessimistic because inverter maintenance is falling rapidly, and trough cleaning is becoming automated, so areas where TLC cost more are shrinking faster than average.

CSP Maintenance: Some reviewers have dwelt on trough/tracker maintenance and cleaning, but these costs can be cross-checked against parabolic troughs for CSP. NREL's **Line-Focus Solar Power Plant Cost Reduction Plan** (Kutscher, et al., 2010) sets an upper bound. It shows a total O&M cost of 1.5¢/kWh, of which 19%, or 0.285 ¢/kWh is the solar field maintenance cost. But TLC generates ~2.5 times the kWh per area, reducing the upper bound at ~0.115 ¢/kWh. Even that is grossly pessimistic because it includes the fragile CSP receiver-tubes and the high-temperature fluid piping with ball joints, which TLC does not have, and costs have also fallen over the years. But regardless of being pessimistic it shows that the total trough/tracker O&M is at most a value cost of roughly 1¢/W, since adding the second tacking axis should add less maintenance than adding a first tracking axis to a 50% larger area of silicon panels.

CSP Cleaning: Cleaning is much like CSP troughs, regardless of the tracking style, but reduced proportionately to the greater efficiency. The NREL line-focus study shows cleaning as ~3.5% of the total O&M, or ~0.024 ¢/kWh, so the efficiency reduces this to ~0.01 ¢/kWh, which is insignificant. Even in an area ten times dustier, 0.1¢/kWh would be just an inconvenience and not a show-stopper.

5.4.4 SunShot 2030 – Pathways to 3¢/kWh

SunShot's 2030 target is not impossible for silicon to meet this, but silicon will have to achieve many massive improvements beyond just cost reduction to be able to achieve this. The "Cost Reductions Beyond the SunShot 2020 Targets" summary (Woodhouse, et al., 2016) outlines a path to an LCOE of 3¢/kWh in Kansas City, as shown in **Figure 5.7.**

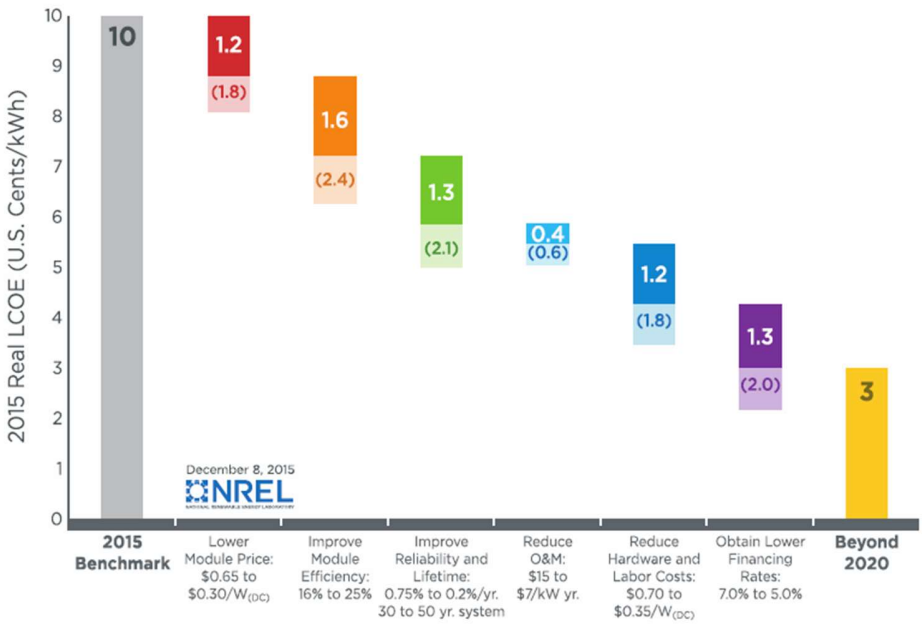

Figure 5.7: *Path Toward 3¢/kWh* (Woodhouse, et al., 2016).

While Silicon PV will have a major challenge, TLC is well positioned. Even with current cells and made at just 1 GW, TLC will meet or beat all three of the top technical targets: TLC will supply **107%** of the module price reduction (at NREL's 15% minimum sustainable margin), **100%** of the reliability / lifetime improvement, and **213%** of the efficiency increase sought. And TLC combines with other solutions: TLC will benefit from tracker, inverter, glass, automation, and installation cost reductions as much as silicon PV does, which should bring LCOE well below 3¢/kWh.

TLC could reach that goal by 2025, five years ahead of the SunShot goal, with a high-quality first world-dominated supply chain. With ample margin to fund growth and low capital to scale, TLC could reach 100 GW by 2030 (with no show-stoppers to even higher volumes).

At 100 GW in 2030 (not even matching the volume of silicon panel production today), even without any cell efficiency improvements for TLC, and even allocating an adjusted 14.3 ¢/W for the tracker (adjusted up 29%), TLC should be able to sell for 22.7 ¢/W, and even after adjustment would only cost 15.7¢/W, a ~35% margin. Include a 5% cell efficiency improvement (half the historic rate) and substrate reuse (already demonstrated in the lab), and a reasonable estimate for 100 GW in 2030 is a module selling for 23.5 ¢/W and costing only 9¢/W to make, for a massive ~62% margin.

- ***NREL System Advisor Model Cross-Check:***

A free download of NREL's solar advisor module (SAM) has been used as well as possible to find the allowable cost to meet the 3¢/kWh goal in Kansas City Missouri by 2030. TLC specific values were used when applicable, values from the NREL study on reaching 3¢/kWh when those could be entered and using the SAM defaults for the remaining fields. SAM indicated that to meet this target in Kansas City Mo., SAT TLC could sell for 48¢/W. This would be a lucrative 37% margin even at the introductory cost of ~30¢/W, and a 43% margin if made at single gigawatt quantity. DAT TLC does even better on costs and margins. This supports that 3¢/kWh by 2025 is realistic for TLC when made at a gigawatt scale.

5.4.5 High Efficiency

Much of the high value comes from high efficiency. TLC is optimized for low cost, yet it claims near-record efficiency as well – how realistic is this given HCPV's traditional losses?

A classic presentation from NREL (Kurtz, 2009) is still a useful reference for this. It shows where the efficiency losses have typically been in CPV, and thus provides a checklist for ensuring that major sources of losses are not overlooked.

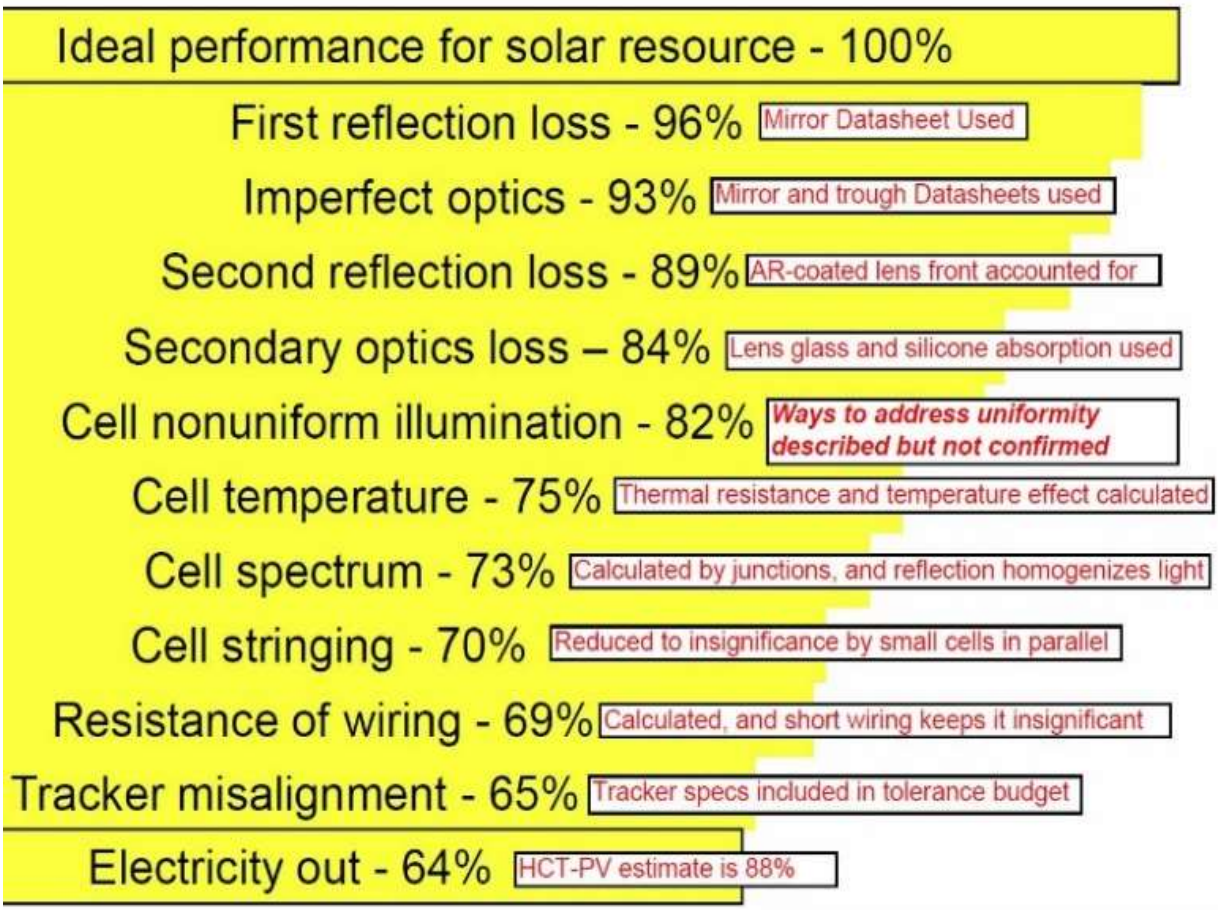

Figure 5.8: *CPV Losses (from S. Kurtz), annotated for TLC.*

Mirrors have lower losses than Fresnel lenses, and TLC lenses have less-sharp illumination angles than typical Fresnel secondary optical elements, the optics are very transparent at the key wavelengths, and with most of the concentration being reflective, the spectrum should be very homogeneous (and small cells are not sensitive (Hayashi, et al., 2015)), so TLC optics should have near-record efficiency. The best commercial cells are used, these cells have low temperature sensitivity and TLC has good cooling, the small cells greatly reduce mismatch losses in cell stringing, and the short electrical path through low-resistance materials ensures that electrical losses are tiny. Trackers are now very accurate, and TLC is not even as demanding of accuracy as heliostats for CSP are.

- ***Possible Extra Losses:***

The major known efficiency uncertainties both involve the cell's conversion efficiency. First, the TLC cell's efficiency at 1500X is based on adjusting the Azur datasheet's efficiency for the concentration's effect on five types of losses, providing a good estimate but which should be validated. Compounding that is the non-uniform illumination highlighted above, not so much at the cell level, but along the row of cells. The initial optical unevenness is comparable in TLC to most Fresnel/Box designs (not as good as an FK secondary, comparable to Silo and far better than no secondary). Small cells are less affected by unevenness than large cells, and TLC's center boost matches its central bus bar, but TLC is already running the cells on the high side of their efficient range concentration. Several ways to ameliorate this, or even to benefit from this, are discussed under the optics section; these could even result in a net gain of a few percent. However, if left unaddressed the Azur cell sheet suggests that this would cost ~~5% for the most impacted cells, and the average would be hit by ~1/3 of this, or a ~~2% power decrease.

5.4.6 Target Market

The initial target market is utility-scale solar on dual-axis trackers in regions that have at least moderate sun quality (DNI higher than GHI). Every major solar market except Germany meets the sun quality, utility-scale solar is the largest and fastest growing market segment (it grew to ~72% of all U.S. installations last year). Single-axis tracking is the largest and most rapidly growing section of utility-scale solar (and so is maintained as an option), but dual axis tracking is starting to become established.

Furthermore, it is not that TLC can't be applied to other markets. It is just that having with the largest and best-financed market already focused on LCOE (value), provides the luxury of focusing on a single market initially. However, utility-scale is a conservative market, so TLC will have to undergo extensive testing, which may be done in smaller, more nimble markets.

In addition to its major cost advantage, there are cases where efficiency will prove key. Consider a developer who has a PPA signed for a 100 MW site, with the interconnection already under construction, and then finds desert tortoises or burrowing owls and has to set aside half the land. The developer can either renegotiate the PPA and risk a lower price or the deal falling through, or try that new-fangled low-cost, high-efficiency technology and meet the already-signed PPA.

Once TLC is proven at a reference site, TLC can undercut silicon on LCOE by whatever it takes to keep TLC's pipeline full, with ample profit to fund rapid growth. At 1500X the world already has a GW/year of cell capacity, and at 40% it already has several GW/yr of mirror capacity and lens over-molding capacity and several tens of GW/year of tracker capacity and lens roll-forming capacity, so there are no technical or financial barriers to extremely rapid growth.

5.4.7 Other Markets

Tracked utility-scale-PV is currently by far the biggest solar market (in the U.S. it is larger than all other markets put together), and it is the fastest growing major market, and it is largely in high-DNI regions where TLC has even larger advantages. TLC will not need to address any other markets for some time. The only related issue is currently single- or dual-axis tracking.

But the world changes and distributed solar or even roof-top solar might become dominant again. And even if utility-scale remains dominant and TLC is successful there, TLC could *defensively* prevent possible competitors from using other major markets as beachheads to gain scale and bring costs down to where they can challenge TLC.

The section on design scaling down shows that TLC can scale down to a few kilowatts at only a modest cost increase, so the limit is really how small one can make a tracker. TLC can thus scale down to community solar and to commercial rooftops, which should confine flat panels to the residential market (and even there, TLC can be used where heat byproduct is useful) and to regions with light quality too poor (too low DNI relative to diffuse light) for CPV. Once TLC becomes dominant, any emerging-technology challengers to it would also have to go through these markets to gain scale, so if TLC meets its performance, durability and cost targets it should dominate solar for a generation.

5.4.8 Production Scaling

- ***Tooling Costs:***

While silicon flat-panel PV's capital cost per watt per year has now come down below $1/W/year, TLC has far lower scaling costs, both in direct scaling costs for TLC production and indirect scaling costs borne by suppliers as a part of their normal costing.

The balance between direct and indirect costs depends largely on the amount of outsourcing, which will increase as production scales. The custom equipment will be molds for silicone over-molding, rollers for roll-forming lenses, and jigs for holding parts during assembly. All tooling is evaluated based on 2500 hours of use per year – comfortable with two shifts four days per week, leaving time for maintenance and avoiding 3rd-shift error rates.

The only expensive initial parts-tooling is for the optics. Estimating from the surface area needing precision machining, mating lens rollers should cost in the neighborhood of $300K each. While $40K has been mooted for a 3-foot-wide roller with a much more complex pattern (Becky Laland, 2009), with another $20K of fittings to a rolling machine, for ~~$120K in total (one set of rollers is good until ~~1 GW/year), in his CPV-17 presentation, Roger Angel quoted **$300K** for a set of glass-sheet-width lens-quality rollers for low-iron glass, so this more-pessimistic figure is used. NREL's CPV costing used hot embossing of silicone on glass, and embossing plates should cost between $2K and $30K (Luce and Cohen, 2010), so **$25K** is allocated; the labor rate for low-cost US manufacturing is not shown, but at $30/hr the 45¢ labor per lens is ~60 per hour, which with 500W lens tile ~125 MW/year for HCPV.

A heated press for substrate lamination is roughly $30K, and this is done as multi-spreader sheets so it is good for 1 GW/year. A receiver assembly jig is probably **$5K** for a set of 10 that is good for around 10 MW/year. The module assembly jig is much larger and should be around **$10K**, and is also good for ~10 MW per year. Miscellaneous trays for holding parts and jigs for placement equipment are allocated **$50K** for 100 MW/year.

In the non-custom equipment, two high-speed CNC dispensers for adhesives will be needed at **2*$25K** a piece; each is good for ~~30 MW/year. A high-speed tungsten carbide saw will be needed for singulating the lens tiles until volume makes custom wire sawing more cost-effective, but sawing is a fast process, and the edges can be rough so the parts-handling time is the limit; 5 minutes per cut would still be 12 kW/hour or >25 MW/year from a **$50K** saw.

A high-speed pick-and-place machine (**$300K** for 10,000 chips/hr (Ellis, 2019)) and a high-speed wire bonder ~**$100K** will be needed to keep costs low for even 10 MW production; each of these could handle ~150 MW/year (limited by parts handling). Moderate-speed placement equipment can be used for the medium parts (e.g., populated substrates) placement; a **$50K** machine is limited by manual loading to 75MW/year.

					10 MW			100 MW			1 GW
Part	Cost $K	Volume MW/yr		10 MW number	**10 MW $**		100 MW number	**100 MW $**		1 GW number	**1 GW $**
Lens/TIR Rollers	300	1000		1	**300**		1	**300**		1	**300**
TIR mold	25	125		1	**25**		1	**25**		8	**200**
AlSiC bonding	30	1000		1	**30**		1	**30**		1	**30**
Rcvr jigs	5	10		1	**5**		10	**50**		100	**500**
Mod jig	10	10		1	**10**		10	**100**		100	**1000**
Tray, etc.	50	100		1	**50**		1	**50**		10	**500**
Dispense	50	30		1	**50**		4	**200**		34	**1700**
Saw	50	50		1	**50**		2	**100**		20	**1000**
Pick-and - place	300	18		1	**300**		6	**1800**		56	**16800**
Wire bonder	100	60		1	**100**		2	**200**		17	**1700**
Medium Parts	50	75		1	**50**		2	**100**		14	**700**
Total					**970**			**2955**			**24430**

Table 5.4: *Table of tooling costs versus production per year.*

Probably additional parts of the manufacturing would be outsourced, but even with all these tasks done internally the estimated total capital cost is $970K for 10 MW/year, and ~$3M total to reach 100M/year. These are tiny numbers by PV industry standards, and the total capital cost of ~$25M for 1 GW/year are roughly *10 times lower* than comparable scaling costs for flat panels.

So where did all of the capital costs go? First, the capital costs really are reduced by TLC, to probably ~~20% of that of scaling silicon PV. But equally importantly, the capital costs of producing the major components: cells, mirrors, trackers, heat rejection, copper sheet and the actual rolling and molding equipment are born by suppliers, and thus already show up in parts costs rather than being capital that TLC manufacturing would have to raise for production equipment.

These indirect capital costs will not impede scaling because the industry already has production capacity for >= 1 GW/year for everything. At ~1 GW/year the CPV cell companies will have to scale production, and the NREL costing shows that to be $1/cm^2/year (0.6¢/W @ 30W /cm^2 = 20¢/cm^2, at 5-year depreciation for $1/cm^2/year. With TLC's ~60W/cm^2 that is only ~1.7¢/W/year, or a tiny fraction of silicon-cell scaling costs.

- ***Raw Materials and Scaling:***

Germanium wafer production scaling is probably the biggest scaling cost, but silicon has shown how wafer manufacturing costs plunge with volume.

Unlike CdTe, where First Solar's few-GW production already uses ~40% of the world's tellurium production and total known Tellurium reserves would enable producing only a few percent of the world's electricity demand, material scarcity should **not** be an issue with TLC.

At < 2 tonnes of Ge per GW (at 60W/cm^2), current Ge production would enable manufacturing >75 GW/year (Wiesenfarth, 2017), and Ge reserves are sufficient for ~20 TeraWatts of TLC. Furthermore, substrate reuse will reduce the germanium consumption markedly, so with TLC and substrate recycling the known Ge reserves are sufficient to meet the entire world's electricity production several times over, and to meet roughly the world's entire energy demand.

Similarly, even if thick GaAs cells are used with no substrate recycling, ~8 tonnes of Ga would be used per GW, and current Ga production would enable ~50 GW per year. However, the bauxite *already* being mined for aluminum could provide over 2000 tonnes per year of gallium (Frenzel at al., 2016), sufficient for 250 GW per year of TLC, with total gallium reserves in bauxite estimated at over 1 million tons (Kramer, 2006), enough to meet the entire world's total energy demand even before substrate reuse.

Indium is more complex because it can be used for top windows and conductors (ITO) and epitaxial layers (InGaAs, AlInGaAs, InP, etc.) well as being part of much-thicker indium phosphide substrates for some cells. Even for CIGS, in which indium is a major component, there is more than enough indium in economically recoverable reserves to meet the entire world's electricity demand (Wadia, et al., 2009), and at more than twice the efficiency and 1500x concentration even a thicker absorber layer would use trivial amounts of indium. However, if 350 µm indium phosphide base wafers are used, and are not thinned or reused, then allowing for a 150 µm wafer kerf TLC would use ~3.23 grams per kilowatt, or 3.2 tonnes per GW. 100 GW/year (the same scale as silicon PV today) would then consume almost half of the world's ~700 tonne/year indium production, or ¼ of the potential by-product indium production, so ***if*** future cells depend on indium phosphide wafers, then indium could eventually become a limit, at hundreds of GW for thick wafers. However, if the wafers are thinned to a few tens of microns then even with indium phosphide wafers TLC could add several TW to the world's electricity production per year and

quickly meet the world's total electricity needs and even total energy needs. And current wafers aren't even indium-phosphide-based, so indium scarcity is a non-issue.

Silver has been cited as a limitation for solar mirrors (Pihl, et al., 2012), but 100 nm is only 1 gram/m^2 or 2.5 tonnes/GW, so the world's current production would be enough for ~10 TW/year of TLC, and silver reserves are again enough to more than meet the world's total energy demand with TLC (besides which aluminum can also be used for solar mirrors with only ~5% loss of reflectivity and minor tuning of the cells). At 50 ¢/gram = 50 ¢/m^2 = 0.125 ¢/W, silver is an insignificant cost as well as not a significant constraint. Even the thicker 200 nm silver of some mirrors is only 2 gm/m^2, or 200 W/gram, which is an order of magnitude better than the 15-20 W/gram typical for silicon panels.

Commodity Materials: TLC uses less glass, steel and aluminum than flat-panel solar, so commodity-material shortages would hit silicon PV before they would hit TLC. The first limit would be float glass – the world produces ~65M tonnes/year, and at 11 kg/m^2 for the mirrors and ~1/8 as much for the lens; TLC will consume ~30 tonnes/MW so today's float-glass production would serve for 2 TW/year. Only a fraction of that is low-iron glass, but most of it would be convertible to low-iron glass. There is no overall shortage of materials from which to make glass, so multi-terawatt per year current capacity *that would limit silicon PV first* is not a concern.

5.5 Conclusion

This chapter has included a detailed analysis of the expected economics of TLC, with costs for high-efficiency, high-reliability TLC modules ranging from ~21 ¢/W at introduction (at 10 MW) to ~16 ¢/W at 1 GW, and approaching 10¢/W if manufactured on the scale of silicon PV. Even after adjusting for capacity factors and tracker costs, TLC has the potential to compete with silicon PV on overall cost in the key PV markets of utility-scale solar in high-DNI areas. Capital costs to scale manufacturing are also much lower than for silicon PV.

CHAPTER 6

6. Simulation and Analytical Validation

6.1 Introduction

This chapter covers the simulation and analytical work done toward validating TLC. It starts by presenting the analytical models on which the module design and the optimization for economics are based. This chapter then covers the 3D-CAD models built to refine the TLC design. These include thermal, optical, electrical, and mechanical models.

6.2 Analytical models

The design and analytical model of TLC are spreadsheet driven. The spreadsheet is based on the design and analysis discussed in the previous chapters of this thesis. As the multi-physics analysis refines the numbers, the spreadsheet is updated so that it continues to drive the CAD model. **The spreadsheet is archived on the online database of research documents, *Savoirs*, provided by the Université de Sherbrooke (https://savoirs.usherbrooke.ca/).**

6.2.1 Financial Tab

The Financial tab allows picking production volume, cell type, rating DNI, exchange rates, volume cost reduction, and tracker inclusion, and calculates the resulting cost per Watt. Most cost are based on quotes from suppliers for parts or materials, with material volumes and process areas calculated from the design itself. For example, the amount of low-iron glass used in the lens tile is calculated from the size and shape of the lens tile in the modeled design, with the width of the lens tile calculated from the focusing capability of the trough. In general, process costs are area based rather than volume based, with the base cost for a process derived from the closest process for which a referenceable cost has been found.

6.2.2 Half-Trough Shape Tab

The half-trough shape tab accepts the pre-bent mirror width, focal length and center offset of a half-trough, and calculates a matching parabola and the coordinates of the inner and outer edges of the trough mirror. Setting the focal length to an RP3 trough shows that the "1641 mm" mirror width of an RP3 inner mirror is the width of the glass before bending, and matches the RP3 map from a paper on RP3 mirror accuracy (Meiser, 2014), with a center offset of 40 mm.

6.2.3 Rim Angle and Receiver Slant

This tab takes the trough shape and calculates the rim angle and the optimal receiver slant based on fundamental optical principles and geometry.

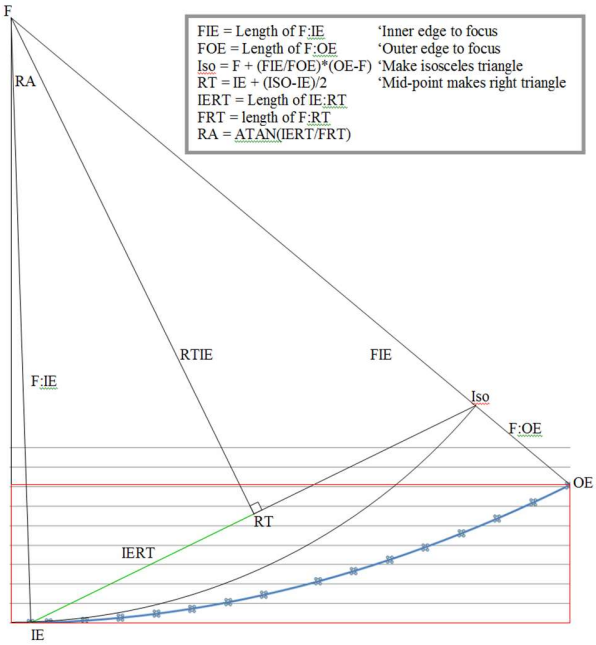

Figure 6.1: *Figuring out the Exact Rim Angle.*

6.2.4 Primary Optics Tab

This tab accepts mirror, mounting and tracking errors both on the trough's focusing and along the trough included, along with a 'spare' error budget, and calculates the trough's focusing. The 'spare' budget is currently set so that in the focusing dimension the trough tolerances are 3.1 times more

relaxed than the most accurate 82x solar thermal troughs and 1.7 times more relaxed than a typical 70x trough. The 'spare' budget is set to more than twice that along the trough's axis.

This tab also calculates the mirror cost and weight. Efficiency is currently in two spectral ranges (top two junctions, and non-limiting 3rd junction); this should be extended to more ranges.

6.2.5 Lens and Receiver Size Tab

The Lens and Receiver Size tab uses the trough's focal width and rim angle to calculate how wide the lens tile is. This involves the lens thickness and refractive index, and the depth of the lens-front curve. This tab then uses the thickness and density of the glass density to calculate the mass of the lens, and then uses the glass cost per kg to calculate the cost of the lens material.

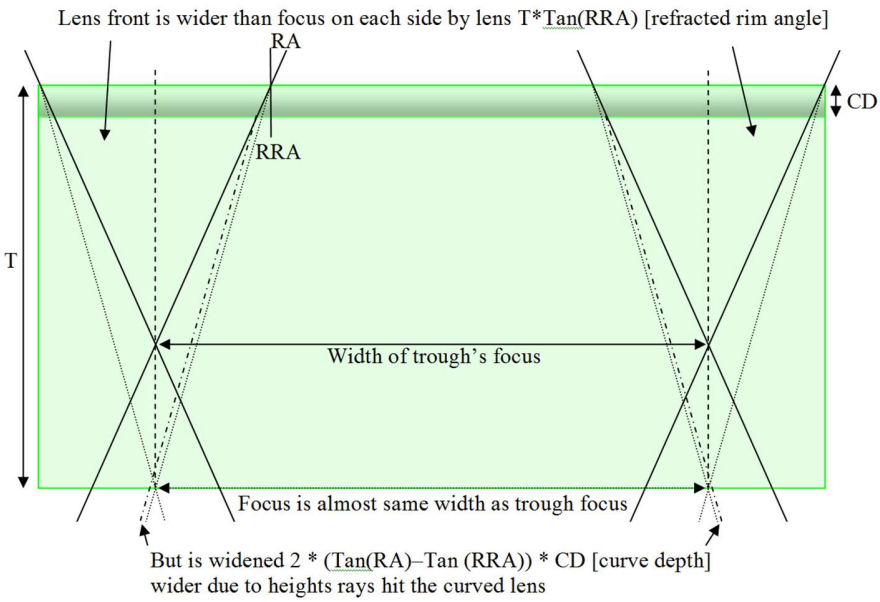

Figure 6.2: *Lens Width from Trough Focal Width.*

Roll-forming cost is calculated from the cost difference between patterned (roll-formed) and un-patterned flat-panel module covers; up to a GW this is multiplied by the lens thickness to account for the heating energy, but by 10 GW this is assumed to be in the float line. AR coating cost is based on flat panels, and lens-tile singulation cost is based on silicon boule wafering per-area cost.

- ***Focal displacement:***

Because the lens front is flat in the direction of the trough's focusing, the lens front does not increase the concentration in the direction of the trough's focus. The higher refractive index of the lens (relative to air) does, however *displace* the trough's focus by reducing the angle at which the rays from the trough are converging.

The Lens and Receiver Size tab calculates this displacement. To a first order, the focal displacement is proportional to the lens thickness times the difference between the angle from the trough (the trough's rim angle, or RA), and the refracted angle of those rim rays, or RRA. However, the thickness of the lens is not constant because the lens front is curved and adjusting the thickness by half the depth of the curve minimizes the focal broadening that this effect produces.

The 'Lens and Receiver Size' tab also calculates the lens tile length based on the modules per mirror, and then calculates the lens length based on the number of lenses per receiver and the number of receivers per lens tile.

The 'Lens and Receiver Size' tab also calculates the optical losses in the lens tile glass based on the optical path length in the glass and the transparency of the glass. It does this for two wavelength ranges because current CPV cells have excess photocurrent in the wavelength range where the glass absorbs more. The glass absorption and the AR coating reflection are from PPG Solarphire data (now Vitro Starphire), adjusted for the transparency of the glass for which pricing was received. This is probably pessimistic because the cost of the Aohong Glass is so low (<0.2 ¢/W) that even if the clearest Starphire glass cost twice as much as the best Aohong Glass, the slight increase in efficiency would decrease the cost per watt.

6.2.6 Lens Curve Tab

- ***Calculating the Lens Focal Length:***

  - *Lens Tile Rolled from a Sheet (<v2.0):*

The Lens Curve tab until analytical model spreadsheet v1.9 assumes that the lenses, the cone-row cavities and the body of the lens tile are roll-formed in standard-thickness glass sheet, and that the lens focal length thus depends on the height of the cones. Since the shape of the cones depends on the focal length of the lenses, this requires 'convergence' by manually copying the value from one

cell to another to let the spreadsheet update without a loop, but a few iterations are all that is necessary to get the lens shape stable to <1 μm. Also included in the lens focal length is the height of the cell (including solder) relative to the height of the substrate (including adhesive and copper).

- *Lens Depth Independent of Cone Height (v2.0+):*

In version 2.0 of the unified TLC analytical model, the lens calculations were modified to allow separate refractive indices for the lenses, bulk of the tile, and the cones (a separate cone refractive index had been hacked into several previous versions). Since in volume roll-forming would be done from liquid glass rather than glass sheet, and since in early prototyping the lenses as well as the cones will be over molded, the dependence on a constant sheet thickness was also removed.

It has been found that the cones handle under-focused light better than over focused light, so the ability to have the lenses' tightest focus partway into the cone was also added to the Lens Curve tab. This is currently a hand-tuned parameter based on COMSOL ray tracing, but a way to approximately calculate it has been found and will be added. This comprises noting that the worst-case light is that which comes from the edges of the trough, which maximize the Petzval effect and thus produces the shortest focal length, it also has the bluest color, which also minimizes the focal length. By calculating the focal length for this worst-case light, the amount of subsequent hand tuning (done via ray tracing) can be minimized.

- *Petzval Effect:*

The Petzval effect of the rays from the trough rim is a critical factor. Rays from the trough rim come into the lens at a wider angle to the surface normal and hence see a focal length shortened by the cosine of their angle to the lens curve, so the optimal lens focal length is adjusted by the square root of the maximum Petzval effect to match a 'geometric average' ray.

- *Lens focal width*

Based upon the lens focusing properties, including the sun's optical width, the mirror and tracker errors, the lens thickness chromatic aberration, and the Petzval effect, the lens curve tab then calculates the lens focal width, which is fed into the cone calculations in the following tab.

- ***Calculating the Lens Curve:***

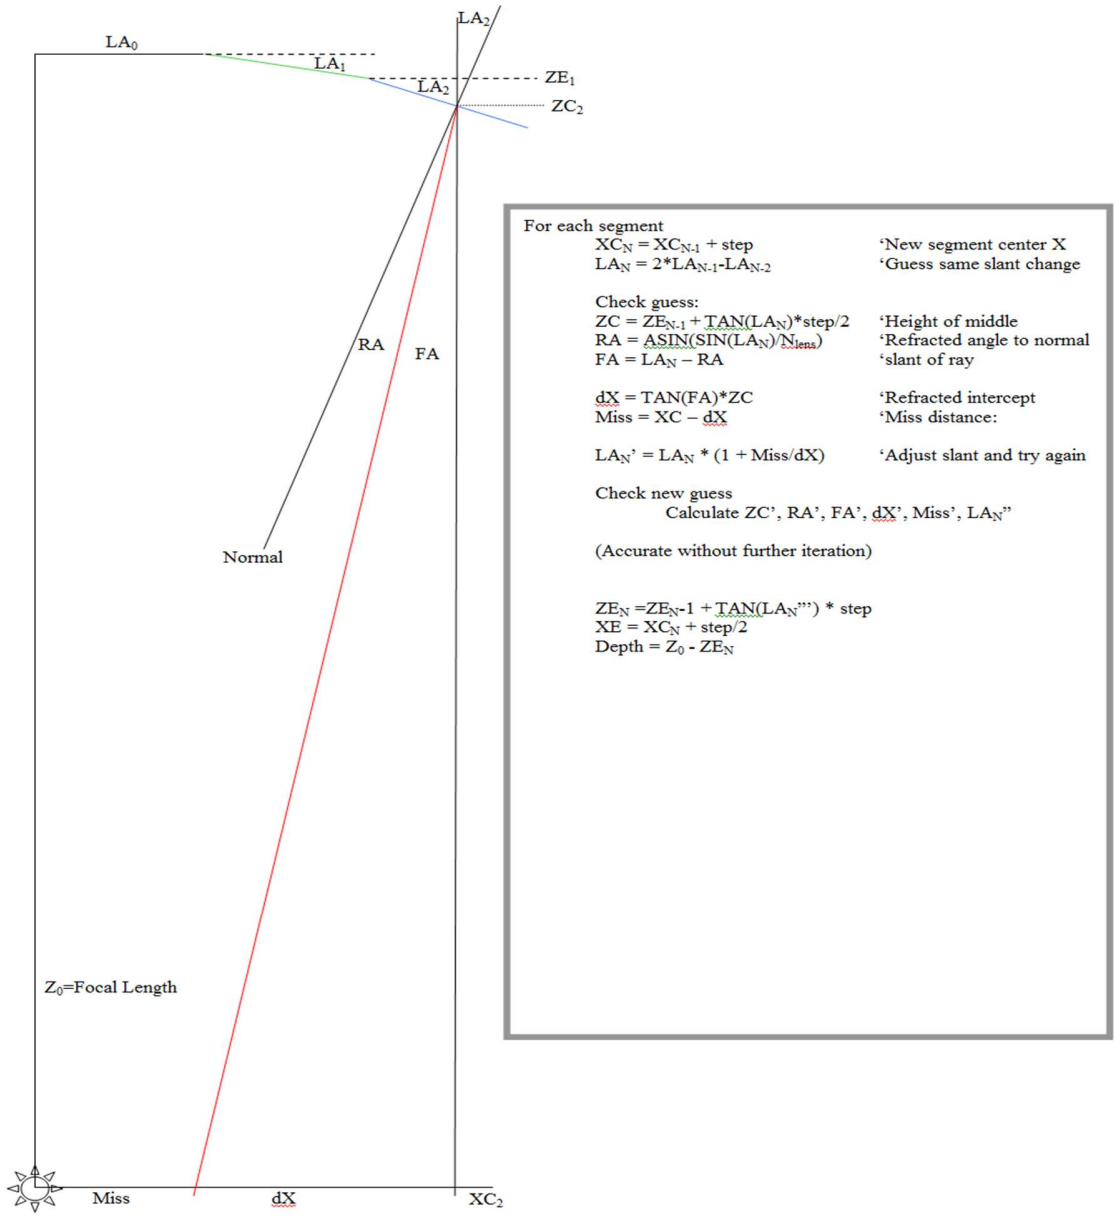

Figure 6.3: *Calculations for the lens curve.*

The lens shape evolves as a segmented curve that always refracts the average ray to the focal point. The focal width around that point is calculated and takes into account the chromatic aberration of the lens and the Petzval effect as well as the sun's size and the trough inaccuracies.

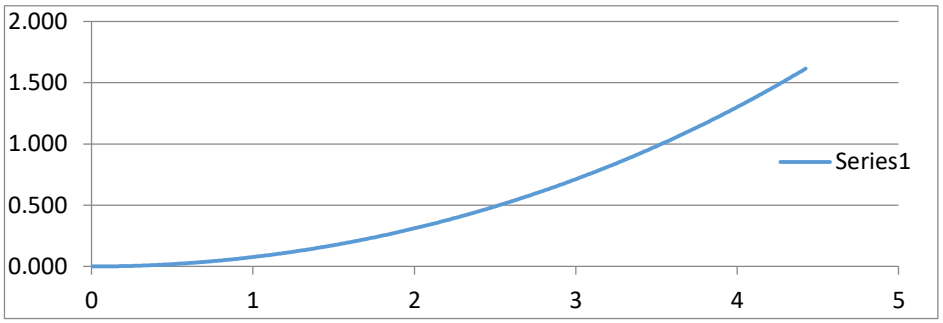

Figure 6.4: *Curve of half-lens (Both axes in millimeters).*

6.2.7 TIR Cones Tab

The TIR Cones tab accepts a cell size and uses the angle of the rays as focuses by the trough and the lens to calculate compound parabolic curves that further concentrate the light onto a cell-sized rectangle. The final concentration, the number of cells per row and the cell length can be specified, along with the alignment accuracy and the coupling thickness. The minimum angle for light hitting the cell after one reflection can also be specified to avoid shallow angles that the cell's AR coating can't handle effectively. The analytical model spreadsheet then calculates the CPC curve to keep light from going beyond the edge of the cell after the rays pass through the optical coupling.

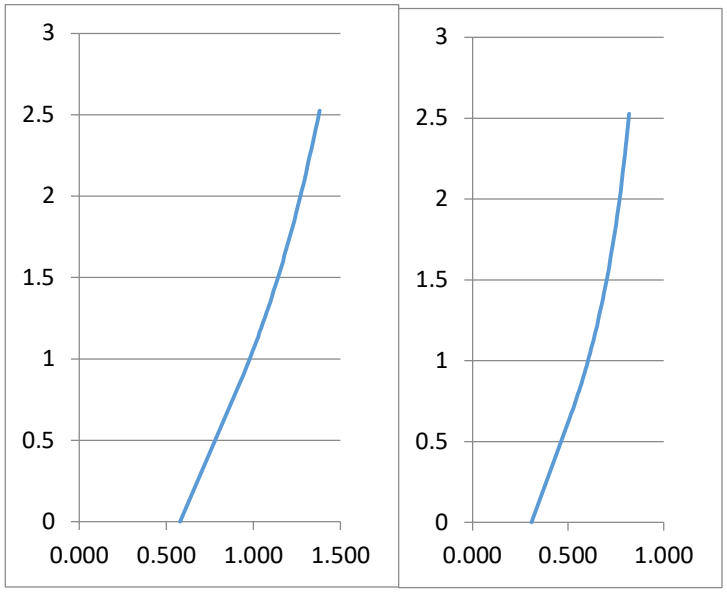

Figure 6.5: *Initial TIR CPC cone shape (in millimeters), half-cone x-section on both axes.*

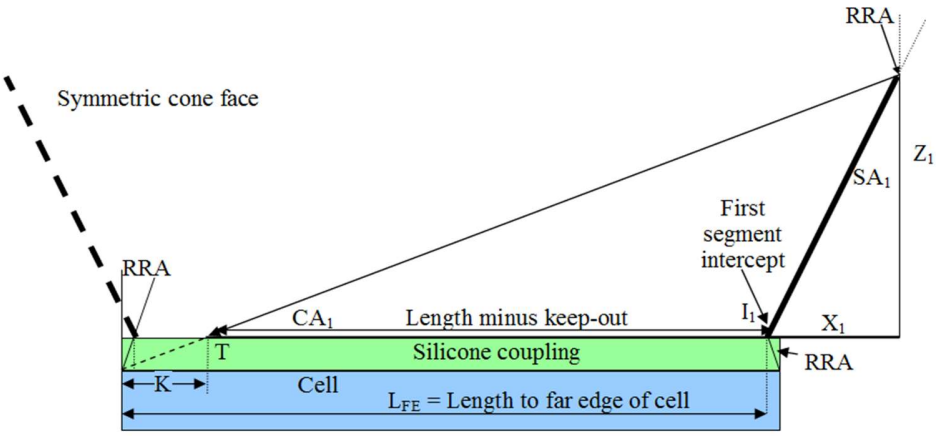

Figure 6.6: *Calculating the first TIR facet.*

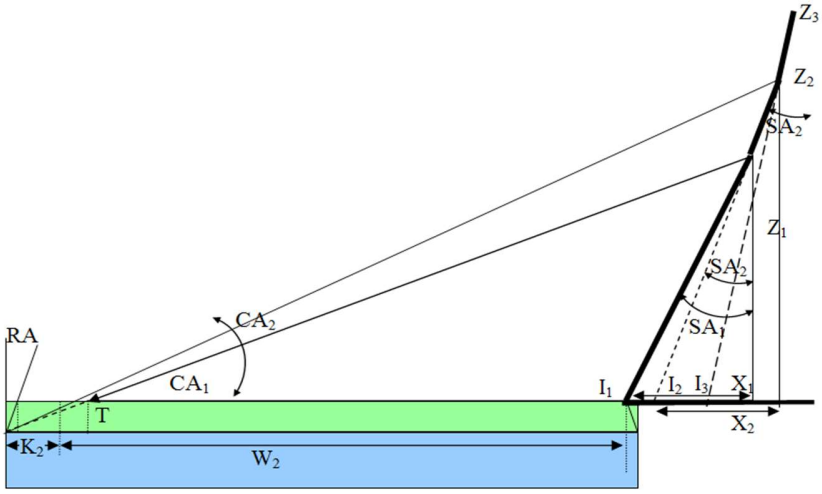

The next TIR segments can then be calculated

Loop:

$SA_N = SA_{N-1} + Step$	'Step to angle of next mirror segment
$I_N = I_{N-1} + X_{N-1} - Z_{N-1}*Tan(SA_N)$	'Find where next mirror would intercept cell height
$CA_N = 90-RRA-2*SA_N$	'Minimum angle of light on cell from segment N
$TanTan = Tan(SA_N)* Tan(CA_N)$	
$X_N = I_N * TanTan / (1-TanTan)$	'Calculate how far light gathering is extended
$Z_N = X_N * Cot(CA_N)$	'Angle and distance from cell edge gives height

Figure 6.7: *Calculating subsequent facets.*

The TIR cones tab then calculates the volume and the cost of the silicone and the cost of the over-molding, using the NREL CPV costing study (Horowitz, et al., 2015) as a reference for costs.

6.2.8 Substrate Tab

The Substrate tab is a relatively simple one. It allows an extra width (for heat spreading for the end cells, or for electrical connectivity in sawn substrates) to be specified, as well as a singulation kerf, and also the costs and thicknesses of the heat spreader, isolation and top-copper layers and the bonding and plating costs for the substrate sheet. It then calculates the cost of the substrate, including cutting cost.

6.2.9 Cell & Attach Tab

The Cell & Attach tab allows specifying the cell dicing kerf and the wafer cost and yield, and it then calculated the cell cost. The Cell tab also allows cell V_{MPP} and I_{MPP} and raw efficiency and temperature sensitivity to be specified (based on the manufacturer's datasheet for the closest cell size for cells made with the same process, and adjusted for TLC cells size and gridlines in the "Small Cell" tab), and it the grabs the delta T from the thermal calculations (copied to prevent looping), and calculates the efficiency based on the optical efficiency for the current-limiting junctions. This is then adjusted by the resistance of the substrate and interconnect (from the Electrical Tab) to get the overall electrical efficiency. This is calculated using the reference DNI from the Financial tab to get the current and the thermal load.

The Cell & Attach tab then calculates the cost of the optical coupling (here because it is cell-sized), solder, wire bonding, cell placement and diode.

6.2.10 Electrical Tab

The Electrical tab allows specifying the number of bond wires per cell, and the bondwire size and resistivity. It then calculates the resistance of the bond wires, top copper, substrate and interconnect strip, and also the module's terminal wires. It then uses the cell's maximum power point current and voltage at the calculated temperature to calculate the electrical losses based on the resistance. The module terminal wire and connector costs are also calculated.

The Electrical tab then calculates the short-circuit current and open circuit voltage of the module, based on the number of cells in parallel per receiver and the number of receivers in series per module, and calculates efficient module packing onto a tracker for a given inverter voltage range.

6.2.11 Thermal Conduction Tab

The Thermal Conduction tab calculates the conduction through the cell and solder; spreading in the heat spreader is then calculated using a slightly pessimistic short-cut, and then straight conduction is calculated again through the high-voltage isolation and the metal lid base. Currently the thermal resistance of the fins to atmosphere is taken from the Heat Rejection tab.

Because the heat to conduct is slightly dependent on the cell efficiency and the cell efficiency depends on the temperature, a manual copy convergence is needed to prevent Excel from looping.

6.2.12 Housing and Heat Rejection Tab

Because the heat-sink lid is part of the module housing, this tab also includes structural elements such as braces and supports as well as the heat rejection. Several alternatives cooling architectures are costed in this tab, based on the amount and price of materials used and the estimated process costs, but only the favored architecture's cost is exported back to the rest of the analytical model.

This tab uses an empirical formula for heat transfer coefficients combined together by curve fitting to data points from an online calculator and adding on a wind-speed term. The main heat transfer has been confirmed in COMSOL (the fin thicknesses were tuned until the worst-case cell temperature was comparable to that in the NREL Fresnel CPV reference). However, the wind direction effect is not yet confirmed. This is a priority to replace with estimates from COMSOL, that include windspeed and tracking angle, and later with real experimental data.

6.2.13 Module Assembly Tab

The module assembly tab calculates the cost of the various adhesives needed to assemble the module (although the optical coupling is applied at this time, it is cell-sized and so is in the cell attach tab).

The assembly non-material cost is calculated from the comparable steps in the NREL CPV costing study. The process for estimating this cost is described in detail in the Module Assembly section of this thesis.

6.2.14 CTE Stress Tab

The CTE Stress tab calculates the average CTE from the thicknesses and stiffnesses of the bonded layers. This average CTE is then used to calculate the stress that each material undergoes in matching the average CTE on cold winter nights and hot summer days on sun. This tab is due for revision once the heatsink lid is tested sufficiently to confirm the intended range of material thicknesses.

6.2.15 Equipment Tab

The Equipment tab allows specifying the dedicated equipment needed to produce TLC modules, the production volume achievable per piece of equipment, and the cost of each piece of equipment. The equipment tab then calculates the total cost of the equipment as a function of annual production volume. It then uses this to estimate capital requirements for the equipment to reach various production levels.

6.2.16 Tracker Tab

The Tracker tab contains the tracker cost. Information such as how the mirrors pack on the tracker frame and the tracking speed and wind limits will be added for a variety of trackers.

The Tracker tab also checks for tropical noon tracking issues – it calculates the tracking error based on the maximum tracking speed and allows comparing that to the maximum acceptance angle, which has been updated following the COMSOL ray-tracing.

6.3 Multi-Physics Modeling and Analysis

6.3.1 Modeling and Simulation Overview

Much of the initial TLC work focused on modeling in Excel and then simulation in COMSOL to allow optimizing the TLC optics and cooling. After modelling in Excel and/or Fusion 360 3D CAD software, the following analyses have been done in COMSOL multi-physics software (green dot) and/or Excel (black dot):

	Optical	Thermal	Mechanical	Electrical
• Trough mirror	●●			
• Lens tile glass	●●	●●	●●	
• Lens tile silicone	●●	●	●	
• Mechanical lens adhesive		●●	●●	
• Optical coupling silicone	●	●	●	
• Cells	●	●●	●	●
• Wire Bonds	●	●	●	●
• Diode		●		●
• Top copper		●●	●	●
• Under-copper insulation		●	●●	●
• Under-cell solder		●●	●●	●
• Wiring (Tabs, module ends)			●	●
• Heat Spreader		●●	●●	●
• Thermally-cond. Epoxy		●●	●●	●
• Lid		●●	●●	
• Fins		●●	●	
• Mechanical Epoxy			●	

The optical analyses include:

- Focusing
- Chromatic aberration
- Spectrum
- Miss-pointing
    - Efficiency / Acceptance Angle

    - Concentration off-cell when light misses cell

The thermal analyses include:

- Fin area and heat rejection
    - No wind
    - CSTC
    - Effect of trough altitude and azimuth angles
- Cell heating
- Copper heating
- Lens heating (with glass and with acrylic)
- Off-cell heating when light misses cell

The thermo-mechanical analyses include:

- Lens->Silicone->Cell Array-> Copper Substrate
    - Lens warping
    - Stress in mechanical epoxy
- Cell->Solder->Copper->Isolation->Lid
    - Stress on cell stays compressive
    - Stress in solder – ensure that it can last >50 years
    - Stress in lid attachment

The electrical analyses include:

- Grid lines - Voltage drop vs position on cell
- Bond wires
- Top copper and base copper
- Voltage drop vs cell position
- Terminal wire

- ***What to Model, and Why:***

The high-level design was first modelled with spreadsheets, but these did not have enough detail to pinpoint regions of stress concentration, non-typical conditions like focal intensity when the module is miss-tracked and could even have contained invalid simplifying assumptions.

The highest risk was thought to be thermo-mechanical stress; it is full of subtleties like stress concentration initiating cracks and temperature cycling propagating them, making long-term reliability strongly dependent on stress levels. Numerous alternatives for materials and design features are available for most areas, with the only highly constrained high-risk area being the cell-to-solder-to-copper region, so this was the most critical region to model first. This has passed its first modeling and initial-stress analysis.

Next on the risk scale was the details of the optics. While the multi-stage optical principles being able to reach sufficient concentration has been confirmed by rigorous analysis (Cooper, 2014), and while an earlier single-axis-tracked design was confirmed by ray-tracing, the current design had not been ray-traced. The optics have now been modeled in COMSOL and ray-traced, eliminating this concern. This even included details such as lens heating, as well as heating of the regions that the focus can land on when the module is miss-tracked.

Next on the list was the heat rejection. The basic fin performance has been confirmed in COMSOL, removing that risk, and cooling versus tracking angle was measured on a prototype heatsink. Stress on cooling a lens tile bonded to a prototype of the heat-sink narrowed down the list of manufacturing options to multi-fin segments bonded to a one-piece 400-series-stainless-steel lid (while a one-piece aluminum lid appears possible, the stress-balancing it requires is too risky for initial manufacturing).

Finally, a full 3D CAD model was used before prototyping to ensure that all positions (especially thicknesses) matched up properly.

6.3.2 Simple Heat Flow

The simplest useful model was to confirm the effectiveness of the heat spreader. This has been calculated by the analytical model spreadsheet with shortcuts thought to be slightly pessimistic, and the COMSOL analysis confirmed that that is indeed the case.

The lens and silicone cone were then added. The geometry is slightly conservative (no bulges on the cones), and the heat loading in the lens and cone is also pessimistic (all heat applied to the centerline, which creates additional lateral resistance). The results agree very well with the analytical model spreadsheet.

- ***Sanity-check Geometry:***

The cell is a simple germanium rectangle, 0.9 mm by 0.648 mm by 190 µm thick, which sits on a similar-sized rectangle of SAC305 solder ~60 um thick. On top of the cell sits a silicone cone 2.5 mm tall, tapering outward as it goes up (the final cones were 3.7 mm tall).

The cell sits on region 2.775 mm x 8.848 mm. Directly below the cell was an AlSiC slab area 1 mm thick (which has since been replaced by ½ mm of copper), which sits on 0.3 mm of thermally conductive adhesive and then the 1.5 mm thick stainless steel of the heat-sink face. On top of the cone is the glass lens (modeled in multiple parts merged to contiguous glass).

All of the materials have known thermal conductivities - several choices are available for the adhesive, so the 1.6 W/mK currently in the analytical model spreadsheet is used. By symmetry only ¼ of a cell and its surroundings is analyzed, providing a convenient cross-section.

- ***Heat Flow:***

The average cell receives ~0.75 W/mm^2 typical case (after optical losses and too-long IR reflection, and with some power converted to electricity). The opposite face is the base of the heat rejection – its temperature is set to 300K to make the temperature differences easy to see.

- ***Results:***

The temperature differences between the components match the spreadsheet's predictions, with a slightly lower temperature difference within the substrate due to removing the shortcuts that the spreadsheet uses.

It can be seen that the surface of the cell is in good agreement with the spreadsheet's prediction of a temperature ~12K above the heat sink, and a lens temperature ~10K above the cell temperature, and that the lens is hottest near the middle but is also only ~10K above the cell surface temperature.

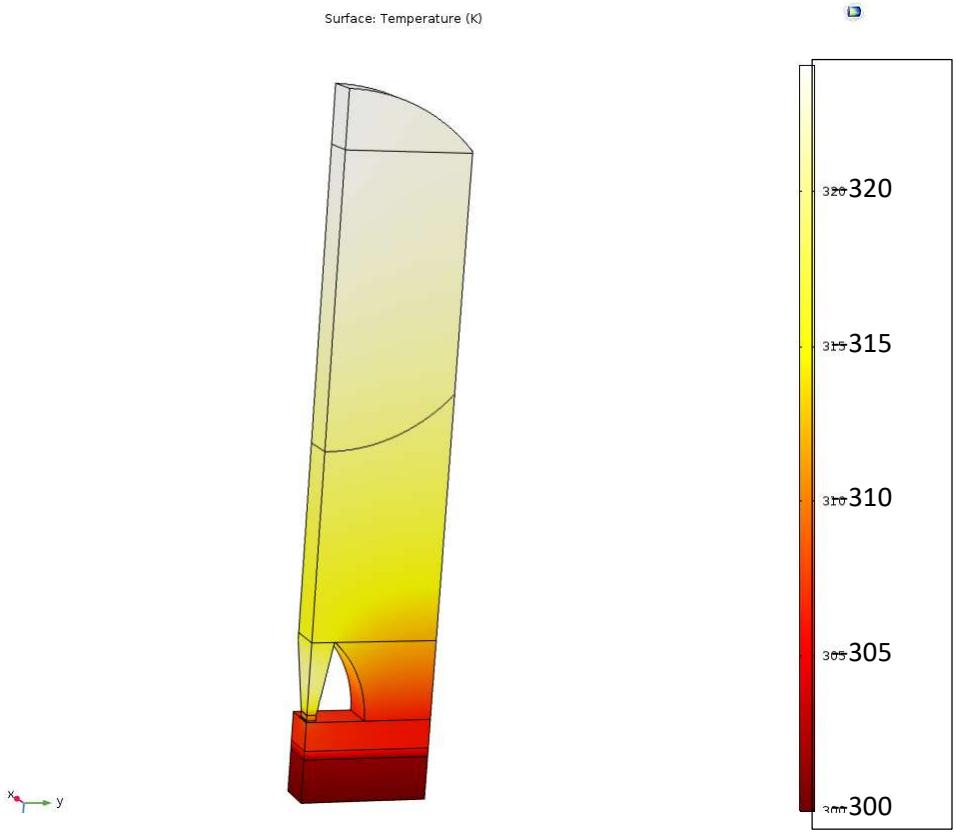

Figure 6.8: *Thermal Analysis of TLC Unit Cell (heatsink set to 300K).*

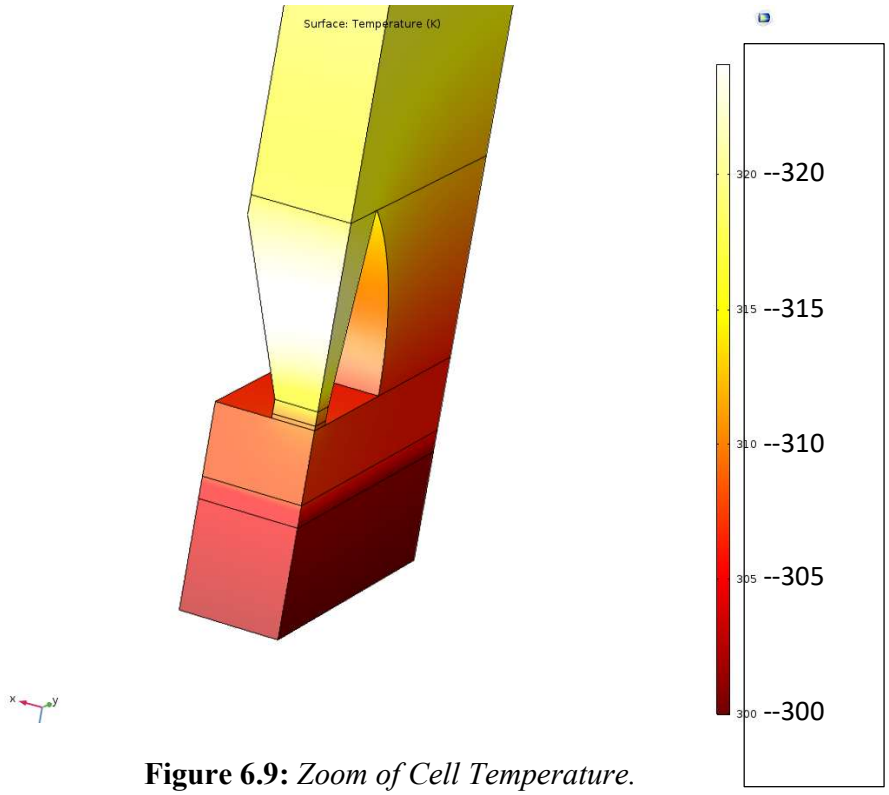

Figure 6.9: *Zoom of Cell Temperature.*

- ***<u>More Detailed Thermal Performance:</u>***

Temperature is critical to accelerated aging, and since aging is exponential with temperature, the near worst-case of full sun (1000 W/m^2 DNI) at tropical noon with zero wind was modeled, with an ambient temperature set at 300K (~27ºC), which makes the temperature-above-ambient easy to see on plots. 42% cell efficiency, 60 µm of SAC305 solder, and 70µm of 2.2 W/mK thermally conductive electrical isolation to the heat rejection were modeled. Trough reflectivity of 97% and lens absorption of 5% (mostly in the unusable infrared) were modeled; with lens heat as additional cell heating (in addition to the 58% of the light that is not converted to electricity becoming heat). The uneven illumination from across the trough's focus was modeled as 2000X on the central 12 cells per row and no illumination on the outer 4 cells, an average of 1500X.

In accordance with the then-current TLC design, the heat rejection was modeled as folded aluminum fins on a 1.5 mm thick alloy-410 stainless steel lid (CTE = 9.9 ppm/K, thermal conductivity 24.6 W/mK). The fins were 140 mm wide (projecting 20 mm less on one side to avoid shading the mirror) by 72 mm tall and 1.7 mm thick, on an 8.85 mm pitch. Emissivity was set to 0.77 for dark anodized aluminum sheet, and fin thermal conductivity was set to 201 W/mK. The fins were connected by receiver-width fold zones at both the top and the bottom, with the rest of the width open. The convective heat transfer was modeled as flowing between parallel plates, and the radiative heat transfer was suppressed for the exterior fin faces (since they will be next to other fins). COMSOL's built-in heat transfer functions were then used to model the heat transfer from the array of illuminated microcells to the 300K ambient air.

The central cell of a row reaches ambient + 50.3ºC, with the coolest illuminated cell 3ºC cooler. Even at 2000X peak concentration, the substrate next to the cell is only ~1.5ºC warmer than between the rows, confirming the rapid heat spreading from the microcells. TLC's fin orientation is worst when the sun is only 25º above the horizon, but the maximum DNI then is only about 700 W/m^2, so the TLC cell temperatures are slightly cooler than at 1000 W/m^2 at noon.

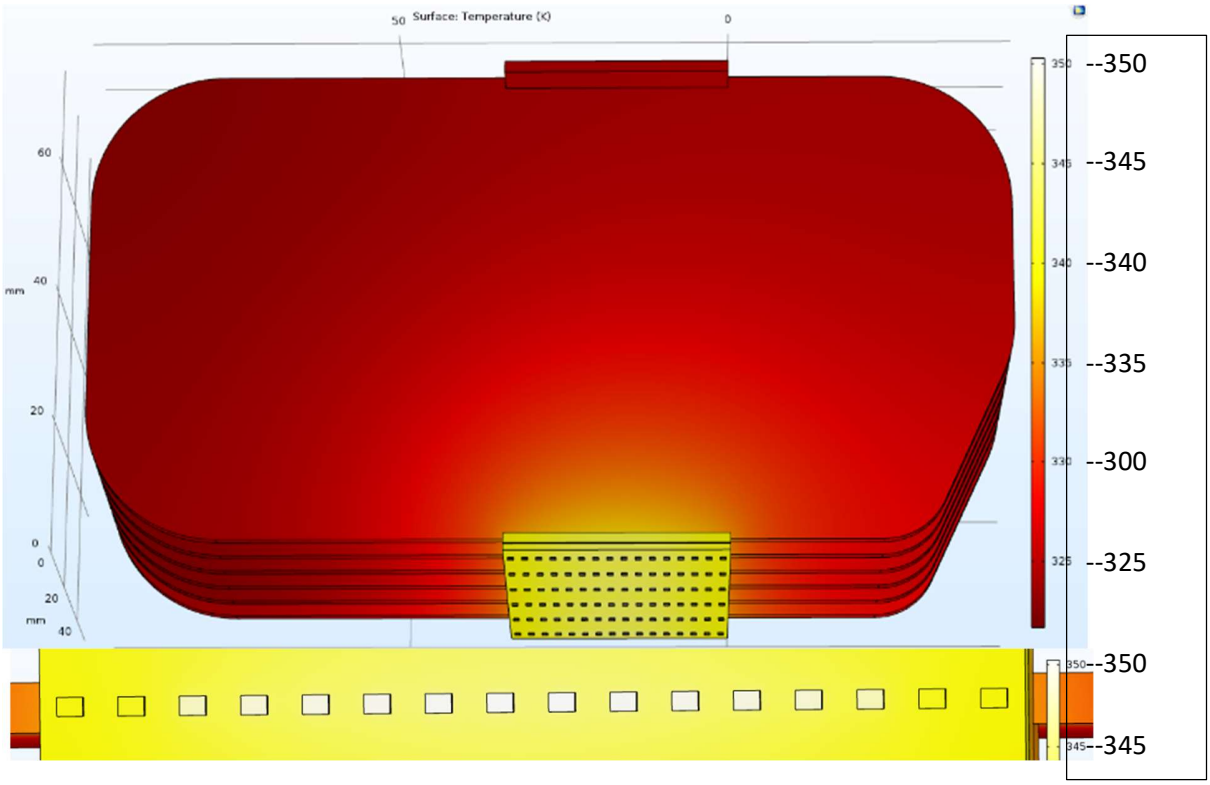

Figure 6.10: *TLC Cooling in COMSOL model: Overall (top), zoom on cells (bottom).*

- ***<u>Simplest Useful Stress Model: (needs update):</u>***

The key stress is expected to be in the solder underneath the cell, and this was be analyzed to a first order by modelling a simple ¼-cell stack-up of cell on solder on copper and modeling a temperature change from no stress at solder solidus at ~220ºC, to maximum stress on a cold winter night (typically -30ºC but modeled at -40ºC for Sherbrooke).

- Model:

The cell is treated as a simple germanium rectangle, 0.9 mm by 0.648 mm by 190 µm thick, which sits on a similar-sized rectangle of SAC solder ~60 um thick, and for this first simple model only the 0.5-mm-thick copper directly under the cell is modeled. The ~260K temperature change from solder solidus to a cold winter night will be modeled first.

- Expected Results:

Unlike DBC/AlN, hard copper is very strong even under tension, so there are no worries about the integrity of the copper spreader itself, as long as it is hard copper so that it will *not* relax and then harden at low temperatures and then expand to put stress on the cells.

The thin, relatively soft solder has a higher CTE than the thicker, stiffer cell or the much thicker and stiffer copper, so the faces of the solder are essentially fixed. The copper's CTE of ~17 ppm/K dominates, and is ~6 ppm/K below the solder; for ~1500 ppm of relative solder shrinkage at 260K temperature change.

With a 50 GPa modulus this would put ~75 MPa on the solder if it were a large sheet. Fortunately, the cell size is small, but even so the solder is 10 times broader, and 15 times longer than it is thick and at least in the middle it will plastically deform to relieve excess pressure. I'm expecting the solder to plastically deform to release all but its after-aging yield strength of ~35 MPa of the stress at the middle and deform progressively less toward the edges where the shear stress lags. Without the non-linear materials physics the solder deformation will not be shown, but a stress map will show where deformation will occur.

The cells have the lowest CTE and shrink the least, so the cell stress should be compressive. The brittle cells are very strong under compression so the only cell concern on cooling is fracturing a layer off the bottom of the cell where the solder binds the cell to the copper. However, the CTE mismatch is not large, and the cell will sit in the solder's shear lag zone (He, et al., 1998) so this should be very far from being an issue.

The basics of this were been confirmed via COMSOL before the changes from AlSiC to copper, and while this should be reconfirmed, the stress on the solder is lower on copper.

- ***Rewarming Model:***

The solder should relax to leave roughly its yield strength in residual tensile stress and be will then be in equilibrium when re-warmed to roughly 70ºC warmer than a cold winter night, or roughly 40ºC in most climates. The solder will then build stress at higher temperatures and will start to put tensile stress on the cell. Assuming that the solder behaves this way right to the edges of the cell produces a worst-case bound on the stress, so the second run with this model will start with equilibrium at 40ºC and warm the solder to ~100ºC (grid down on a hot summer day).

- Expected Rewarming Result:

A large area of solder would build ~20 mPa of stress, but with the cell and solder in the shear lag zone the stress in the solder should not exceed the yield strength. The cell is >3x thicker than the solder and the copper will take much of the force, so the stress in the cell should be ~5 mPa. Since that is only ~4% of germanium's tensile strength, stress testing is below optics in priority.

- ***Optical Model (completed):***

The next step was to use COMSOL to ray-trace the optics to confirm the optics and analyze details like acceptance angles and the initial focal unevenness.

Multiple lens rows were modeled, with the lens and cone curves spreadsheet driven. The light from the trough was modelled as reflected sunlight to get the correct angular distribution of the incoming light. Initially air-to-glass reflection and internal absorption were calculated rather than modelled.

The optical materials have refractive indices modeled for the wavelength extrema and for the central wavelength for each junction so that chromatic aberration can be evaluated. Absorption and mirror reflectivity will be coarsely modeled as modifying the initial spectrum of the sunlight.

As expected, the lens front focused the light from the full angular range into the mouths of the cones. It was also confirmed that the cones funnel all single-reflection light onto the cells at an angle of less than 60º to normal, and that almost all two-reflection light reaches the cell (and three-reflection light is not expected). The light distribution is fairly even over a given cell, and matches the trough's natural focal intensity from cell to cell. Since the ray traces focused well, the prototype processing was started.

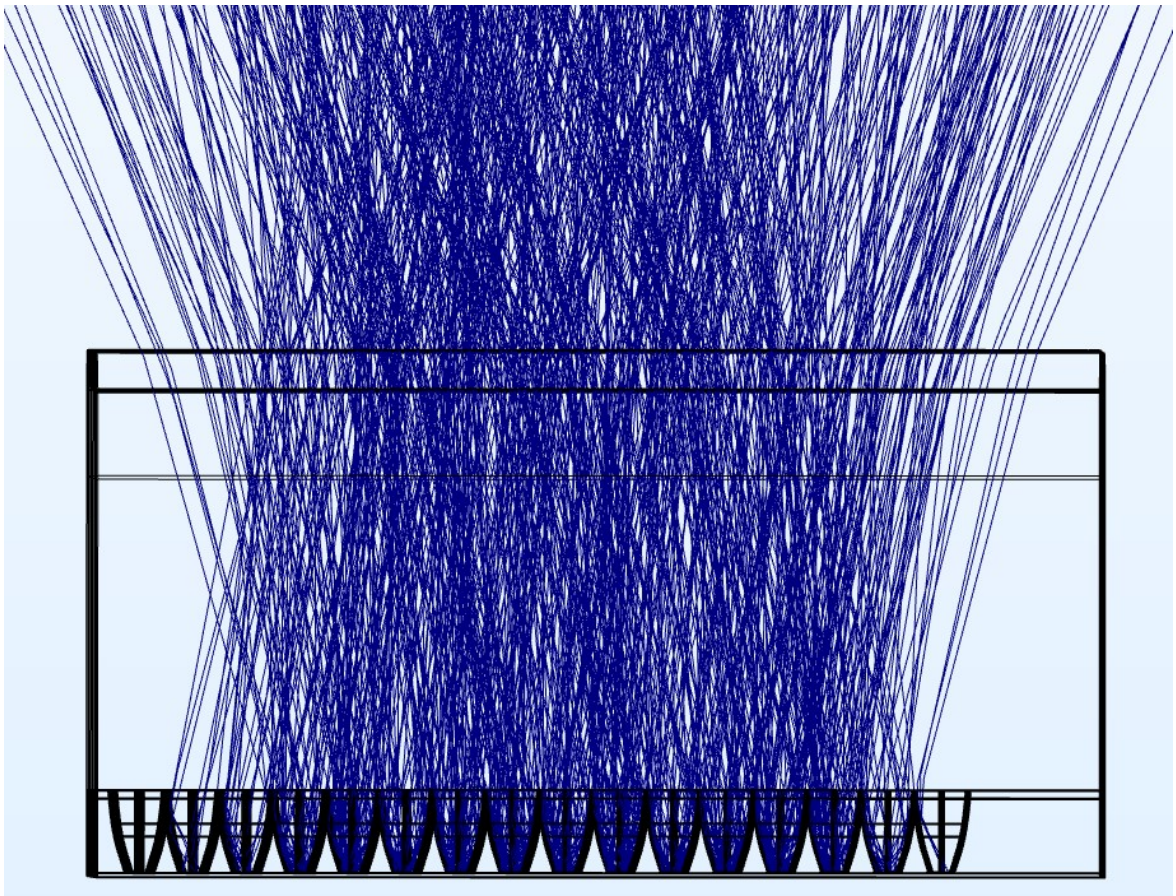

Figure 6.11: *Ray-traced focus* ***across*** *the trough's focus*

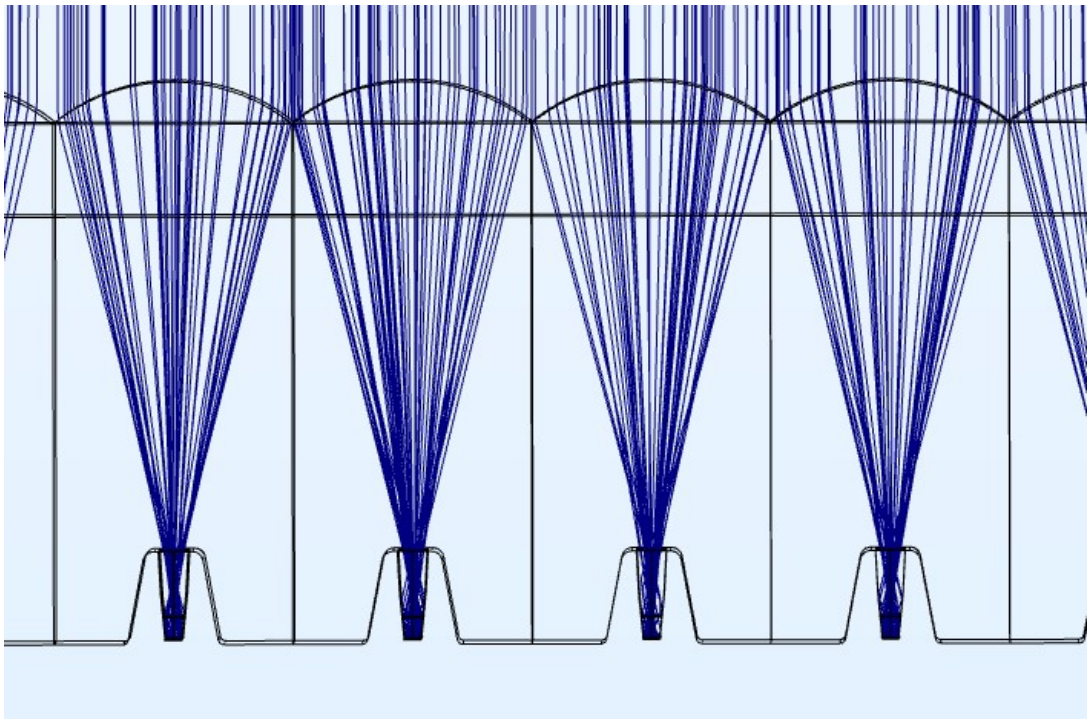

Figure 6.12: *Ray-traced focus* ***along*** *the trough's focus*

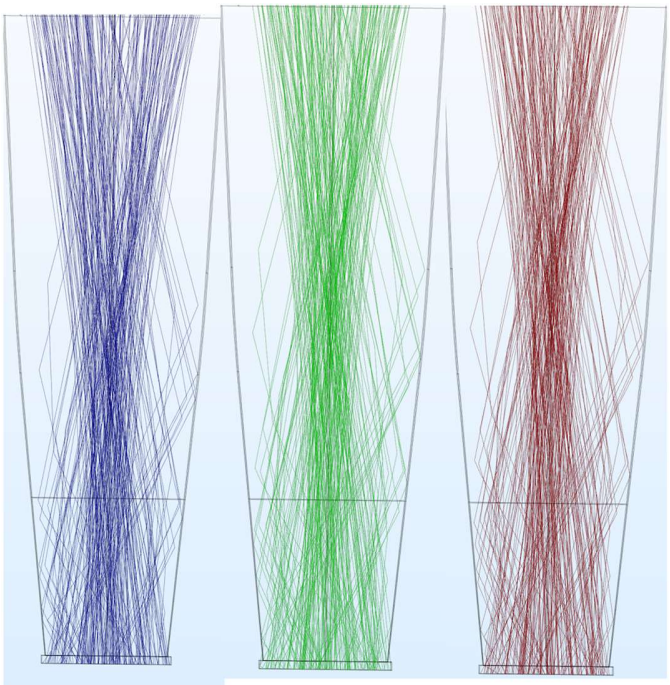

Figure 6.13: *Ray-traces at the center wavelength of each junction*

6.3.3 Full Thermal Stress Model, Flexible Adhesives

As the prototyping progressed, the thermal stress models were refined. The lens tile glass is the bulk of the stiffness of the whole module, so its coupling to the other parts is key. The plan was to couple the lens to an aluminum lid with stiff structural epoxy, and even with thick epoxy and a 50K temperature change, the module is so long that only the very ends will be in the shear lag zone.

While the cell is decoupled from the lens through the soft silicone, the lens was expected to be coupled to the thick copper substrates with structural epoxy. This would slightly increase the stress on the solder, and significantly decrease the stress on the cells. Even a hard epoxy is much softer than solder, but the CTE mismatch is much smaller, and the temperature change will only be ~50K, and for a 50 µm bond line the shear lag zone should be on the order of a centimeter. That is smaller than a receiver, so the expectation was the epoxy to bind the spreader to the lens but for there to be some displacement toward the outer edges.

However, while modelling showed that balancing CTE-mismatch stresses could keep stress on the glass compressive while almost eliminating warping, the balance was sensitive. With the too-thick heat-sink stock obtained for prototyping, and without end-caps to convert shear stress to compressive, and without toughened glass, in the first bonding experiment the heatsink sheared a layer off the lens glass upon cooling. Hard adhesives were therefore replaced with flexible adhesives for the first mini-module prototypes.

Flexible adhesives proved successful, so the initial full-length module will use flexible adhesives. Over the ~120K temperature swing from winter night to summer on-sun, the ~14 ppm/K CTE difference from the glass would amount to +/- ~0.7 mm if the aluminum heatsink were monolithic, which would require unacceptably thick (cost and thermal resistance) adhesive. Therefore, the monolithic alloy 410 stainless-steel lid was brought back, with narrow (receiver-length) fin extrusions attached with flexible adhesive. The steel's CTE of 9.9 ppm/K cuts the differential expansion to the glass to 1.3 ppm/K and an acceptable +/- ~0.07 mm, and the steel's ~13 ppm/K difference from the aluminum fins produces +/- 25 um shifts. The steel's 7 ppm/K difference from copper shifts the receiver corners by an additional +/- ~15 um relative to the steel, for ~80 um total shift, which thermally-conductive silicone can handle. The glass-to receiver adhesive can also be flexible, but should be significantly stiffer than the glass-to-steel adhesive to avoid the steel shifting the receivers relative to the optics (the cones will flex to handle residual shifts).

- ***Modeling Assembly Temperatures:***

The hardnesses and the cure temperatures of the lens-to-substrate, substrate-to-housing, and lens-to-housing can be optimized to reduce the stress on the cells and on the lens. To minimize bulge of the silicone cones due to their high CTE, the receivers-to-lens cure temperature is expected to optimally be ~60ºC, or about 40K over a typical near-room-temperature assembly; the lens and copper CTEs are similar enough that even if the copper is not constrained by the lens the shift in the optical alignment is only ~ 0.15 µm/K, this will produce a ~6µ shift in the alignment, which the cones will flex to handle.

It should be acceptable to have the substrate-to-lid and lens-to-lid temperatures be the same (although if needed the lens-to-lid adhesive could be injected later). The optimum if hard adhesives are used is expected to be near the maximum operating temperature, or ~~80ºC, so that the higher CTE of the lid ensures compressive force on the lens tile.

- ***Temperature Cycling***

With the higher-CTE solder sandwiched between the stiff low-CTE cell and the stiff mid-CTE copper, the solder will bear the brunt of stress from temperature cycling, but optimizing the assembly temperatures should keep the stress minimal. The solder should be in equilibrium at roughly 50ºC warmer than a cold winter night, or ~20ºC in most climates, and a large area would rebuild a full 35 MPa of stress on a hot summer day. With the copper's CTE of ~17 ppm/K and the cell's CTE of ~6 ppm/K, most of the force goes into the cell, but the cell is ~3X thicker than the solder and thus should have at most ~10 MPa of tensile stress at the highest temperatures, which is far below the tensile strength of 135-150 MPa. The actual stress should be even lower because the cell is entirely in the shear lag zone of the solder, and should thus only see a few MPa. Thus, the solder should not put significant tension on the cells (and the copper constrains it if it tries to).

Once the solder relaxes it will see +/-60K temperate swings in the field, for a 900-ppm change or 45 MPa for a large sheet and very roughly 10 MPa for a cell-sized region (being entirely in the shear lag zone). This is well below its yield point, so no further bulk relaxation is expected.

Repetitive stress effects are especially hard to predict because stress can accumulate at crack tips. However, in accelerated aging testing of SAC solder joints, solder between low-CTE resistors bigger than cells, survives ~3000 rapid thermal shock cycles between -40ºC and 125ºC (Nurmi,

2005). That is 15 times more cycles than the 200 cycles from -40ºC to 90ºC said to be sufficient for ten years in flat panels (Kurtz, et al., 2013), so no problems are expected even in 50 years.

6.3.4 Thermo-Optical Model

The optical model ray tracing was used in understand where the focus can land when the trough is miss-tracked. Surprisingly, when the lens focus misses the cones it is relatively diffuse (at most a few hundred X) when it reaches the receiver. Analysis showed that no matter where the focus lands everything is sufficiently thermally conductive enough to conduct the heat to the heat rejection without overheating.

6.3.5 Electrical Model

An electrical model in Excel showed that the resistances are small enough that a COMSOL model is unnecessary. However, an electrical model will be useful when a custom cell run is affordable.

6.3.6 Heat-Rejection Model

The heat-rejection model was used to optimizing the fin size and spacing for low cost. This showed that cooling comparable to standard Fresnel-lens CPV can be achieved with 62% less aluminum.

6.4 Final Design and Conclusion

This chapter has covered some of the extensive analytical work that has been done to validate the proposed TLC design, and has shown that, while there are still challenges to resolve, the basic TLC design is ready to prototype.

Except for the except for the optical calculations, the description of the spreadsheet tabs in the analytical model is kept high-level because it is based on straightforward calculations (e.g., steady-state thermal conductivity through constant thickness materials, mass of material times cost per mass, etc.). The complexity and utility of the analytical model is much more in how the tabs of the spreadsheet interact as a whole than in any given equation or tab. The analytical model spreadsheet itself is provided with the thesis, and is being made available for readers of this thesis.

The analytical model predicts that if the optics can be fabricated within tolerances, a module with a RioGlass mirror, Azur 3C44 microcells, Aohong ultra-low-iron lens tile glass, optical PDMS, an insulated metal substrate with the cells directly on 0.5 mm-thick copper, two 25-micron gold bond wires per cell, an aluminum-finned heatsink, and affordable adhesives, should have a module efficiency of roughly 39% under standard test conditions. For a 1.7-meter-long, 32-receiver module, the open circuit voltage would be roughly 87 V, the short-circuit current would be roughly 12.3 A, and the maximum power output (at 1000 W per square meter DNI) would be over 1 kW.

Parameter	Value
Module aperture area	2.69 m^2
Module length	1720 mm
Mirror width	1649 mm
Heatsink width	120-150 mm
Lens tile width	43 mm
Mirror mass	30 kg
Sealed module mass	12 kg
Receivers per module	32
Microcells per receiver	96
Cell size	0.648 x 0.9 mm^2
Cell type	Azur 3C44
V_{OC}	86.7
I_{SC}	12.3
Watts at 1000 W/m^2 DNI STC	1060
Concentration (with high refractive index silicone)	1500 X

Table 6-1: Module parameters if module is fabricated as designed

While a long-term goal is to fabricate a full module to production specifications, there were many unknowns in fabricating the module optics and assembling a whole module, so the immediate goal was to use readily available resources and processes to fabricate a two-receiver mini-module, and test it on sun to prove out the concept and to reveal any hidden showstopper flaws in the design. The prototyping focused on the dense array of cones that funnel light to the parallel-cell receiver. As the most-novel design feature, this had the highest uncertainty, and it was not known whether a usable cone array could be made with equipment affordable for a limited-budget prototype. This prototyping is detailed in the next chapter.

CHAPTER 7

7. Prototyping and Experimental Validation

This chapter covers experimental work done toward validating TLC; it details the prototyping of individual parts and processes, including the building of test equipment used, and concludes with details of building a physical prototype of a TLC mini-module and putting it on sun, and analyzing the resulting optical efficiency.

To leverage funds from Terra Firma Innovations, three Canadian/Quebec funding programs were applied to: Programme Innovation, IRAP, and MITACS, and all were accepted. With modest funding in place, prototyping has proceeded on a tight budget.

7.1 Details of TLC Prototype/test at the University of Sherbrooke

The prototyping / test steps at the university are presented in chronological order. In many cases additional tests are planned for when time permits.

7.1.1 PDMS optical silicone RI ~ 1.415:

Dow Sylgard 184 was selected because it is widely used as optical coupling in CPV and is readily available at the university. DOWSIL 1200 OS worked as an adhesion promoter on glass.

Polyethylene showed no inhibition of silicone cure, so it is used for molds for silicone.

7.1.2 Dispersing blocks:

Dispersing blocks were cast to allow seeing the path of light emerging from a lens or cone. A tiny amount of fine abrasive was dispersed in Sylgard. 32 mg of 3 µm alumina abrasive was mixed into 5.5 g of liquid silicone, and one drop (~100 mg) was mixed into 22 g of liquid silicone, which was cast into a block. The roughly 1 part per 40,000 (by weight) of 3 µm white (alumina abrasive) particles clearly shows the path of a laser pointer's beam:

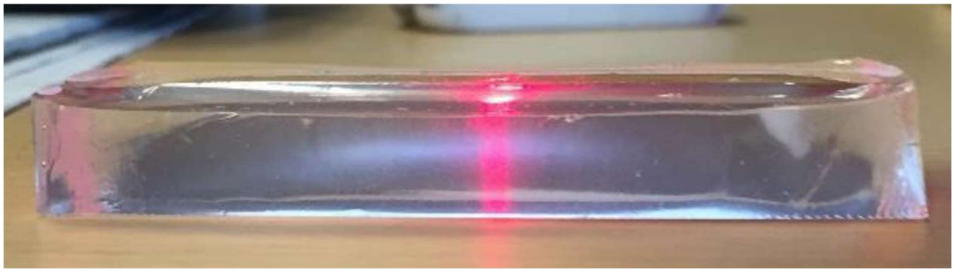

Figure 7.1 Dispersing block shows the path of light

7.1.3 Heatsink stock:

While in production an optimized heatsink will be needed, a custom extrusion would be far too expensive for early prototyping. A quick search found heatsink stock with very similar fin spacing and fin height, and although the fins are roughly twice as thick as needed, the effect of that can be accounted for. The heatsink base is also around four times as thick as is needed, but that can be mechanically thinned. An 8-foot length of heatsink stock was acquired (enough for roughly a dozen mini-modules) and two module-width pieces were cut from the length of stock. Testing of the malleability of the aluminum showed that the aluminum would have to be de-tempered before it could have flanges folded from it, so for the first prototype heatsink lid flanges will be brazed on.

7.1.4 Pre-existing molds for shrinkage and surface roughness testing:

Since polyethylene is relatively easy to mold and does not inhibit the cure of optical silicone, it is the plastic used for the first cone molds. Polyethylene shrinks upon solidification and with further cooling, so the mold for the cone molds is pre-stretched to account for the shrinkage. While the shrinkage of polyethylene can be looked up, different types of polyethylene (especially different molecular weights) shrink by different amounts, with a range from 1.5% to 3.5%. While the cone molds will flex and stretch to accommodate shrinkage mismatch along the lens tile, there is no width-wise accommodation and being off by 1% in the shrinkage would shift the cones at the end of the row by ~170 μm, which is significant on the scale of a microcell.

To more precisely determine the shrinkage, a sample of the polyethylene was injected into a pre-existing mold using roughly the temperature and pressure that we expect to use for molding the cone molds. The polyethylene part produced was extracted from the mold and then partially

reinserted firmly against one side of the mold, and the gap from the other side of the mold was carefully measured. The measured shrinkage was 2.54%, so the cone array master and the cone mold mold designs were oversized by this percentage.

When forced into a mold the injected polyethylene starts to solidify as it flows through the mold, increasing its viscosity enough to produce surface distortions. The tests with a pre-existing mold confirmed that the polyethylene scorches if overheated initially to compensate for this, and that preheating the mold is a better method. Since the mold's temperature has to be kept near polyethylene's melting point, heaters and thermocouples have been included in the mold design.

Some of the molding tests were also used to prepare parts of the surface roughness test. Small rectangles were cut from a standard surface roughness sample plaque in these rectangles were placed in the pre-existing mold to see how well the surface roughness transferred to the polyethylene. These were subsequently used in the TIR-versus-smoothness testing.

7.1.5 QDEC oversized CPV microcells:

Custom tandem cells would cost ~$40,000 for a cell run, and while this would provide many tens of thousands of cells, we only need a few hundred cells early prototyping. We therefore acquired ~200 small QDEC cells made by a company associated with the university. While the QDEC cells are considerably bigger than production TLC cells would be, this will be useful in early prototyping because it provides an oversized target for cell-to-cone alignment (e.g., it could cover for shrinkage variation in the polyethylene cone mold).

A paper on the potential for thermal runaway when one cell set apparel cells is under illuminated (M. Steiner, G. Siefer, and A. W. Bett, 2013) was worrisome enough that as soon as the QDEC cells were acquired, potential for this thermal runaway mode was tested. A cell was placed on copper tape adhered to a heated stage, and leakage current was measured as a function of voltage and temperature. At all normal operational temperatures (up to 100°C) a QDEC cell's open circuit voltage (V_{OC}) drove far less current than the 217 mA/mm^2 safe limit for the production cell size, and considerably less than the 70 mA/mm^2 safe limit for the larger prototyping cell size.

Even extrapolating to the temperature that would be reached if the prototype were left on sun and the grid was down and all light was converted to heat (~ 150°C), both cell sizes are in a safe zone

(Norman, et al., 2020). However, this may not apply to the expendable silicon cells used for early prototypes, so with the silicon CPV cells just having been received, the test will run with silicon cells. The QDEC cells will also be retested out to 150° remove the extrapolation.

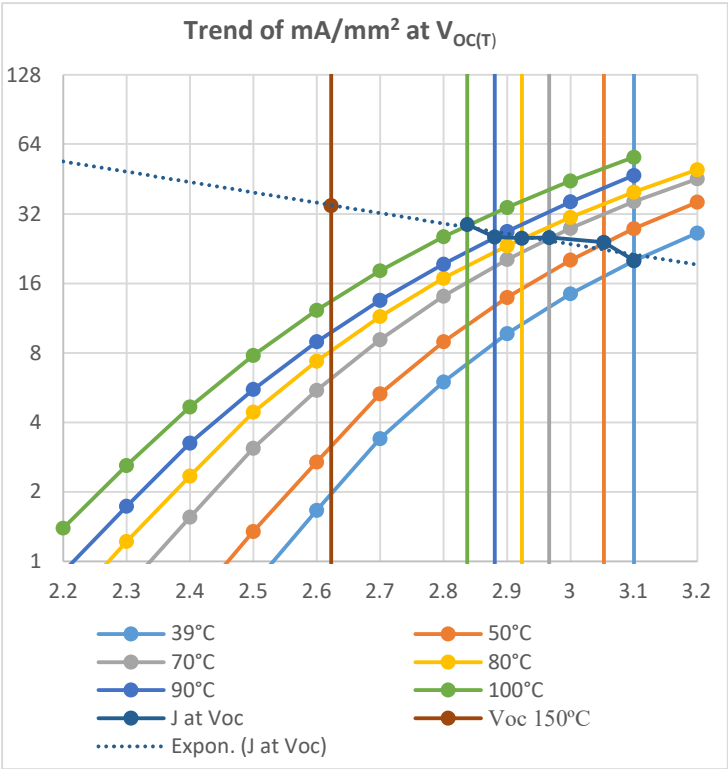

Figure 7.2: Confirming parallel microcells' resistance to thermal runaway (published in (Norman, et al., 2020))

7.1.6 Ultra-low-iron glass:

The prototype is designed to use 19 mm ultra-low iron glass, which is the thickest standard glass used in tall buildings. This glass is thus readily available by the container load at a low cost per kilogram, and we were able to obtain two 305 mm x 305 mm samples at no cost. As soon as the samples were received the thicknesses were measured and found to vary by up to 50 µm within a sample (the faces are not quite parallel), and by up to 100 µm between the two samples. The lens mold design was adjusted to accommodate the thickest point on the thicker sample, and when a thinner glass block is used the difference is made up with optical silicone. Lens-tile-sized glass blocks were cut using a tile saw with a water-cooled blade.

7.1.7 Three-wavelength laser:

While a normal laser pointer can be used to excite silicon cells, high-efficiency cells such as the QDEC cells need three junctions excited by different wavelengths. An elegant hack was suggested – a green laser pointer starts with a 'red' laser in the range absorbed by the middle junction and uses this to drive an infrared laser of wavelength suitable for the bottom junction, and this infrared laser is run through a frequency doubling crystal to produce the green light which fortuitously is at a wavelength suitable for the top junction. Thus, by removing any infrared filter one should get a laser with three wavelengths that are a match for the three junctions of a QDEC cell. A quick test was made using a bare QDEC cell on a workbench with the laser held by hand, and it worked.

The green laser used is more powerful than lasers in the eye-safe range, but measurements in the optics lab showed that if the current is kept not far over the lasing threshold the optical power is low enough to be safe (but never look directly into a laser in any case!). A test set up was designed and built that allows tests with limited laser power in the open and enables higher-powered measurements when a protective shroud is closed.

7.1.8 TIR-versus-smoothness tester:

TLC's cones use TIR (total internal reflection) to guide and concentrate the light onto the cells. With perfect surfaces TIR would be perfect when light comes to a surface at any angle shallower than the critical angle, which can be calculated from the index of refraction of the material. But surfaces are not perfect, and even imperfections as small as 1/10 of the wavelength of light start to allow light to leak when the incident angle is within a few degrees of the critical angle.

The closer to the critical angle one can come, the higher concentration one can achieve, so we want to know for a given surface roughness how close to the critical angle we can come before the cones leak a significant amount of light. Surface roughness is very complex, so the best way to determine this is to try reflecting light from a similar surface.

A total internal reflection tester was designed and built. A cylindrical front allows light to enter perpendicular to the front surface (to avoid refraction changing its direction) and yet reach a test surface at any desired angle. To get test surfaces as similar as possible to the cones, polyethylene samples were molded (at two different temperatures) against standard roughness coupons and

silicone test surfaces were then cast against the polyethylene samples. This closely replicates the silicone surface that would be produced from a cone mold that was itself molded against with a master cone array polished to the same roughness as a given standard roughness coupon.

This showed that a surface finish of 4L (lapped to 4 microinches, or 100 nm surface roughness) was very close to as good as a sample cast against a smooth glass sheet, so the target when polishing the master cone array is at least as smooth as the 4L sample surface.

Figure 7.3: One of the three TIR tester setups (by N. Caillou)

7.1.9 Silicon oversized CPV microcells and receiver-sized cells:

The expendable silicon CPV cells have been received. Just like the precious QDEC cells, these silicon cells are supposed to have solderable backs and wire-bondable fronts. We have ~1500 cells with part of a bus bar for an easy wire bond target, and roughly 15,000 cells with just the gridlines – although we refer to these as mechanical cells, we expected to be able to land a wirebond on each grid line and use them as good cells.

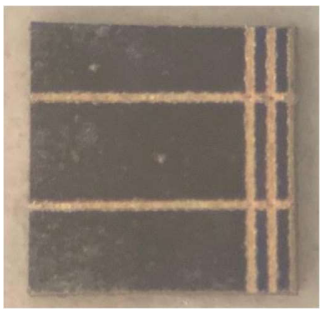

Figure 7.4: Silicon CPV cell with bus bar on one side

Once the 3IT reopens, we will repeat the thermal runaway testing that was done on the QDEC cells to find out if there are concentration of voltage limits beyond which we cannot test with silicon cells, and we will also flash test the QDEC cells to confirm their efficiency.

We also have roughly 60 receiver-sized silicon CPV cells. A single such cell can be used to measure the light through a receiver-sized region of lenses, cones, or lenses and cones. We can also use 32 such cells to emulate a full TLC module for profiling the trough's focus, although this is not planned for near-term testing.

7.1.10 Thermally conductive prepreg and copper laminate:

While a standard IMS process produces a usable receiver, it puts a layer of relatively low thermal conductivity (compared to copper) dielectric isolation between the thick copper backplate and the backs of the cells. It also does not take advantage of the copper backplate as an electrical conductor and instead carves the top copper into an interdigitated pattern of positive (cell back) and negative (cell front) conductors. Some IMS manufacturers have a process called "selective dielectric removal" (SDR) that would allow the cell backs to be directly on the copper backplate, but this is an expensive process that in production would add 20% to the total TLC module cost.

However, having all cells in parallel in a regular array opens up the possibility of lower-cost ways to remove dielectric and even to partially pattern the top copper. The currently-most-promising of these is to start with a sheet of top copper on prepreg (dielectric isolation) on thick copper and to saw straight grooves across the substrate to selectively remove strips of top copper and dielectric where the cell rows will go, and then to use photoimagable soldermask to define the cell sites. While this disconnects the individual rows, these can be reconnected with copper strips.

The materials for testing this process have arrived, along with thermally conductive prepreg that will be used to line the inside of the heatsink lids for the mini-modules. The sawn substrates are not in the critical path and experiments will be when time is available (a description is under "Improved materials and equipment for retesting"), the thermally conductive prepreg will soon be used to line a heatsink lid, and the thermal resistance of this combination will be measured.

7.1.11 Ordinary-IMS receiver substrates:

The standard IMS receiver substrate has been designed and fabricated. The final substrates added a few features to the substrate drawing below – it has all the features of the standard-IMS receiver (as shown below), but also has a 17th set of cell sites added at the top of 5 of the rows, which will allow silicon cells for alignment even when QDEC cells are used for the 16 "real" cell sites in each row, and it has room for two diodes for the high current of silicon cells. Dummy cell sites with varying sizes were added in other blank areas to allow determining the optimum pad size for cell-centering from the surface tension of the solder during reflow.

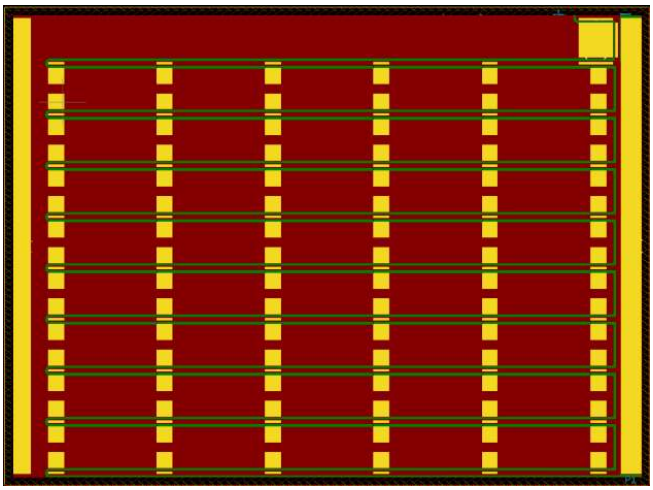

Figure 7.5: Standard IMS substrate for 6 x 16 array of cells

The expendable silicon cells were already ready, allowing cleaning, solder paste dispense, and reflow procedures to be worked out. 1000 cells were then hand-transferred to 3D-printed cell trays, from which six substrates were populated using automated pick-and-place.

7.1.12 Alignment test stand for testing with laser

The alignment test stand comprises a laser holder and a base. The laser holder has lenses that focus the laser's beam to a few-hundred-micron width. The final lens should then collimate the beam to maintain the tight focus at any distance, but this is not working well so the height will have to be carefully controlled. An arm that allows angling the laser to the prototype pivots around a point at

the same level as the lenses. The base will hold the prototype in the laser beam's path and has rotational and translation stages that allow aligning the prototype accurately to the beam.

This is designed to test acceptance angles and misalignment tolerance of lenses and cones, and we will initially use it to measure the lens acceptance angles using the lenses-only glass block.

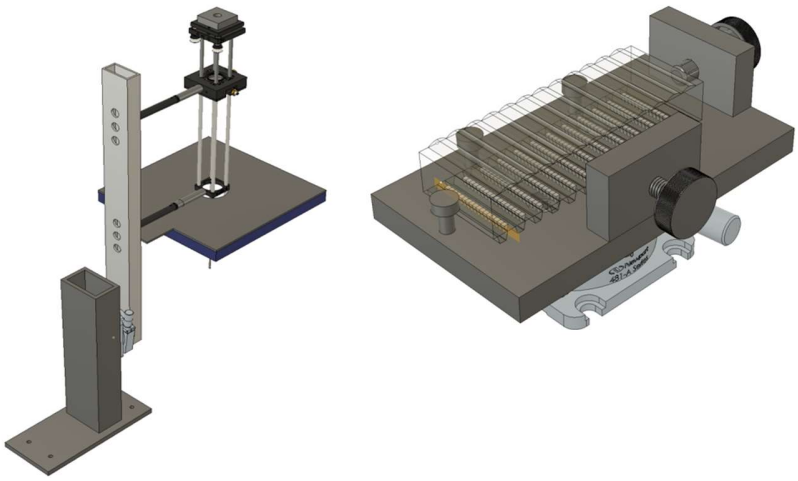

Figure 7.6: Alignment test stand (left), and prototype-holder base(right)
(This has been worked on by several interns; the present PhD candidate participated in the design.)

7.1.13 Temporary Mirror:

Figure 7.7: Temporary parabolic mirror to test mounting and tracking

Due to Covid delays in the ordering process and in shipping, topped off by a longshoremen's strike in Montreal, the crate of RP-3 mirrors (primary concentrator plus ample spares in case of breakage), kept getting pushed out. Eventually a temporary mirror was hacked from thin glass on bendable plywood to allow verifying the mounting and tracking means. This was replaced when the real mirrors finally arrived.

7.1.14 Brazed-flange heatsink lid:

With the parabolic mirror we can test at high enough concentration that a sample needs to be well cooled, and the TLC module's heatsink lid is designed specifically for such cooling.

The first heatsinks use pieces cut from an off-the-shelf extrusion that has fins thicker than optimal. The thicker fins slightly constrict the airflow, and the fins are also slightly shorter, so the width has been increased to achieve the planned cooling capacity.

The long-term plan is to fold the heatsink lid's flanges from the heatsink material itself, but for simplicity we initially removed the excess material and brazed on pre-folded flanges. The first try burned through a flange (providing an expendable heatsink for performance measurements), but by using two torches on lower heat the second try succeeded.

The thermally conductive prepreg that will be used to electrically isolate the module from the lid has arrived. Pressing it on a blank heatsink will be pre-tested, and we will then press prepreg into the flanged lid.

We will then use thermal grease to mount a black-painted bare receiver substrate plus a thermocouple inside the lid and test the flanged lid at the focus of the trough. Weather permitting, we can check the cooling performance in a range of wind and temperature conditions.

Once we are done with a one-sun testing of the glass block with silicone lenses and thermal bridges and large-cell receiver, we can mount the heatsink lid on this set up and repeat that testing at the focus of the hand-steered trough mirror.

7.1.15 Lens and thermal bridge mold:

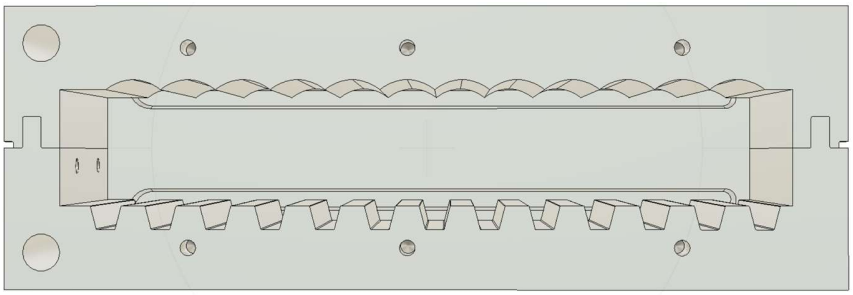

Figure 7.8: Lens and thermal bridge mold

The parts for the lens mold contracted to an EDM shop, and after numerous Covid-related delays were finally received. The optical surface was then polished to a mirror finish (by Etienne Léveillé):

Figure 7.9: Polished surface for lens molding

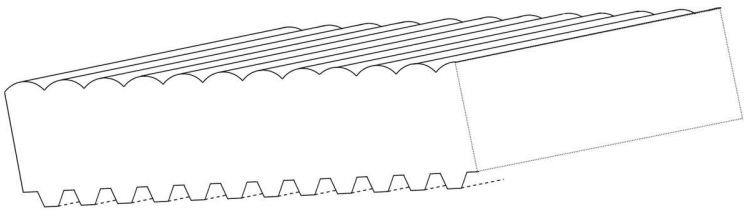

Figure 7.10: Glass block with silicone lenses and thermal bridges

Both lenses and thermal bridges were overmolded on a glass block (but no picture before the cones seems to be available) – see above sketch.

- *Lenses-only block:*

A lenses-only block was placed block on a sheet of thin white paper and the focus was observed on and through the paper. Without cones, chromatic aberration gave a slight tinge of color to the edges of the focus, and as expected this was barely visible.

More measurements can be done with the collimating tube. Under sunlight the lenses should produce a series of short 0.965 mm wide (9.2X) focal lines on the paper at 0º and +/- 24º (matching the trough's rim angle, with a slightly tighter focus in between these angles (minimum of ~0.9 mm). A CCD image sensor (or a cheap camera) affixed to the lens tile might be able to measure the ~50 μ of colored edge expected.

We will then use the receiver-sized silicon cells to measure optical efficiency by measuring I_{SC} of a large-area cell under sunlight, measuring the same cell with the lenses-only block interposed in the light path. We will then repeat this using a mask made from dark paper (or plastic) with a laser-cut slit for the focus of each lens; by using masks with slits of various widths we can profile the distribution of light in the focus and compare it to ray-tracing expectations.

We are likely to end up with some early blocks in which the lenses are good but not great. We will deliberately get such a block dusty and put fingerprints on it to see how easily it is cleanable (and thus how much care we have to take to keep the best blocks and future prototypes clean).

- *Block with thermal bridges:*

With over molded thermal bridges on a third glass block bare substrates were used to test the receiver-to-lens-tile attachment and the receiver-to-receiver interconnection. The block had copper tape applied to the thermal bridge on each end and the thermal bridge in the middle. Mechanical epoxy was then dispensed onto the non-copper thermal bridges and electrically conductive epoxy was dispensed onto the copper-covered thermal bridges (this is done last to reduce the chance of getting conductive epoxy where it is not wanted). Two receiver substrates were then mated to this block. The conductivity between them confirmed the interconnection process.

7.1.16 Cone mold mold:

We received the cone mold parts after months of mostly-Covid delays. The cone plates that form the master cone array have a fairly smooth EDM surface, and our goal is to improve this finished

to roughly 100 nm Ra with hand polishing, but due to the delays we tested its use before the laborious polishing of the cones. We molded a series of cone molds while experimenting with polyethylene temperature and mold temperature until we had a recipe for molding a complete cone mold without scorching the polyethylene. Once the cone plates are polished, we will fine-tune this to get the smoothest cone-mold surfaces.

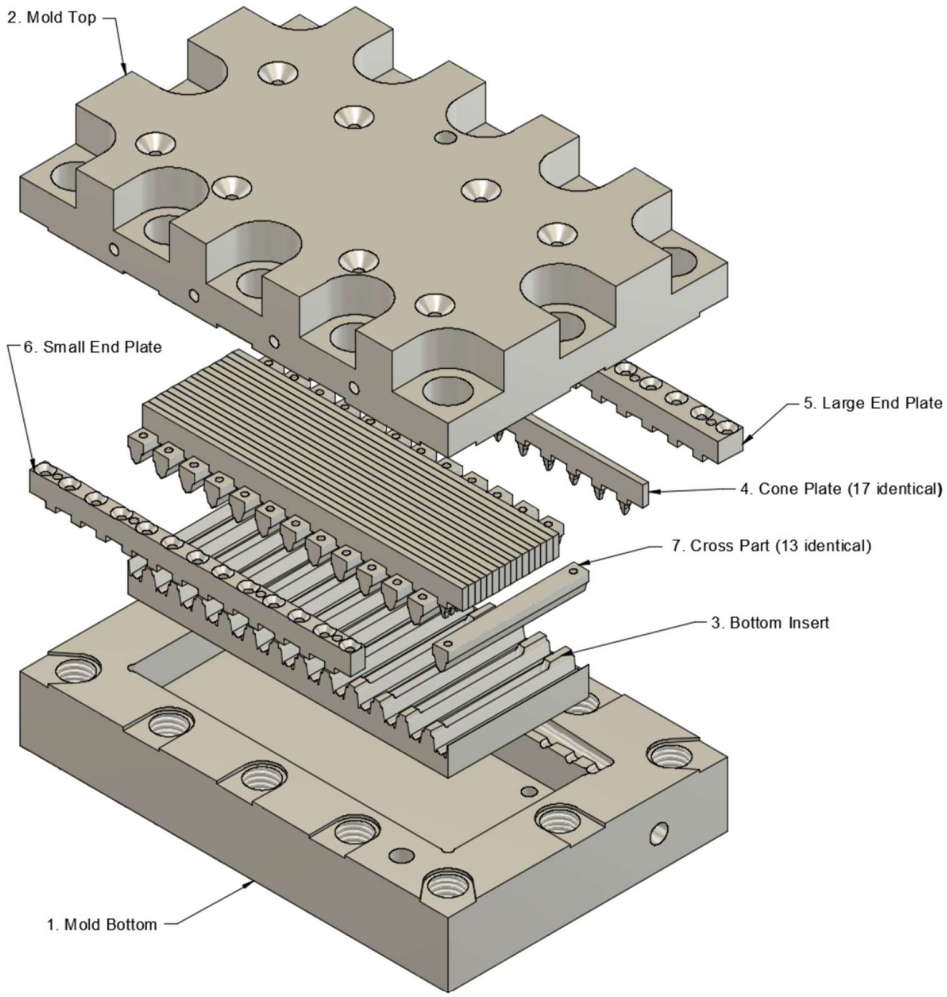

Figure 7.11: Cone mold mold

We used the best cone mold to mold an array of cones on a bare glass block. We tested how well the cones funnel the light by using the large-area-cell photocurrent measuring device. To get smoother surfaces, we Parylene coated the master cone array and reflowed the Parylene.

7.1.17 Self-aligning cone mold:

The purpose of the cone mold mold is to mold a cone-array mold that comprises a series of cone-row molds with slightly flexible linkages between them. This flexible mold not only makes demolding easier, but it automatically aligns to the thermal bridges with no CTE match issues even along the 1720 mm length of a full-module lens tile. While thermal expansion still has to be watched across the ~ 40 mm tile width, this is a 40X reduction in sensitivity. These cone-array molds can themselves be molded either in a thermoplastic polymer such as polyethylene (which does not inhibit the cure of addition-cured silicones), or in a setting polymer such as urethane, which would then be Parylene-C coated to prevent cure inhibition of the silicone. Since the molds are themselves molded, they can be very low cost so a slow-curing silicone can be used for the optics, even in production, without requiring a massive investment in molds.

The first usable self-aligning cone mold is shown below, along with a sketch of how it mates to a lens block:

Figure 7.12: Self-aligning Cone Mold (showing cone cavities) (EL & NC).

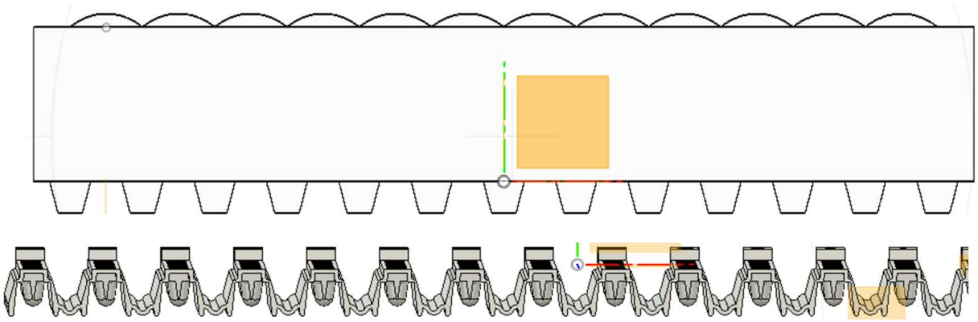

Figure 7.13: Cone Mold Flexes to Mate to Thermal Bridges

7.1.18 First lens tile:

With the useable cone molds now producible (albeit with unpolished cones), the work on the cone molding itself started. The proof-of-concept test went surprisingly well!

Trimming the silicone overflow away with a surgical scalpel allows a view of the narrow parts of the inter-cone gaps, which can be seen (in Figure 7.15) to taper to well below the 25 µm width of the gold wire. After trimming, the very narrow section of gap is distorted; this may be a trimming artifact or may be from the very thin polyethylene cone mold wall in this area being distorted. However even at 25 µm, the 1.2% loss of light through the gaps would be acceptable.

We vacuum degassed silicone and then repeat the test on one of our lens-tile blocks, in parallel with test-polishing the cone plates to try for have a cone-array mold that produces sufficiently high-quality cones. We will then combine these to mold an array of polished-surface cones on the back of the glass block that has both lenses and thermal bridges. This will produce a "lens tile" that comprises the second and third stage optics of a TLC mini module. Since the test setups will be ready, we will immediately be able to test this with the large-area photocurrent measurer to measure the optical efficiency, and to compress it against white paper (on sun and in the laser test jig, and with a dispersing block). Note: polishing proved difficult, so the on-sun testing proceeded with un-polished cones.

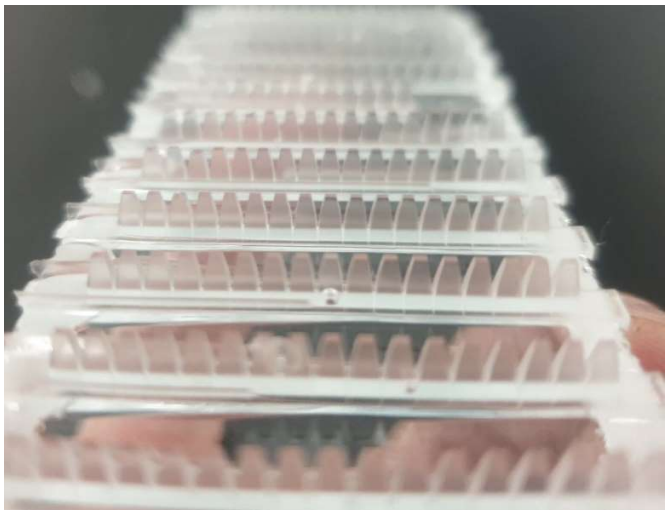

Figure 7.14: Array ~3 mm silicone cones with mold from unpolished master array (NC)

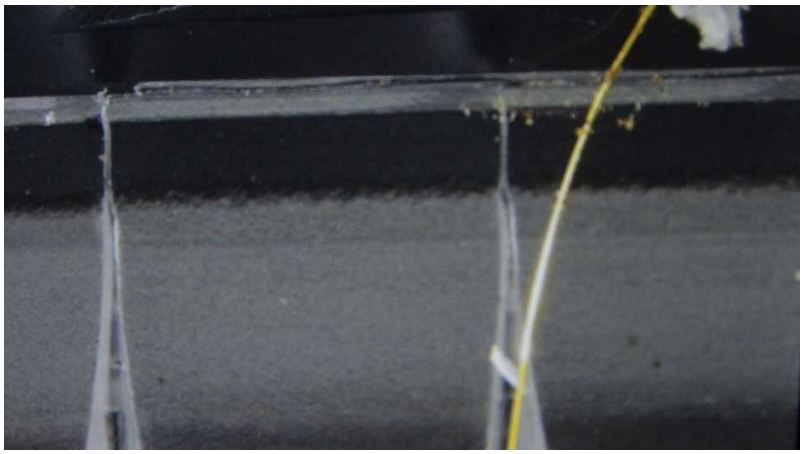

Figure 7.15: Side view of cones with 1-mil (25 µm) gold wire for scale

7.1.19 RP-3 Trough mirror:

After many months of delay, the parabolic mirrors finally arrived. Our immediate need is only one mirror (plus a few spares), but shipping is almost as much as the mirrors, so we ordered ten to have spares and some for future use as well (e.g., profiling the focus of multiple mirrors, and, later, a multimodule prototype). The temporary mirror on the tracker was replaced with a commercial mirror, so we have a 2-axis-tracked RP-3 mirror ready to mask and then use. We can also hand-steer a mirror to briefly flash a sample receiver with up to 50X concentration.

7.1.20 Receiver substrate populated with cells and diode, interconnected:

Once the cleaning, solder paste dispense and reflow processes were worked out, a panel of six receiver substrates was assembled with the expendable silicon cells. However, the subsequent step of wirebonding the cells' gridlines to the pads on the substrate could not be made to work reliably enough to electrically interconnect the receivers, leaving us with nice-looking mechanical receivers which were useful for debugging the module assembly process. We used several arrays to develop the receiver-to-cone alignment process, and a process for holding a newly-placed receiver in place while its adhesives set. We used several more to debug the receiver-to-receiver interconnections through copper strips on the lens tile and found that we had to add insulation to prevent electrical epoxy from flowing where it is not wanted (proper dispensing equipment would avoid this issue);

this left a ridge on the receiver that, with the adhesive used, increased the thermal resistance to the heatsink by six-fold. We then assembled the first opto-mechanical mini-module and found that while the thermal path was not great, it was acceptable for initial testing.

With the season for outdoor testing approaching, we then took a calculated risk and used our irreplaceable QDEC cells to assemble two functional receivers. While we had AZUR wafers on order to cut even higher efficiency cells from, AZUR is the company that years ago had the issue with parallel cells, and we had no assurance that AZUR's new cells will work on our receivers.

While the QDEC cells proved easier to wire-bond to than the silicon cells, we still only got ~ 90% yield on wire-bonds due to using an automated wire-bonder with aluminum wire rather than manually bonding with gold wire (not a production concern –the electronics industry makes 15 trillion wire-bonds per year, many with gold wire). Fortunately, each cell has two wire-bonds, so this left only a few cells on each array unconnected.

7.1.21 Receiver-sized silicon cell test set up:

Along with the QDEC-cell-sized silicon cells we also received some whole-receiver-sized silicon CPV cells. We used one for a large area photocurrent measuring device by mounting the cell in a holder that provides one wire for electrical connectivity to the cell back and another connected to the cell front. Since the short-circuit current (I_{SC}) is almost perfectly proportional to the light on the cell, this makes an accurate way of measuring the total light reaching a receiver-sized area.

Testing using a large-area silicon cell confirmed that the optical efficiency of the first lens tiles was quite low: only about 30% instead of the 95% that a perfectly made lens tile should have.

7.1.22 Artificial Sun:

An "artificial sun" was set up – an inexpensive ultrabright flashlight mounted at a distance from a sample calculated so that the flashlight's optical diameter matches the sun's optical diameter. At night (with no stray light to block), this set up works very well for qualitative tests, and allowed the effects of experiments to be seen, and, in many cases, to be photographed.

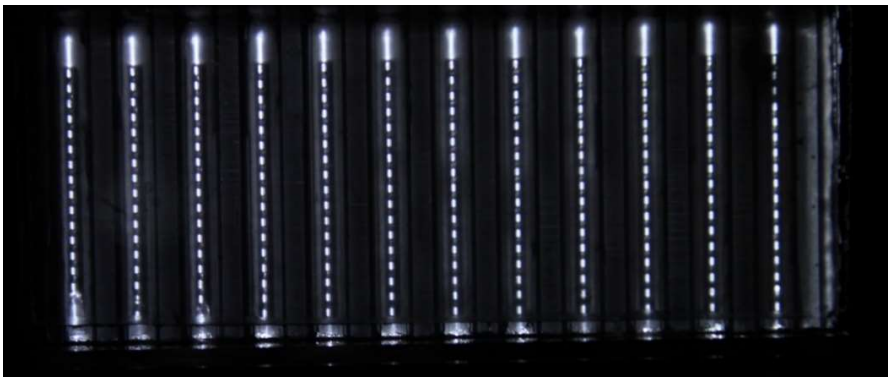

Figure 7.16: The lenses and cones guide the light to microcell-sized focal spots

In addition to the cone roughness, several other significant causes of loss of light were found. One of these is already resolved– while the first lens tile (made before the lens mold was polished) shows significant light leakage across almost the entire bottom of the lens tile, lens tiles made after the lens mold was polished leak very little light through the bottom of the tile. Light leaking between the cones within a row proved minimal, and examination under a microscope showed that the inter-cone gap was less than 25 μm), confirming that making the steel master cone array as separate plates was successful at preventing significant inter-cone gaps.

The three main causes of loss found were bubbles, optical decoupling, and cone roughness. Bubbles proved more impactful than expected because in the lens tiles with complete cone arrays, large bubbles were quite common, affecting more than 20% of the cones. Most bubbles are near cone tips (a mold-filling issue). Cone arrays where the liquid silicone had been more thoroughly vacuum-degassed after filling the mold had fewer bubbles, but the degassing process had foamed away enough liquid silicone that some of the cones lost contact with the back of the glass block, and these cones were missing from the lens tile. In production short sections of cone-mold can be filled with pre-degassed liquid silicone while already under vacuum, for no air to trap in bubbles.

Optical silicone is soft enough that it normally optically couples to a surface it is pressed against, but a coupling agent was needed to fill the pits in the rough cone tips. To avoid permanently gluing the cone to the test apparatus, liquid optical silicone was used without a curing agent; a needle was used to transfer a tiny droplet onto the tip of each test cone. When the lens tile was illuminated with the artificial sun, the coupled cones were much brighter from certain angles.

Testing a lens tile made with a cone mold formed after two cone plates had their cones laboriously hand-polished, the focal spots from polished cones were much brighter than spots from other cones, showing that the main loss of light was from cone surface roughness. Under diffuse illumination these cones appear much darker because their surfaces leak much less light:

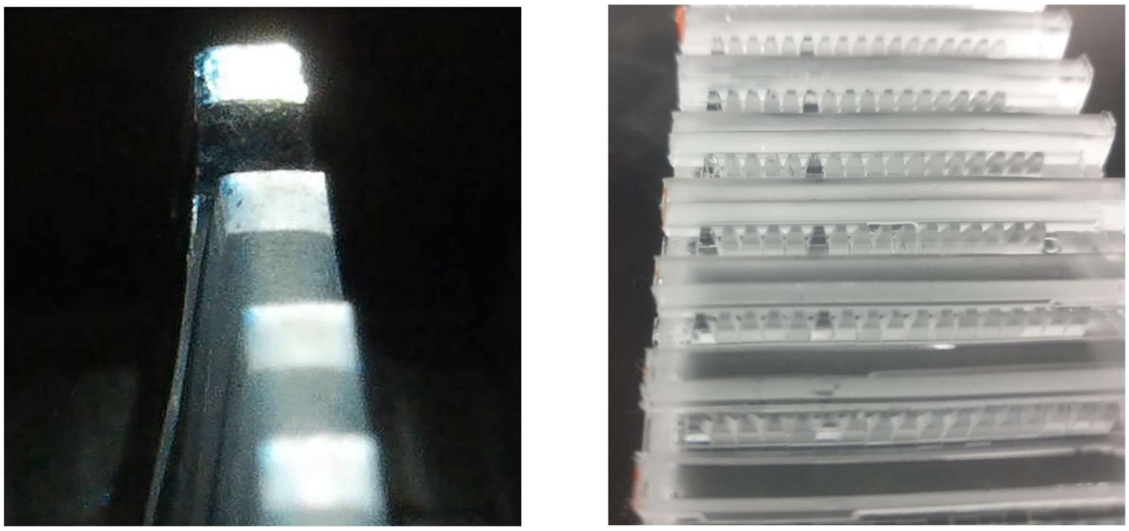

Figure 7.17: (left) Polished cone at top has bright tip and polished cones leak less light (right)

7.1.23 First electro-optical mini-module:

Smooth optical silicone can optically couple to a smooth surface under slight pressure, but the current cones are far too rough for this to work, so a tiny trace of liquid optical silicone is used as a coupling agent. The prototyping cells are oversized, so the bondwire should not interfere with the coupling, but even in a production module the bondwire is so small compared to a cone that the soft silicone will simply mold itself around it, as shown below.

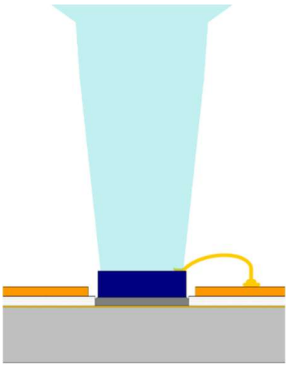

Figure 7.18: Roughly to-scale sketch of cone and bondwire on cell

With only two functional receivers, we used just one functional receiver for the first mini-module so that we would get two shots (for the second receiver we used one of the silicon cell arrays). After several practice runs without adhesives, the 3IT Tresky system proved able to align receivers to a lens tile to roughly 25 µm accuracy. We dispensed adhesives and bonded one functional and one mechanical receiver to a lens tile. The assembly went well, so we had a TLC mini-module ready to test (shown below, with wiring for thermocouples to monitor the temperature).

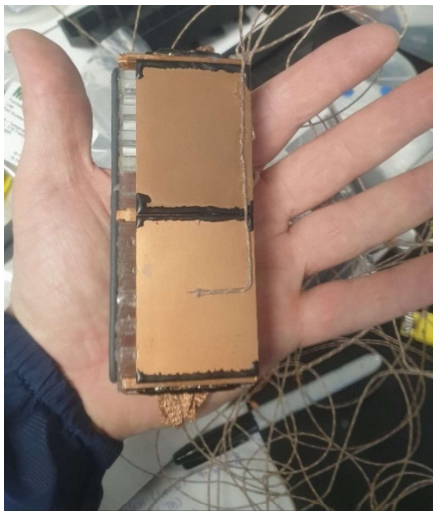

Figure 7.19: Receivers mounted on a lens tile (N. Caillou, 2021)

7.1.24 Instrumentation and Data acquisition system (DAQ):

Critical data are voltage and current, temperatures at various points on the prototype, ambient temperature and wind speed/direction, DNI, time of day, and tracker altitude and azimuth.

- *Voltage and current:*

Voltage and current data are key to measuring the power output of a receiver under a range of conditions. A single TLC receiver is hard to measure accurately because the current is as just high as a whole module while the voltage is only 1/32 as high. Fortunately, even I_{SC} measurements tolerate voltage drops of a few hundred mV, and for other current measurements we can tolerate higher voltage drops (as long as we measure voltage); and voltage can be measured at zero current where resistive voltage drops are zero.

With a current of ~10A at the voltages of interest of 2V to 3V (with tandem cells), we want a load with an equivalent resistance controllable from ~0.2 to ~0.3 ohms. While we could hack this from

a length of 14-gauge copper wire (~ 8 milliohms per meter) and a sliding contact, it will be more convenient to make a programmable load with a small PCB with a programmable power transistor very close to the prototype, with the current measurement made on that PCB. This measurement PCB will need protection from the sun's focus (e.g., a small mirror) and a heatsink to dissipate up to ~30W. The target was to be able to vary the load at 10 Hz to be able to capture an I/V curve in a few seconds. The target accuracy for voltage is 5 mV, and for current it is 25 mA when testing full modules and 2 mA when testing individual cells.

The data acquisition system worked well for temperatures from the various thermocouples, but the programmable load did not work for the current and voltage, so this must be reworked at some point. However, a simple digital multimeter can take the key measurements in the meantime.

- *Temperature (including ambient) and wind speed:*

Temperature data has several uses: watching temperature rise may allow saving a prototype from destruction (the target is to be able to raise an alarm in at most one second and preferably 0.1 seconds); since cell conversion efficiency depends on temperature, tracking the temperature near the cells allows us to estimate this efficiency hit and hence how large *other* losses are (recording at 1 Hz); and temperature influences module life and reliability, which need to be understood. The DAQ has 8 channels for thermocouples (lens front, lens back, substrate under cells, substrate between cells, heatsink base, heatsink tips, ambient or rejected air, unassigned).

Since cell temperature depends upon wind speed and direction as well as ambient temperature, we also need to record wind speed and direction.

- *DNI:*

The input energy to a module is proportional to the DNI (and to the mirror-with exposed), so DNI is critical to calculating efficiency and to understanding how much light is lost and where it is lost. There are several DNI sensors on site, so ideally, we can pick up the signal from one of them, but it may be simpler to just have our own dedicated sensor on the tracker itself. DNI can change fairly quickly, so the target is recording at 1 Hz.

- *Time of day, and tracker altitude and azimuth:*

Time of day and tracker altitude and azimuth are less critical – they have a secondary effect on tracker accuracy and structure distortion. However, if we record them then we can compare

efficiency to how close we are to proper alignment, which will make determining acceptance angles much easier. Time of day is also useful in diagnosing afterward how long something was on sun before some changes were observed.

7.1.25 Mirror and prototype holder mounted on tracker:

The trough mirror is huge compared to the mini module used in the early testing. The mirror has a parabolic curve across its 1582 mm width, and is straight along its 1700 mm length, so a whole mirror focuses a 1582 mm x 1700 mm region of sunlight onto a ~ 40 mm wide by 1700 mm long focus at roughly 40 suns concentration.

A 40 X focus is dangerous – it can blind someone almost instantly, it can rapidly burn skin or set clothing on fire, and it can also destroy prototypes or measuring equipment **and that focus will be invisible unless there is something at the focus (and even when a prototype is present, the focus floats around invisibly until the tracker is on sun in the focus is on the prototype).**

The first mini-module prototype is only ~120 mm long, and we lose roughly a centimeter on each end of the focus from sunlight spreading, and we want another centimeter on each end for mounting tolerance, so initially we only use ~ 160 mm or ~ 10% of the focus. For safety when setting up a test, and to allow short duration testing, the usable 10% of the mirror has a sliding mask that allows us to "flash" the prototype with sunlight for controlled intervals.

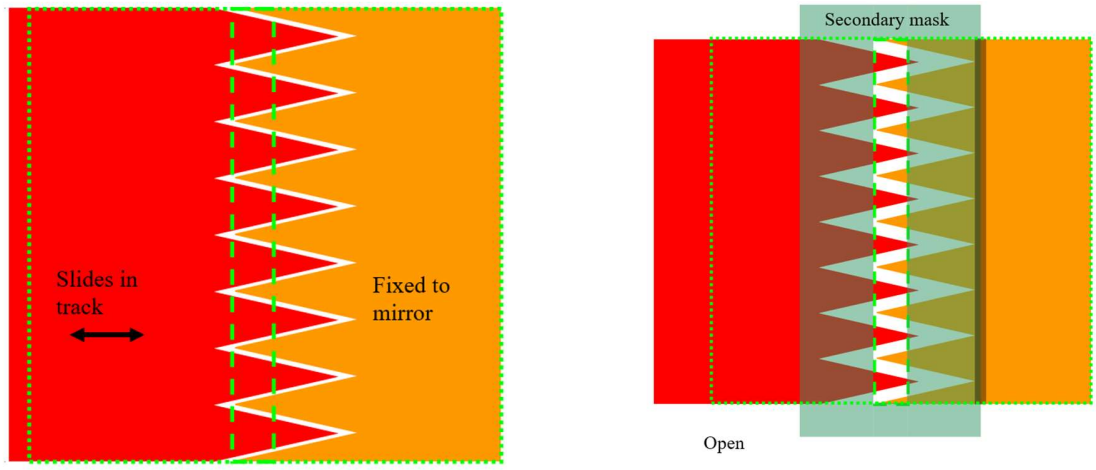

Figure 7.20: Sliding mask for safety and for short-duration tests

The mask's interdigitated triangles allow testing at any concentration ranging from zero up to 2/3 of the mirror's ~ 48 X concentration, while providing relatively even illumination of the lens tile from the full range of angles of the mirror. The mask allows going beyond 2/3, but the lens tile becomes less evenly illuminated. 2/3 concentration, or ~ 32 X on the surface of the lens tile, produces ~ 300 X concentration in the lenses and ~ 850 X concentration in the cones and on the cells). Additional masking could control the angular range from which light reaches the lens tile.

7.1.26 First mini-module (populated lens tile with heatsink)

Figure 7.21: Two-receiver mini-module mounted in heatsink (N. Caillou, 2021)

Mechanically, the bonding of the first electro-optical module to the heatsink went well, but the thermal path was much worse than desired. Insulating epoxy (to prevent electrical epoxy from flowing where it is not wanted) had left ridges on the receiver backs, and these ridges and a thermocouple wire acted as standoffs. The thermal epoxy used bond the lens tile to the receiver heatsink also had 250 μm glass beads, along with roughly half the intrinsic thermal conductivity we had planned for. Together these increased the thermal resistance to the heatsink 6X to 10X beyond the expectations for a production module (proper dispensing equipment and a better thermally conductive adhesive would prevent this issue).

The ends of the heatsink were left open for easy access for wiring to reach test equipment. Under one sun testing, this first mini-module showed an optical efficiency that matched the expectations from optical testing of the lens tile with the large-cell silicon sensor.

7.1.27 Mini-module on sun

As the mini-module was being assembled, the on-sun-testing platform was readied, with a series of problems fixed or bypassed. The tracker was not working in automatic mode, so we installed a manual control box. The data acquisition system worked for temperature, but not for voltage and current, so we ran the voltage and current wires to an external multimeter. We reinstalled the mirror on the tracker and used a bare heatsink to test the focus; the mirror focused 8 centimeters beyond where it was supposed to, so we adapted the mounting system accordingly. The mask to control the focal intensity kept catching the wind and threatening to blow off, so we cut its size in half and attached most of it directly to the mirror for support. Finally, the test platform was ready.

May 24, 2021 was a beautiful sunny day, and everything went well. With the mask closed, we mounted the mini-module at the mirror's focus, and tracked the mirror on sun. At 10:40 AM we opened the mask an inch, exposing the prototype to 5X concentration from the trough mirror (~140X concentration at the cone tips). Our TLC mini-module was on sun!

The temperatures, voltage and current were all within their expected ranges, so we cautiously opened the mask another inch. The cells got a bit warmer than expected, but not dangerously so, so we opened the mask to 3, 4, 5 and then 6 inches (32X concentration from the trough, and about 850 X at the cone tips). We measured the key values of voltage, current and cell temperature at

each mask opening level, and everything functioned as expected. We then deliberately mis-tracked by measured distances in each direction and got rough acceptance angle data.

Having completed all the scheduled tests, we then open the mask beyond its normal limits to get to ~1000 X concentration at the cone tips. However, due to the sub-optimal thermal path from the cells to the heatsink (the increase in thermal resistance discussed under "First mini-module"), we had to stop that test after a few minutes when the cells reached 100°C.

The masked-mirror efficiency calculations include the loss from the mask edges blocking some light. Each triangle of the mask has two edges, one of which blocks light from just the edge of the mask, and the other of which blocks light from the mask edge plus two passes through the mirror and then hitting the back of the mask), so the average light blocked is from the mask plus one pass through the mirror. The edge-loss area is constant until the mask is fully opened, at which point it decreases linearly to half its original value as the mask is over opened and the triangles from only one side of the mask are present in the usable area.

The area of light lost is the thickness of the material times the sine of the angle of the light surface. Sines are approximately linear out to the rim angle of the trough, so the average sine on a parabolic curve is very close to one third of the maximum sine. The thickness of the mask includes the 4 mm corrugated plastic and the two layers of black gorilla tape, while the 4 mm thickness of the glass is divided by the refractive index of 1.52 (since this decreases the angle of the light inside the glass). The corrugated plastic was fairly tight against the glass; if a gap were present, it should also be included in the calculations. The slant of the edge along a triangle also introduces two cosine effects, but one increases the area, and the other dilutes the effect, so they cancel each other.

Results:

Optical efficiency:

The optical efficiency was in line with expectations based on early testing of lens tiles on large-area silicon cells. Estimated optical efficiency rose from 41.5% when the mask was almost completely closed to 43.1% when in the mask was fully opened, and 44.8% when the mask was fully over-opened. The light is uneven when the mask is over-opened, and it is not recorded whether the receiver under test was on the brighter or the dim side of the focus, so the last data point has higher uncertainty.

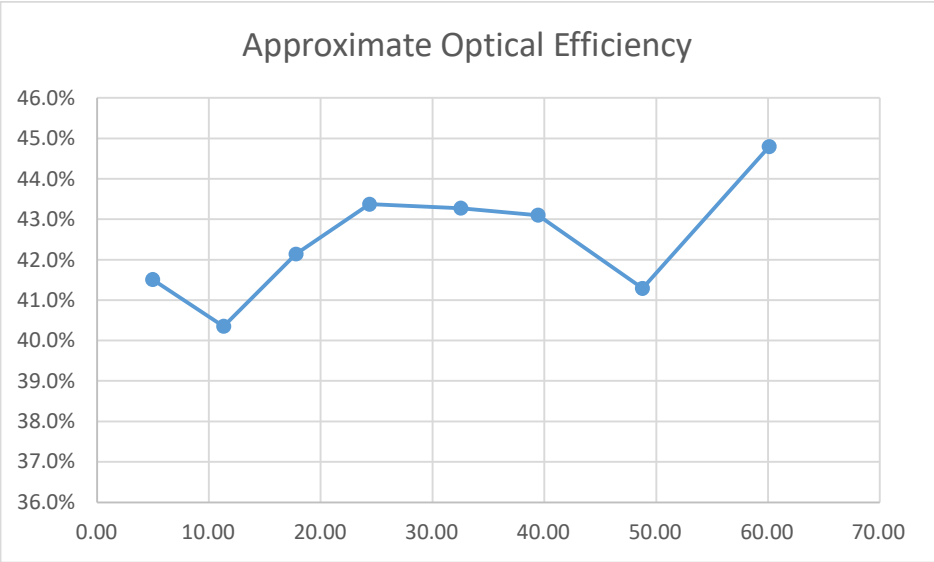

Figure 7.22: Optical Efficiency vs Optical Energy (Watts)

The horizontal axis is in Watts of optical energy on the usable aperture) because it is the DNI in Watts/m^2 times the mask opening area (that illuminates the prototype) in square meters.

While the optical efficiency is significantly lower than the target for a product, it is not unreasonable for the first prototype of a new CPV architecture. Work is now in progress to track down where the light is lost, with the cone surface roughness being the primary prime suspect.

Acceptance-angles:

Acceptance-angle data was also collected by tracking the tracker to maximize the photocurrent, and then miss-tracking the tracker while measuring the photocurrent relative to the maximized photocurrent. A tape measure was used to measure how far the reference point on the tracker moved relative to the fixed scaffolding, and the distance was then converted to milliradians based on the distance from the center of the tracker to the reference point.

The acceptance angle loss mechanism for the altitude tracking is well understood (the edge of the light simply misses the receiver entirely beyond the acceptance angle, so it was merely sanity checked with three data points. This confirmed that roughly 85% as much light was still received even when miss-tracked by 0.44° (very large compared to the 0.1° typical tracking error of a modern 2-axis tracker).

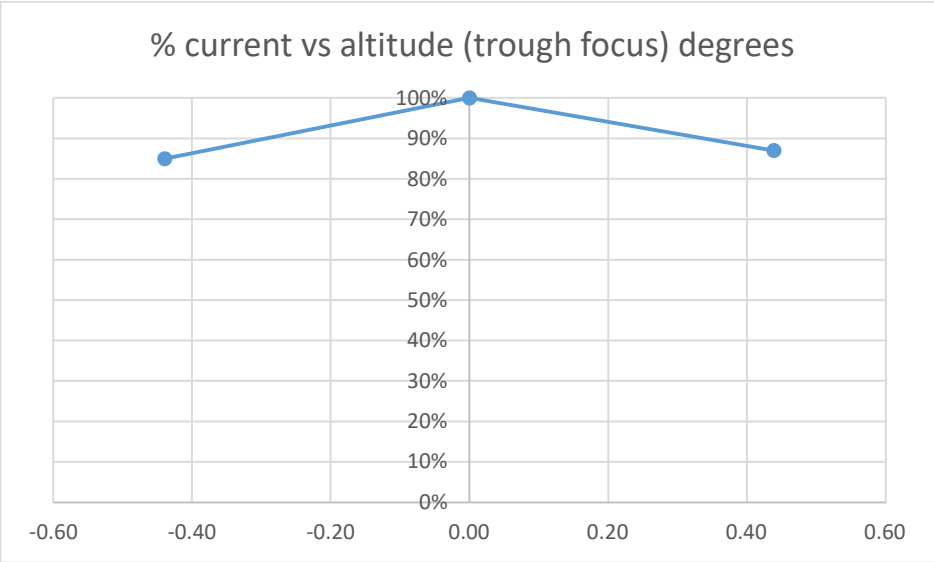

Figure 7.23: Altitude Acceptance versus miss-tracking (degrees)

The azimuth acceptance angle was critical to measure because the cone tops have been broadened in the azimuth direction (after the acceptance angles were last calculated from ray tracing), so more data was gathered on this axis. 99% as much light was received out to a miss-tracking of 0.44°, and roughly 90% was received out to a massive 1° of miss-tracking (due to the cone top broadening). Extra acceptance angle on this axis is welcome because it relaxes the cone-to-lens alignment requirements.

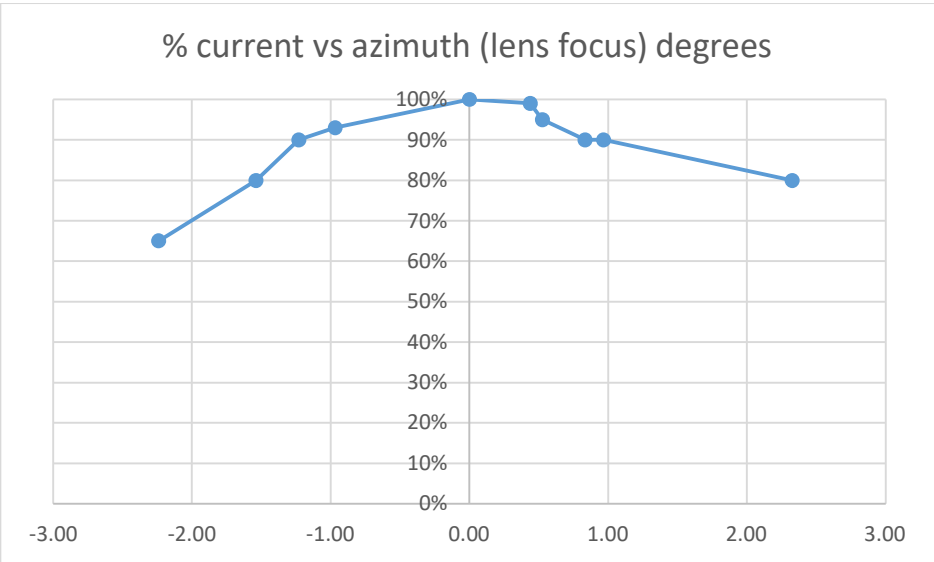

Figure 7.24: Altitude Acceptance versus miss-tracking (degrees)

7.1.28 Well-polished cone-plate

To be able to confirm that the cone surface roughness is of being the main source of lost light, one cone plate was laboriously hand polished with better polishing equipment (using the oscillatory tool with polishing sticks designed specifically for such polishing and being careful to avoid contact between the strip holder at the cone tips). A cone mold was then made with the polished plate next to unpolished plates in the master array.

This cone mold was initially used to mold test cones on a glass slide, and it will subsequently be used to mold cones on a lens tile. Initial images of the cones on a glass guide slide were taken from the back of the glass slide; the focal plane of this microscope was too narrow to see the whole cone, and the plane was just above the tip of the cone. (Pictures from Etienne Léveillé.)

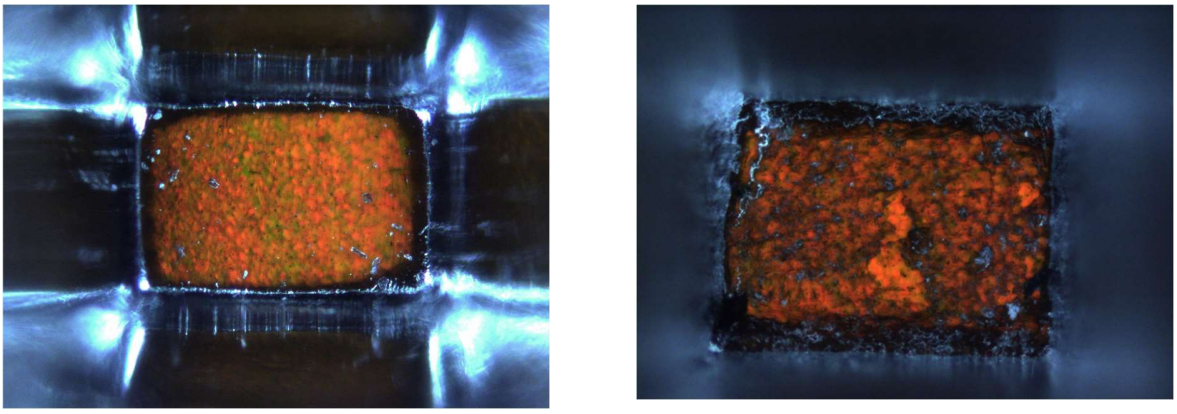

Figure 7.25: Polished cone (left) and unpolished cone (right) (from Etienne Léveillé)

A deep-plane stereoscope was also used to inspect the cones and the whole sides of the cones are similar to what we can see on the edges surrounding the bottom. All the sides of the polished cones also turn Kapton-colored as the cones tips optically couple to the Kapton film that the cone tips are pressed against. In contrast, the rough cones walls stay hazy whitish even when coupled!

7.1.29 Camera Image Sensor:

To visualize where light through a lens tile was going, a camera was acquired and disassembled to extract the image sensor from it. Initial experiments with uncollimated light on the glass slide show that we can see where the light strikes the sensor, and by controlling the intensity of the light

so that we do not over saturate the center pixels, we will be able to use image analysis software to extract data on how much the lights strikes each pixel of the sensor (the “ghost” image to the left of the polished cone is a multiply reflected image of the bottom of the cone, which proper alignment should eliminate).

Figure 7.26: Uncollimated light through polished (left) and unpolished (right) cones (Pictures from Etienne Léveillé)

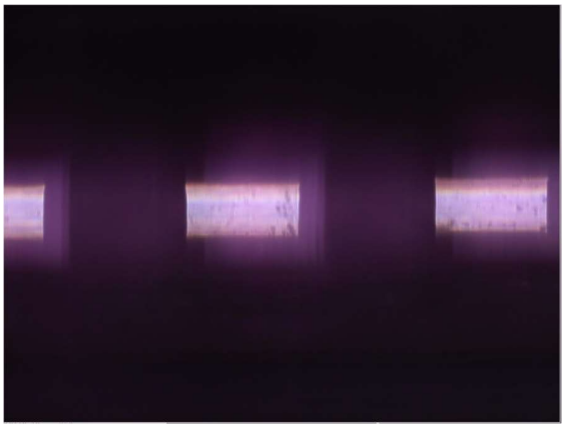

Figure 7.27: Sensor fits multiple cones for comparison under same conditions (from N. Gariepy video)

Two cones will fit simultaneously on the sensor, so we can ensure that the same camera settings are used for comparison images by fitting the two images in the same frame. This will allow us to do comparative test between polished and unpolished cones to measure the effect of polishing and coupled and uncoupled cones to measure the effect of improved optical coupling by using liquid optical silicone as a coupling agent.

The camera sensor also has video output, so we will be able to capture the effect of changing the alignment of the lens tile to the light source, and thus get acceptance angle data for the lens tile itself. We will also be able to increase the compression of the cones and watch the coupling improve, and thereby determine how much compression is needed. Having determined the angular and compression ranges of interest, we can then capture still images at those key points (or note the corresponding video frames of interest) for more detailed image analysis.

7.1.30 Laser Scan on Silicon-Cell Receiver

To test the effect of cone polishing on optical efficiency, receiver substrates were assembled with small arrays of individually-wired silicon PV cells. The outputs of individual cells were recorded while a laser pointer was scanned across and along the array. The scan of the edge of a bare cell was used to profile the laser beam itself; since the cell has sharp edges, the transition from zero to maximum photocurrent is from the beam profile. This confirms that the laser beam was fairly narrow, with ~50% of light within 100 um of the center of the beam, and ~85% within 200 um.

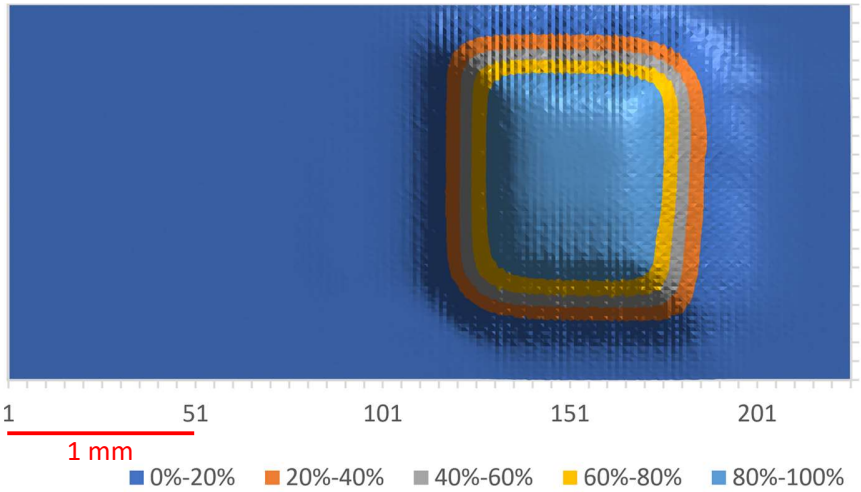

Figure 7.28 Scan of a bare cell to profile the laser itself

The populated cell sites were chosen so that some cells were under unpolished cones on a lens tile, some were under cones from Parylene-reflow-polished master cones, and some were under cones from mechanically-polished master cones. The scans were done first on a bare receiver, and then repeated with a lens tile mounted on the same receiver. Since a lens tile should guide the light to

the cells, dividing the photocurrent at each point with the tile present by the maximum bare-cell photocurrent produces a map of the optical efficiency of the lens tile, with a laser-beam-width taper zone at the edges of the lens and the edges of the cell.

The first map was with an older lens tile (from the same era as the lens tile that went on sun in the first TLC mini-module), with unpolished cones, and the laser beam normal to the lens tile. Figure 7.29 shows a lens-to-lens length (8.85 mm) by cone-to-cone width (2.09 mm) region of the scan:

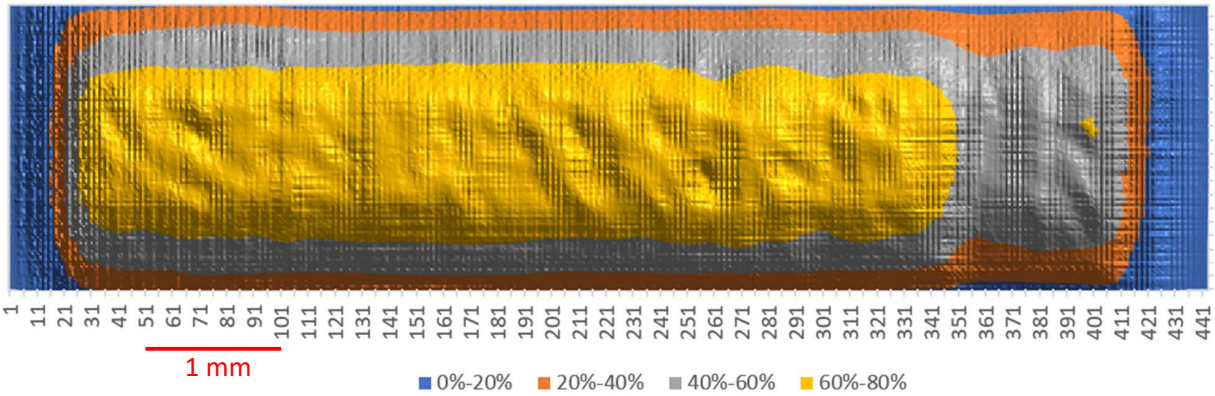

Figure 7.29 Optical efficiency map through a lens tile with an unpolished cone

With this older lens tile the optical efficiency peaked at 76% near the center, and dropped to 66% at cone tip edges (grazing a cone wall near the cell), and, even after accounting for the laser beam width by adding the current of an adjacent cell, dropped to only 34% near the cone-to-cone divide.

Several cone-polishing methods were tried on cones of the master cone array from which cone-molds are molded. Of these, reflowing a Parylene coating was the simplest and was fairly successful even on the quite-rough steel cone plate, while carefully hand-polishing the tiny cone surfaces produced the best overall surface but was the most difficult and distorted the cone shape noticeably. If the initial steel surface were smoother (without the occasional ~ 10 um deep pits), CNC-polishing should produce a true mirror finish with acceptably low distortion, and a thin Parylene coating should also reflow to a very smooth final surface.

Maps of lens tiles made with polished cones show significant improvement – a cone cast from a reflowed Parylene master cone (figure 7.30) had a broad peak reaching 91%, and a cone cast from a hand-polished master cone (figure 7.31) had an even-broader peak reaching 89%:

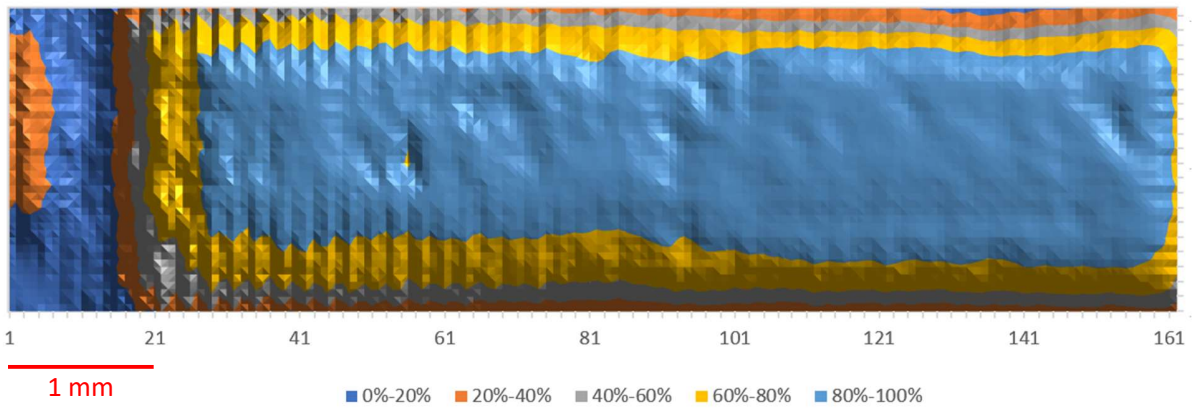

Figure 7.30 Optical efficiency map with a Parylene-polished cone

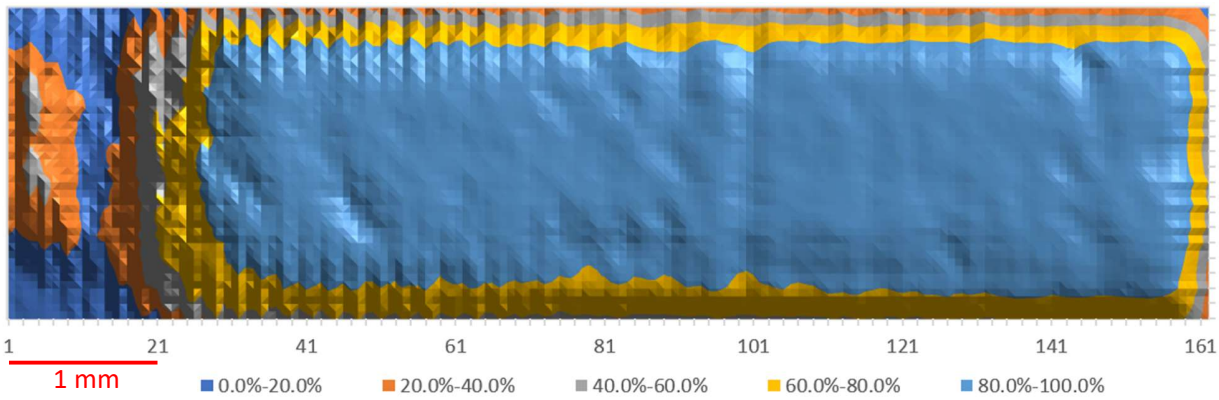

Figure 7.31 Optical efficiency map with a hand-polished cone

7.1.30.1 Two-cell Scan

The odd peak seen on the left side of each lens-tile scan was traced to a 1.2-degree error in the laser beam angle. This is larger than the no-loss acceptance angle in that direction, so it is impressive how tolerant the optics were to such an error. When the beam angle is corrected, as can be seen in Figure 7.32 the odd peaks are gone.

The Parylene- and mechanically- polished cones were scanned together; they are adjacent cones, so by adding their photocurrents light otherwise lost due beam width at the cone-to-cone kerf is captured from next cell, isolating the loss from the cone-to-cone kerf (although still measured with beam-width resolution). This isolated cone kerf loss can be seen in center of Figure 7.32.

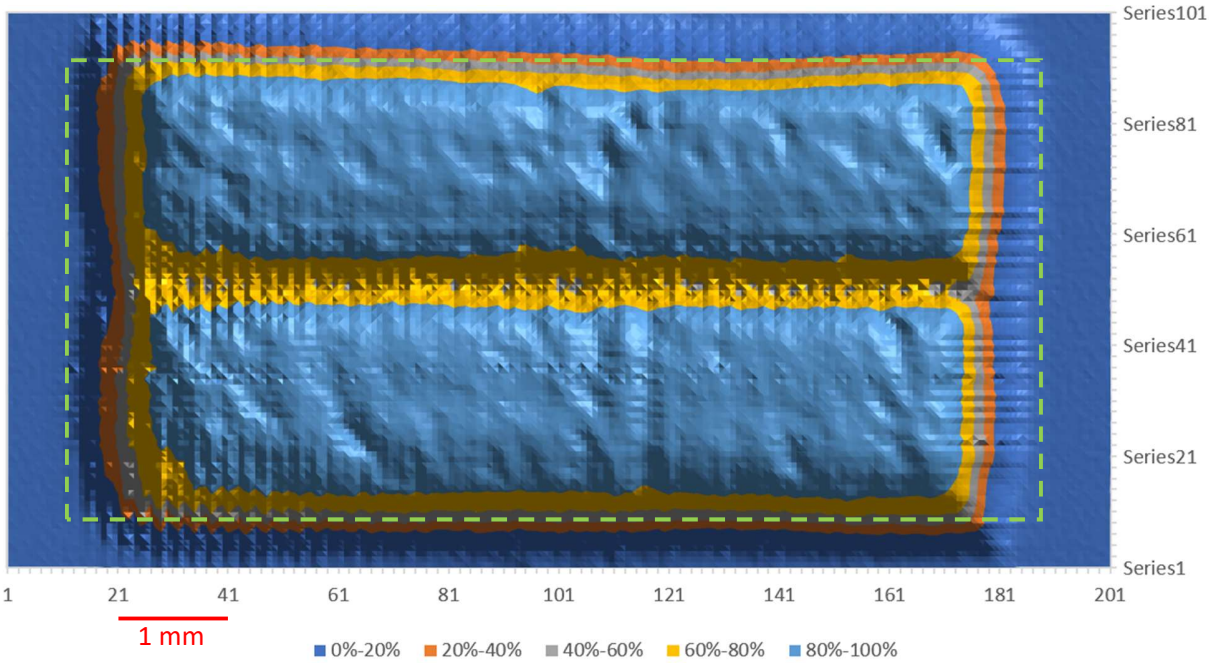

Figure 7.32 Combined optical efficiency map with adjacent Parylene-polished (top) and hand-polished (bottom) cones

The green rectangle outlines a lens-spacing by two-cone-spacings region of the map. Both the Parylene-polished cone (top) and the hand-polished cone (bottom) now have broad, flat peaks that reach 91% optical efficiency. While this is a great improvement over early lens tiles, the combined map shows that the loss at the molding kerf between the cones is significant; the cone-kerf dip costs ~6% of the light. A similar 2-cone map with unpolished cones shows that that older cone-kerf dip cost ~16% of the light. Although their effect is not similarly isolated here, it can be seen that the molding kerfs between lenses also contribute significantly to the optical losses.

7.1.30.2 Scan at a Rim-angle Slant

The previous maps were with the laser nominally normal to the lens tile. When a TLC module is on sun, light comes in to the lens tile from a trough with a 25-degree rim angle, or at all angles up to 25 degrees. The tiles were therefore scanned with the laser beam at various angles up to this angle. The scans were taken before the 1.2-degree error on the other axis was corrected (the shift of the aforementioned odd peak in the at-a-slant maps led to the diagnosis of that error).

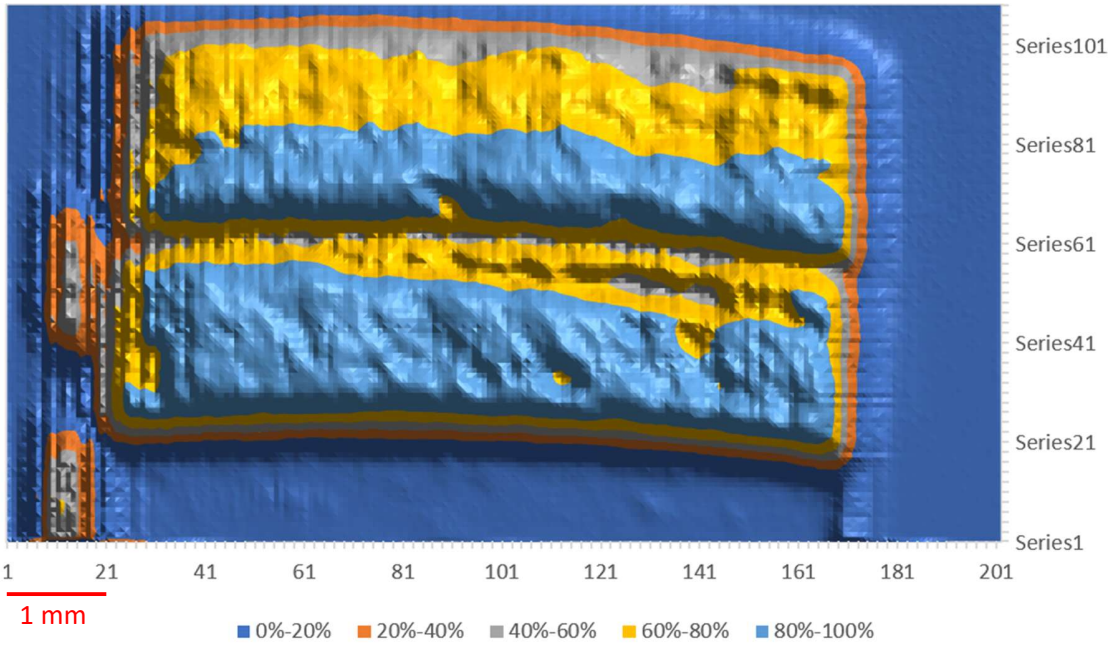

Figure 7.33 Combined map of the same lenses/cones with the laser at a 25-degree slant.

As can be seen in Figure 7.33, the cone made from the hand-polished master cone (bottom), which reaches a peak of 95%, does much better overall than that from the Parylene-polished cone (top), even though the latter reaches a peak of 96%.

7.1.30.3 Lens-to-Lens Transition

The lens-to-lens transition (kerf plus higher slant into cone tops) can be seen in Figure 7.34, and even in the newer lens tile costs around 8% of the light.

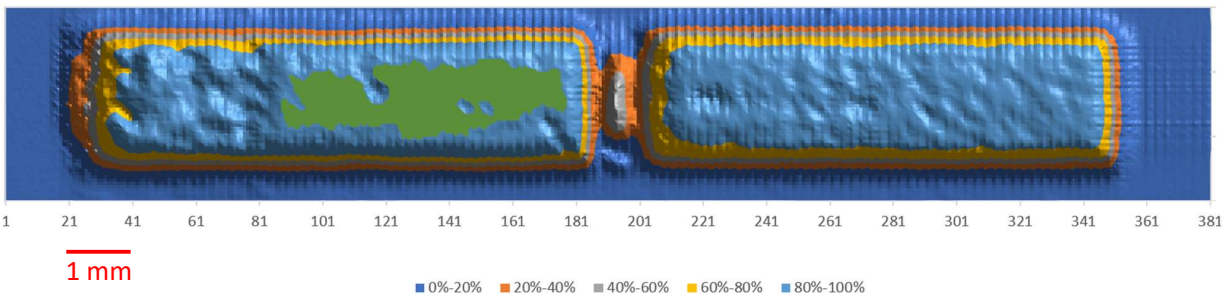

Figure 7.34 Combined map of two lenses.

7.1.30.4 Other Optical Losses

When a lens tile that is well-coupled to cell is viewed straight on, the surface of the lens is almost uniformly quite dark (cell-colored). In the prototype that went on sun, the surface is grey and non-uniform, indicating that the optical coupling from the cone tips to the cells does not provide good coupling (likely due to a module assembly misunderstanding).

While the primary mirror looked clean, the safety mask prevented it from being freshly cleaned before testing the prototype. Its reflectivity could easily have been reduced a by several percent.

7.2 Conclusion

This chapter has covered some of the extensive work that has been done to validate the TLC design, and has shown that, while there are still engineering challenges, the basic TLC design works.

A tremendous amount was learned in the experimentation phase the project. TLC is in many ways a more complex design than a typical point-focus CPV system, and TLC proved difficult to prototype (especially the cones) and to diagnose (especially the module optics). Much of this is budgetary – while a production tooling budget would have allowed for single-point-diamond-milled optics for individual prototypes or single-point-diamond-milled master arrays for producing multiple prototypes, affordable EDM-cut tooling produced useable but much rougher components.

Testing module optics without the cones (lenses only) showed that the unexpected loss of light was largely due to the cones, but the tiny size of the cones made it very difficult to diagnose where the light ended up, and multiple diagnostic setups were designed and tested. These showed that the surfaces of the cones were not guiding the light through total internal reflection as designed and as expected from the Multiphysics modeling. Extensive hand polishing, and reflowing a conformal Parylene coating, both produced much-improved results and allowed confirming that the lower on-sun efficiency was significantly due to the rough prototype surfaces.

Other discrepancies from expectations were typical of a typical prototyping project. Critically, no unforeseen issues were found in assembling a prototype and putting it on sun at the focus of a trough mirror. The low optical efficiency of the current prototype presents engineering challenges in manufacturing, especially the master cone array from which the cone molds are themselves molded (which the tested polishing processes already partially address).

CHAPTER 8

8. CONCLUSION

8.1 Summary

The design part of the project was a success. The analytical model met its goal of enabling the design of TLC modules for minimum cost for production volumes, and, after numerous design changes guided by the analytical model, a design was achieved that had very low materials cost, and low estimated process costs. While future high-volume process costs involve many assumptions, materials cost dominate PV. With materials costs justified by references, this analysis that showed that if the chosen materials can be used as designed, TLC has the potential to beat not only the then-current silicon PV pricing, but even the US Department of Energy's goal for silicon PV pricing in 2030 (Woodhouse, et al., 2016). This result (Norman, et al. PVSC, 2019) is vital to CPV, which has been dismissed by many as uncompetitive due to current CPV architecture costs.

The modeling part of the project was also successful. Ray tracing in COMSOL showed that the chosen low-cost materials could, if shaped as designed, achieve near-record optical efficiency and high concentration with adequate acceptance angles, and heat flow analysis in COMSOL showed that TLC could achieve comparable cooling to Fresnel/box CPV while using significantly less aluminum (Norman et al., 2018). The analytical and COMSOL models allowed exploring the parameter space around the initial design, which revealed that the design was fairly close to the center of a broad optimum for mirror width, mirror rim angle, and lens glass thickness. This led to another published paper on optimizing the TLC optics for lowest total projected cost (Norman, et al., 2019). It was, however, found that to achieve the target of 1500X concentration with adequate acceptance angles, a higher-refractive-index silicone would be needed, and the initial high-refractive-index silicone was a phenyl-silicone that is less durable than methyl silicone; a nano-titania-loaded methyl silicone should offer the best of both worlds (Huang, et al., 2017), so a process for preparing it should be tested.

The prototyping and testing of components and proposed manufacturing processes was mostly successful. One of the key questions arising from the earlier phases was whether the massively

parallel microcell array would induce thermal runaway in under-illuminated cells, as suggested by articles uncovered in a literature search (M. Steiner, G. Siefer, and A. W. Bett, 2013). A simple test for this was devised and the experiment carried out; this confirmed that, for the QDEC cells being used in early prototyping, this is not a problem. This novel simple test can be widely applied to parallel-cell CPV architectures so it was published as a paper (Norman, et al., 2020).

Molding of the cones on the back of the glass lens tile was identified as the biggest manufacturing challenge. A way of forming a master-array of densely packed cones and then molding flexible cone molds from the master array was devised. Prototyping of the proposed cone molding process showed that the basic process was practical, leading to a paper accepted for the CPV-17 conference proceedings (Norman, et al., 2021). But electron discharge machining of the master array left much rougher surfaces than anticipated, and experiments on smoothing the cone surfaces are ongoing.

Prototyping of the heatsink was less successful. While a thermally-adequate heatsink was machined from off-the-shelf aluminum heatsink stock, the metal was several times too thick and the heatsink lacked end caps to convert shear stress on the glass lens tile to compressive stress. Therefore, after attaching this heatsink to a lens tile with hard epoxy, thermomechanical stress cracked the lens tile upon cooling from the bonding temperature (if the heat were gathered as a byproduct, the fins would not be needed, and the problem would not exist). While sticking more closely to the design might be sufficient to resolve this, the team proposed a number of alternative attachments means for the aluminum fins, of which several were combined to make an improved low-stress design.

Design, modeling, and even component prototyping can miss unknown failure modes, so putting a prototype on sun was a key final step to catch any remaining unknown unknowns. Due to medical issues, the present candidate was less active in physical side of this than had been anticipated but led the design of the testing and test equipment, participated in the preparatory work in an advisory role and also physically participated in the final testing.

Wires for thermocouples, excess isolating epoxy, too-large standoff beads, and the use of a flexible-but-lower-thermal-conductivity epoxy as the thermal interface led to a far more resistant thermal path designed, but the prototype still survived on sun at up to two thirds of the maximum concentration. Results were consistent with the optical efficiency anticipated from rough cone surfaces and from the higher and thermal resistance to the heatsink. No unexpected problems were

observed, and the prototyping module survived the testing in spite of repeated exposures to uneven illumination. An additional paper is in progress on the on-sun testing of the first TLC prototype.

8.2 Original Contributions

The novel design of the optically coupled lens and cone optical stages at the focus of a parabolic trough was analyzed, modeled, prototyped, and tested, leading to four published papers. The proposed process for molding dense arrays of tertiary optical elements on the back of the lens tile was also analyzed, modeled, prototyped and tested, leading to an additional accepted paper. The current PhD candidate led this work and is first author on all five papers.

- Proposing the TLC concept, with versions for both single-axis and dual-axis trackers
- The lens-tile optics, and potential low-cost, high-accuracy production methods for these
- The gap-free flexible-cone coupling of the lens tile to the microcells
- The compact massively-parallel microcell-array receivers
- Quick test for cell suitability for parallel-cell-array receivers
- The receivers interconnecting in series as they are bonded to the lens tile

8.3 Research Perspectives

In addition to the ongoing work on achieving smooth cone surfaces, on verifying that AZUR 3C 44 cells are suitable for massively-parallel microcell arrays, and on improving the thermal resistance from cells to heatsink, there are several research directions that should be pursued. Two of these are the logical steps toward commercialization: first, confirming that smooth cone surfaces will achieve high optical efficiency and then single-point-diamond-machining cone plates for a master array, and second, a full-module prototype, which would confirm that no novel failure modes arise in a full-length, 32-receiver, 1 kW module.

There are also two significant design variations that should be explored. First, the discovery that internal micro-tracking of a single-axis-tracked TLC module can be done without micro-tracking the cooling or the receivers can dramatically simplify single-axis-tracked TLC. A second design

variation that fits well with single-axis-tracking but could also be used with dual-axis-tracked TLC would be to gather the waste heat from the cells as a heat byproduct. Both 400-series stainless steel and especially titanium are great CTE matches for the low-iron glass lens tile, and turbulent flow provides sufficient heat transfer to cool heat is spread out to the intensity of a trough's focus, so together, these can eliminate the cost and CTE-mismatch issues of aluminum fins without the complexity of micro-channels and can also replace the high cost of dual-axis tracking with the low cost of single-axis tracking. Since titanium tubing is already used in thermal desalination, this suggests that a single-axis-tracked TLC-based CPV/desalination system could be highly economically competitive in regions where desalination is valuable (which tend to be good regions for single-axis-tracked TLC).

8.4 Sommaire

La partie conception du projet a été un succès. La feuille de calcul a atteint son objectif de permettre la conception de modules TLC à un coût minimum pour les volumes de production et, après de nombreuses modifications de conception guidées par la feuille de calcul, une conception a été réalisée avec un coût des matériaux très bas et des coûts de processus estimés faibles. Alors que les futurs coûts de processus à haut volume impliquent de nombreuses hypothèses, le coût des matériaux domine le PV, donc avec des coûts de matériaux justifiés par des références, cette analyse a conduit à un article publié (Norman, et al. PVSC, 2019) qui a montré que si les matériaux choisis peuvent être utilisés comme prévu, TLC a le potentiel de battre non seulement le prix du silicium photovoltaïque alors en vigueur, mais même l'objectif du département américain de l'Énergie pour le prix du silicium photovoltaïque en 2030 (Woodhouse, et al., 2016).

La partie modélisation du projet a également été couronnée de succès. Le lancer de rayons dans COMSOL a montré que les matériaux choisis pouvaient, s'ils étaient façonnés comme prévu, atteindre une efficacité optique élevée et une concentration élevée avec des angles d'acceptation adéquats, ce qui a conduit à un article publié (Norman et al., 2018), et l'analyse du flux de chaleur dans COMSOL a montré que la TLC pouvait obtenir un refroidissement comparable à celui de Fresnel/box CPV tout en utilisant beaucoup moins d'aluminium. Ensemble, la feuille de calcul et le modèle ont permis d'explorer l'espace des paramètres autour de la conception initiale, ce qui a révélé que la conception était assez proche du centre d'un large optimum pour la largeur du miroir,

l'angle de la jante du miroir et l'épaisseur du verre de l'objectif. Cela a conduit à un autre article publié sur l'optimisation de l'optique TLC pour le coût total projeté le plus bas (Norman, et al., 2019). Il a cependant été constaté que pour atteindre l'objectif de concentration 1500X avec des angles d'acceptation adéquats, un silicone à indice de réfraction plus élevé serait nécessaire, et le silicone initial à indice de réfraction élevé était un phényl-silicone moins durable que le méthyle silicone; un silicone méthylique chargé de nano-oxyde de titane devrait offrir le meilleur des deux mondes (Huang, et al., 2017), donc un processus de préparation est actuellement en test par un étudiant en master.

Le prototypage et les tests des composants et des procédés de fabrication proposés ont été pour la plupart couronnés de succès. L'une des questions clés découlant des phases précédentes était de savoir si le réseau de microcellules massivement parallèle induirait un emballement thermique dans les cellules sous-éclairées, comme le suggèrent les articles découverts dans une recherche documentaire (M. Steiner, G. Siefer, and A. W. Bett, 2013). Un test simple pour cela a été conçu et l'expérience réalisée; cela a confirmé que, pour les cellules QDEC utilisées dans le prototypage précoce, ce n'est pas un problème. Ce nouveau test simple a conduit à un article publié (Norman, et al., 2020).

Le moulage des cônes à l'arrière du carreau de verre a été identifié comme le plus grand défi de fabrication. Un moyen de former un réseau maître de cônes densément emballés, puis de mouler des moules à cônes flexibles à partir du réseau maître a été mis au point. Le prototypage du processus de moulage du cône proposé a montré que le processus de base était pratique, conduisant à un article accepté pour les actes de la conférence CPV-17 (Norman, et al., 2021). Mais l'usinage par décharge d'électrons du réseau maître a laissé des surfaces beaucoup plus rugueuses que prévu, de sorte que des expériences sur le lissage des surfaces des cônes sont en cours.

Le prototypage du dissipateur thermique a eu moins de succès. Bien qu'un dissipateur thermique adéquat ait été usiné à partir d'un stock de dissipateur thermique en aluminium standard, le métal était plusieurs fois trop épais et le dissipateur thermique manquait de capuchons d'extrémité pour convertir la contrainte de cisaillement sur la tuile de verre en contrainte de compression. Par conséquent, après avoir fixé ce dissipateur thermique à un tuile de lentille avec de l'époxyde dur, la contrainte thermomécanique a fissuré le carreau de lentille lors du refroidissement à partir de la température de liaison. Bien que s'en tenir plus étroitement à la conception puisse être suffisant

pour résoudre ce problème, l'équipe a proposé un certain nombre de moyens de fixation alternatifs pour les ailettes en aluminium. Cependant, si la chaleur était collectée en tant que sous-produit, les ailettes ne seraient pas nécessaires et le problème n'existerait pas.

La conception, la modélisation et même le prototypage de composants peuvent manquer des modes de défaillance inconnus, donc mettre un prototype sur le soleil était une étape finale clé pour attraper toutes les inconnues restantes. En raison de problèmes médicaux, le candidat actuel a été moins actif sur le plan physique que prévu, mais a dirigé la conception des tests et des équipements de test, a participé aux travaux préparatoires à titre consultatif et a également participé physiquement aux tests finaux.

Des fils pour thermocouples, un excès d'époxyde isolant, des billes d'écartement trop grandes et l'utilisation d'un époxy flexible mais à conductivité thermique inférieure comme interface thermique ont conduit à un chemin thermique beaucoup plus résistant, mais le prototype a toujours survécu au soleil jusqu'aux deux tiers de l'ensoleillement maximum. Les résultats étaient cohérents avec l'efficacité optique attendue des surfaces coniques rugueuses et de la résistance thermique plus élevée du dissipateur thermique. Aucun problème inattendu n'a été observé et le module de prototypage a survécu aux tests malgré des expositions répétées à un éclairage inégal. Le test au soleil du premier prototype TLC est susceptible de conduire à un article supplémentaire.

8.5 Contributions originales

La nouvelle conception des étages optiques à lentille et cône couplés optiquement, au foyer d'un creux parabolique a été analysée, modélisée, prototypée et testée, ce qui a conduit à quatre articles publiés. Le processus proposé pour mouler des réseaux denses d'éléments optiques tertiaires à l'arrière du carreau de lentille a également été analysé, modélisé, prototypé et testé, ce qui a conduit à un article accepté supplémentaire. Le doctorant actuel a dirigé ce travail et est le premier auteur des cinq articles.

- Proposer le concept TLC, avec des versions pour les trackers à un axe et à deux axes
- L'optique lentille-carreau et les méthodes de production potentielles à faible coût et de haute précision pour ces
- Le couplage à cône flexible sans espace du carreau de lentille aux microcellules

- Les récepteurs compacts à matrice de microcellules massivement parallèles
- Test rapide de l'adéquation des cellules pour les récepteurs de réseaux de cellules parallèles
- Les récepteurs s'interconnectant en série lorsqu'ils sont collés au carreau de lentille

8.6 Perspectives de recherche

Outre les travaux en cours visant à obtenir des surfaces coniques lisses, à vérifier que les cellules AZUR 3C 44 conviennent aux matrices de microcellules massivement parallèles et à améliorer la résistance thermique des cellules au dissipateur thermique, plusieurs directions de recherche doivent être poursuivies. Deux d'entre elles sont les étapes logiques vers la commercialisation : premièrement, confirmer que les surfaces lisses des cônes atteindront une efficacité optique élevée, puis des plaques coniques à usinage au diamant à un seul point pour un réseau maître, et deuxièmement, un prototype de module complet, qui confirmerait qu'aucun nouveau mode de défaillance n'apparaît dans un module pleine longueur de 1 kW à 32 récepteurs.

Il y a aussi deux changements de conception importants qui devraient être recherchés. Tout d'abord, la découverte que le micro-suivi interne d'un module TLC à axe unique peut être effectué sans micro-suivi du refroidissement ou des récepteurs peut simplifier considérablement la TLC à axe unique. Un deuxième changement de conception qui s'adapte bien au suivi à axe unique mais qui pourrait également être utilisé avec la TLC à suivi à deux axes serait de collecter la chaleur résiduelle des cellules en tant que sous-produit thermique. L'acier inoxydable de la série 400 et en particulier le titane sont d'excellents matchs CTE pour le carreau de lentille en verre à faible teneur en fer, et l'écoulement turbulent fournit un transfert de chaleur suffisant pour refroidir la chaleur est répartie à l'intensité du foyer d'un creux, donc ensemble, ceux-ci peuvent éliminer les problèmes de coût et de décalage CTE des ailettes en aluminium sans la complexité des micro-canaux et peuvent également remplacer le haut coût du suivi sur deux axes avec le faible coût du suivi sur un seul axe. Étant donné que les tubes en titane sont déjà utilisés dans le dessalement thermique, cela suggère qu'un système de CPV/dessalement basé sur la TLC à suivi unique pourrait être très compétitif sur le plan économique dans les régions où le dessalement est précieux (qui ont tendance à être de bonnes régions pour les TLC).

APPENDIX A

Additional TLC Details and Variations

A.1 TLC Durability Compared to Silicon PV

The following table compares the durability of TLC to silicon PV using known failure modes for silicon PV modules. Since silicon PV reliability has received the most study, the risks and impacts are best understood there so a comparison to silicon PV forms the bulk of the reliability assessment; this is based on an International Energy Agency study (Kontges, et al., 2014), supplemented by an NREL study (Kurtz, et al., 2013). The HCPV sections and especially the **TLC section are smaller but more specific.**

TLC was found to be less susceptible to the main causes of silicon module failure than the silicon modules are (largely due to the compact size of the sealed module in TLC). Susceptibility less than or equal to silicon PV is considered a success and so failure modes that fall into that category are not evaluated further.

Category	Phenomenon	Risk and Impact Analysis
Silicon PV	Glass AR-coating degradation	TLC's exposed glass is typically downward facing and gets less dust and does not need as much cleaning as flat panels, and the mirrors themselves take out much of the harsher UV. The exposed-glass area is much smaller in TLC, so a better AR coating can also be afforded In SAT TLC the cover glass is replaceable (it does not touch the cells or lenses), and could be made easily replaceable. SAT TLC does have two more AR-coated surfaces, but the additional surfaces are protected behind the cover glass. Assessment: Much lower impact in TLC (than flat panels).

Silicon PV	EVA discoloring	The optical silicon (the closest thing in TLC to the EVA) is designed for CPV and sits behind on optical path that has ~25 mm of glass (mirror glass and lens glass) and hence is far better protected than normal CPV. Assessment: Low risk and much lower impact in TLC.
Silicon PV	UV exposure for cables	Exane Transit Wire cables are used; these are much higher quality and more durable than standard PV connection wire. Assessment: Much lower risk in TLC.
Silicon PV	UV exposure for junction box	Cable is sealed into glass, with connection covered with aluminum, both of which are UV-proof. Assessment: No risk in TLC.
Silicon PV	UV exposure for backsheet	TLC's aluminum is UV-proof. Assessment: No risk in TLC.
Silicon PV	Delamination / Cracked-cell isolation	The TLC cells are on a common electrical backplane, and that has CTE-matched isolation and is actively cooled. Assessment: Very low risk in TLC.
Silicon PV	Corrosion of cell and interconnect	Moisture and fluxes are the main causes of corrosion. In TLC the cells and all internal interconnections are hermetically sealed between glass sheets. The wire bonds (the closest thing to cell-to-cell interconnections) are gold wire, which does not corrode. Assessment: Very low risk in TLC.
Silicon PV	Diode failure	Diodes heating when in use is the main cause of diode failure. The TLC diodes should rarely by used (less differential shading), and are on actively-cooled cold plates and so will not overheat even when in use.

		Assessment: Very low risk in TLC.
Silicon PV	Cell interconnect breakage	A receiver is ~800 times smaller than a silicon module and the cells are ~20,000 times smaller than a wafer, and they and the wire-bonds are hermetically sealed inside the receiver and so should be shock- and vibration-proof. Even if a wire bond broke it would be only ~1% of a receiver's power, so it wouldn't even trigger the bypass diode; and even if it made intermittent contact, at less than 0.1A it wouldn't arc. The receiver-to-receiver connections are also shock-proof and vibration-proof. Assessment: Very low risk and very low impact in TLC.
Silicon PV	Contact failure / j-box	With 10 times fewer junction boxes in SAT TLC (and none in DAT) one can take twice the time per box to make sure that they are done right. If it does occur there are also 10 times fewer junction boxes to check, and in the meantime, there is nothing flammable nearby. Assessment: Very low risk + lower impact in SAT TLC, and no risk in DAT TLC.
Silicon PV	Glass breakage	Mirrors are replaceable at low cost, and there is 30x less other glass area per Watt in TLC, and it is thick glass. Assessment: Much lower risk in TLC.
Silicon PV	Loose frame	No multi-part frame in DAT TLC. With 5x less circumference and 10x fewer frames per watt, this should be much rarer even in SAT TLC. Also, the frame opening is down, rather than up, so rain won't get in, and all sensitive components inside are hermetically sealed.

		Assessment: No risk in DAT TLC, and lower risk and much lower impact in SAT TLC.
Silicon PV	Wear-out	While everything wears out eventually, none of the three main causes of wear-out identified applies to TLC. Assessment: Much lower impact in TLC.
Silicon PV	Clamping and bolting	The module clamps/bolts do not go near the glass in TLC. The mirror attachments are field-proven in CSP. Assessment: Much lower risk in TLC.
Silicon PV	Quick-connector failure	Since fewer connector ae needed in TLC, high-quality 1500V connectors are affordable. Assessment: Lower risk in TLC.
Silicon PV	Lightning	The module area is smaller, reducing the likelihood of a strike hitting a module. A strike that takes out a string would be equally damaging in TLC. A strike that takes out one module would be worse due to the higher output per module. Assessment: Fairly comparable risk and impact in TLC.
Silicon PV	Cell cracks	Should be eliminated by the small, well-cooled cells on a well-CTE-matched substrate. Impact would also be lower due to cells in parallel and lower power per cell. Assessment: Much lower risk in TLC, and lower impact.
Silicon PV	Potential-induced degradation	This has become a serious problem in silicon PV (Kurtz, et al., 2013), as newer inverters allow portions of a system to operate with negative ground. Silicon PV has thin encapsulant between the cells and cover glass, and reverse polarity causes sodium ions to migrate through it to the cells.

		TLC's silicone cones provide a millimeters-thick barrier, which prevent this from being an issue.
Silicon PV Testing	Initial Measuring	Power is averaged over >1000 cells instead of 42 to 96, so except for a serious defect all modules should pass. Assessment: Lower risk in TLC.
Silicon PV Testing	Thermal Cycling	With small cells on sufficiently CTE-matched substrates, thermal cycling should not be an issue. Assessment: Much lower risk in TLC.
Silicon PV Testing	Damp Heat	With all moisture-sensitive components hermetically sealed between glass and aluminum, moisture should not get in Assessment: Much lower risk in TLC, especially DAT.
Silicon PV Testing	Humidity Freeze	Should be no risk in DAT TLC. With all moisture-sensitive components hermetically sealed inside a weatherproof box, moisture should not get in. SAT TLC should bake out any moisture in the outer box as soon as the module is on sun, but an issue to watch for is that if humidity did get in and froze, and the system were to try to micro-track while it was frozen, this could damage the stepper motor. The stepper motor controller should monitor for this (and any other micro-tracking motion failure). Assessment: Lower risk and much lower impact in TLC, with cautionary note in SAT TLC.
Silicon PV Testing	Hot-spot Endurance	Ray tracing confirms that the Winston cones homogenize the light to avoid significant hot spots within a cell, and with all cells in parallel a receiver should be hot-spot-proof.

		Differential shading on a module should be rare in the field (bird droppings do not fall up), but it will occur in testing. The actively cooled diodes should make TLC immune to module hot spots. The only risk would be a short between the two contacts, heat spreader and front metal, but that would only be caused by a manufacturing defect and would be picked up before the receiver is placed on the module. And even a defect that escaped testing wouldn't be a hazard since it would be on a well-cooled substrate hermetically sealed between non-flammable glass and aluminum. Assessment: Very lower risk and no impact in TLC
Silicon PV Testing	Mechanical Load	The mirrors are field-proven, and the module is bolted to a sturdy steel member. In the field this is therefore just a matter of picking the right distance between supports. But the wording of the test might call for the module to be bridged at the ends, in which case the 4-meter length would work against SAT TLC. While this would not be a real-world problem, it might require working with the testers for SAT. Assessment: Low risk in TLC, cautionary note for SAT.
Silicon PV Testing	Cut Susceptibility	With all internal electrical connections far away from the metal shell, this should not be an issue in TLC. Assessment: Low risk in TLC.
Silicon PV Testing	Bypass-diode Thermal	The actively cooled diodes should make TLC immune to the bypass diode thermal test. Assessment: Very low risk in TLC.

HCPV	Cell Degradation	TLC provides active cooling with more heat rejection surface, and hence will keep the cells cooler than traditional CPV in normal conditions. However, in no-wind conditions the cell temperature will be comparable to normal CPV unless an extra heat-pipe run is added. Since cell degradation is exponential with temperature, this provides confidence that the cells' already-low degradation rate can be beaten. And this degradation rate is already an order of magnitude lower than the rate used in either of the above charts and thus the value calculations, so the rate assumed is pessimistic. Assessment: Low risk in TLC.
HCPV	Concentration Hot-spots	The Winston cones homogenizes the light onto the cell so it should not be possible to get significant hotspots Assessment: Very low risk in TLC.
TLC	Miss-tracked Trough	A trough could be miss-tracked so that the focus lands beside the receiver. This is a significant worry because this also occurs briefly when tracking onto or off sun. This could expose parts to a concentration of up to ~25 suns. **DAT**: The internal hazard is low because everything sensitive is well cooled. Externally the fins can take the ~15x focus at that height, so the only known external hazard is stray focus on the module-to-module wiring. This wiring should be strung in the shade of the module back and fins. **SAT**: The internal hazard is low because the receivers' cooling tubes shade the wires, and everything else is either made to take the sun, has coolant flowing in it, or is steel. The external hazard is primarily on the SAT housing. This is clear glass (no issue) surrounded by steel; the steel will be

		painted white with a high-temperature paint that is roughly 95% reflective so the total heating will be little more than a black metal sheet in normal sunlight. This could still heat up well above 100ºC, but the silicon used to seal the glass is good to ~200ºC so this should be OK. However, this will have to be thoroughly tested. **Either DAT or SAT**: A slight additional hazard is the module supports; while these will be metal, the wiring to the inverter will run through them. The focus would be less intense (~20x), so painting these white should also be sufficient. Assessment: Should be manageable but requires careful design and extensive testing.
TLC	Miss-tracked Receivers	Receivers could be miss-tracked (either micro-tracked in SAT, or with the whole trough in DAT) so that the foci land beside the cells. But ray tracing shows that the foci leak from the cone tops, not bottoms, so this does not exceed the trough's primary concentration which everything in the module can easily handle.

Table A.1: *Table of Reliability / Durability Risks.*

A.2 Other Variations of TLC

The version of DAT TLC this thesis focuses on is not necessarily the best TLC. Many design decisions have been made for convenience for prototyping or low costs at introductory scale, or for the target introductory market. The following variations fit within the broad TLC context:

Category/Name/Description	Effects and Reasoning
Mirrors: Lower Rim Angle A lower rim angle would consume less angular budget for the trough and give more to the non-imaging TIR cones.	**Effects:** A lower rim angle allows higher concentration, higher acceptance angle, or a combination thereof. Wider receivers would raise lens and substrate costs. Focal length also grows longer **Reasoning:** This was an early design and is still a viable choice. But RP3 mirrors for introductory scale saves early costs, current cells lose efficiency at higher concentration, and the low-error optics don't appear to need a higher acceptance angle. The receiver has already been widened, gaining much of the reduced handling advantage of wider receivers that a longer focal length would allow.
Mirrors: Wider Mirrors Wider mirrors at the same rim angle would allow larger receivers and also more power per module.	**Effects:** Larger receivers and more power per module reduce parts handling costs. Limit is 5"x7" substrate size for many high-speed assembly machines. Electrical losses rise slightly. Reasoning: For DAT TLC this would interfere with what appears to be the lowest-cost cooling, so if anything, mirrors get narrower.

	For SAT TLC this is still in the design for very large-scale production, but wider receivers could still reach ~1200X with RP-3 mirrors so much of the incentive is gone. However, cells with more junctions will increase cell voltages and lower cell currents, which will drive the ideal receiver size larger. Larger receivers also reduce costs if the coolant feeding proves more challenging than expected. Wider modules would allow eliminating the dual-off-center design with its twin receiver module, saving some costs there.
Mirrors: Other Mirror Types Thin-glass, polymer and metal mirrors are also used for troughs.	**Effects:** Few TLC-specific effects. Lower specularity of non-glass slightly affects 2nd-axis focusing. **Reasoning:** TLC will work with any other kind of mirror suitable for CSP, but glass mirrors are low cost, durable and highly reflective and so are used initially.
SAT Trackers: On-Slope Mounting The tracker axis can be mounted on an equatorial-facing slope to reduce latitude effects.	**Effects:** On a natural slope: Reduces latitude effects and increases annual output. On supports that mimic a slope – hard to gain more than a degree or two without trough supports becoming expensive. **Reasoning:** Mounting on a slope slanted toward the equator has a very similar effect to installing at a site closer to the equator; for example, in at 30ºN an installation on a 5-degree south-facing behaves similarly to a horizontal installation at 25ºN (the effect is not identical – for example mounting on a slant does not change sunrise or sunset times, while a change of site latitude does). Use natural slopes!

	A slope can be mimicked with trough mounts that raise one end, but for long troughs (to limit end losses), this is not thought to be practical.
SAT Trackers: At Latitude Mounting The tracker axis can be mounted at-latitude to minimize end losses receiver tilts.	**Effects:** Allows much higher concentration. Requires a shorter trough to not have one end ridiculously high. **Reasoning:** At-latitude slant is somewhat like mounting at the equator, and while this frees lens focal length and allows both higher concentration and larger cells, these problems have been solved in other ways. However, if future cells go beyond 2000x or go up significantly in cost with higher efficiency, this could make sense as still lower cost than two-axis tracking. This would go well with a module that gets moved N/S as it micro-tracks, eliminating end losses and allowing short troughs.
SAT Trackers: Linear Fresnel Mirror 'Trough' TLC could also be applied to low-cost linear Fresnel reflectors.	**Effects:** Lower concentration, much larger modules and receivers **Reasoning:** Linear Fresnel reflectors (LFRs and CLFRs) need to be very large to obtain high concentration from a low rim angle. However, this could be pursued if cell costs, and thus concentration needed, drop dramatically.
Trackers: Untracked Primary Two-axis-tracked secondary with two-axis-secondary concentration	**Effects:** Lower-cost primary concentration, but lower final concentration **Reasoning:** The single-axis tracker costs twice as much as the mirrors, and an E/W trough with its focal plane at an at-latitude slant would eliminate that cost. This would be a return to Miller and Stephens, but with modern mirrors, TFI's mirror supports and alignment, and TLC's lens tiles. A 'guestimate' is that spherical aberration would reduce the

	concentration very roughly 2x- at today's cell costs this is unlikely to win - 2x would roughly balance savings from eliminating the mirror tracking, and the module would also grow larger, and the benefits of tracked primaries would be lost. But if trackers stop dropping in cost and cells do drop in cost, this could be worth considering.
SAT Module: No Cover Glass: A cover glass should not be needed since the glass-fronted modules are hermetically sealed. A metal shroud behind the receivers could protect the micro-tracking.	**Effects:** Removing the cover glass would remove a 5% loss of light, thereby increasing efficiency and decreasing both cost/W and cost/kWh. **(DAT doesn't use a cover glass.)** However, receivers would be less well protected, and would be exposed to module cleaning **Reasoning:** This was an earlier SAT design, and it is still a viable choice. However, a cover-glass just looks right on a 50-year product, and the value of 50 years outweighs the efficiency cost of the glass. Anti-reflective coatings are also making rapid progress so the efficiency cost of the cover glass is likely to be reduced (and a protected inside at 25x concentration will make such coatings affordable for TLC long before they work for competing architectures).
SAT Module: **Glass Tent:** A cover glass 'tent' could cover the whole trough.	**Effects:** Trading the local cover glass for a cover glass tent over the trough is a wash optically (if same AR coating, and when clean), but it reduces dust sensitivity and also removes the module shell to help balancing the greater glass cost. **Reasoning:** This is well worth looking at for dusty areas. A 'tent' of flat-panel cover glass only takes ~1.4x the mirror area, or about 2¢/W at \$5.50/m^2 (DOE, 2016), plus some framing. While the direct savings would only be about 0.7¢/W, dust would be on a one-pass optical surface for lower

	sensitivity and thus lower cleaning cost. The fin tube would be above the tent so that it still catches the wind. If the mirror backs were sealed, this could allow first-surface mirrors, but for CPV the gain is too small to be worth risking.
SAT Module: **Glass-House**: The whole solar field could be glass-house covered.	**Effects:** Trading the local cover glass for a glass-house is a wash optically (if same AR coating, and when clean), but it reduces dust sensitivity and also removes the module shell to help balancing the greater glass cost. A glass house also would allow lighter mirrors and tracker structure. **Reasoning:** It sounds expensive, but GlassPoint's trough-in-glass-house solar thermal technology (Trabish, 2013) was used to build one of the world's largest solar installations. Similar to the 'tent' above, but this would allow lighter mirrors and lighter tracking, offsetting the greater use of glass. GlassPoint focused on enhanced oil recovery, but could be a natural fit for TLC in very dusty regions. And after the nearby oilfield is played out, the current thermal receiver-tubes could be replaced by TLC modules to use the same troughs for PV. However, the cooling loses the benefit of the wind unless the heat is exported from the glass house (works great if heat is a useful byproduct, such as if greenhouses are nearby. The world currently has >35 billion square meters of greenhouses; filled with TLC at 35% efficiency at 1/3 packing that would be ~4 TW of solar, so the idea is not outlandish.
Receiver Rows: Rows in series Keeps current low	**Effects:** All rows on a receiver are currently in parallel, but they could be in series. If the top copper is a strip under each row, the cells in the row can be wire-bonded to the next strip connecting the rows in series. With DAT this would also

	allow longer receivers, reducing handling costs up to the limits of wire-bonding machines. **Reasoning:** Currently this would increase the diode cost, but with wider troughs, or with 2-part lenses, this could reduce resistive losses. But the module voltage also might become too high (especially with future cells), shortening the inverter strings and requiring enough more strings to raise the cost.
Oxidized AlSiC: Insulator under cells Saves adhesive insulator costs, and high thermal conductivity.	**Effects:** AlSiC can be made with aluminum skin, which can be oxidized to alumina, which is a thermally conductive electrical insulator, and copper foil can stick to alumina. Even without a skin, SiC oxidizes to silica, which is a good electrical insulator, and a good enough thermal conductor for the spreader front. Oxidizing the back to an alumina/silica mixture capable of withstanding ~5000V might be more challenging but should be low cost in very high volume. **Reasoning:** Less suitable process for prototyping, but under consideration for high volume. Also goes along with receiver rows in series due to lower thermal penalty under the cells.
Heat Spreader: Cast Iron Least expensive CTE-matched material.	**Effects:** The lowest-cost material well-CTE-matched to the glass and cells and with tolerable thermal conductivity is cast iron; at ~$1.20/kg, even a 4 mm spreader is only ~0.25¢/W. While a great CTE match at 9.8 PPM/K (better than low-cost AlSiC), its thermal conductivity is only 30% as good. However, a cast-iron spreader could have a copper insert under each cell, which would eliminate the thermal penalty. A 1.25 mm diameter copper wire would quickly spread the heat to its 16 mm cylindrical surface, an area 20 times larger than the cell, and from there it would spread through the cast iron with resistance starting at a low 0.6K/mm and getting

	lower. In high volume stack of 3 mm-thick spreaders could have holes drilled and have chilled work-hardened copper wires inserted; these would expand enough on warming to press the copper into the iron to obtain good thermal contact and would not work-harden further during the days and shrink on a winter night. A similar technique was used for hundreds of years for fixing holes in cast-iron cookware, which encounters considerably more thermal and environmental stress than well-protected receivers. And the 50x large area of cast iron would constrain the copper for an overall CTE close to that of cast iron, or ~10 ppm/K. **Reasoning:** The small cells mean that this should not be necessary, especially as the CTE limit of rolled AlSiC is steadily improving.
Heat Spreader: Silicon Matches borosilicate glass CTE.	**Effects:** A slab of metallurgical-grade polysilicon would be a very low-cost heat spreader matched to a borosilicate glass. **Reasoning:** The best glass is currently a low-iron soda-lime glass with a higher CTE than silicon. Also, the electrical conductivity of the copper is currently useful (however with rows-in-series this advantage would be reversed).
Cells: End Contact(s) on cells Saves reverse-bonding on bond wires.	**Effects:** A cell is enough longer than it is wide that a landing pad on the end of a cell, or even on each end, would cost less area than a bonding-pad-width bus bar along one edge of a cell. A pad on the end could even be narrower since the bond wire would run parallel to its length rather than is width, so only ~40μ would be needed for a 0.7-mil wire that could handle the current. **Reasoning:** The central bur-bar design consumes little active area anyway, but it requires either an extra notch in the lens

	to accommodate it, or reverse-bonded wires to be flat enough to fit under then lens. If these prove a pain, then end pads could be reconsidered.
Lens Alternative: Light-guide Optics Flatter lens alternative with potential similar to two-part lens	**Effects:** Light-guide optics can be much flatter for a given focal length than a lens. This could allow larger cells for a given trough size and concentration. **Reasoning:** Flatter optics allow bigger cells (and probably clearer glass), without much other change (e.g., the same cooling could be used), so this is something that is under serious consideration. The current lens has an 18-mm path through glass and a 2.5 mm path through silicone to get a 0.65 mm cell width, light-guide optics with an average of a 1 cm path through glass and almost no silicone could reach over twice the cell width. This would reduce cell costs and cell placement costs, reduces glass volume and thus cost, totaling potentially as much a 2¢/W in small volume and ~1¢/W in large volume (roughly 10%). But cooling of the cells would be harder unless the light is turned downward, which would be difficult with over-molding on a roll-formed glass sheet. If these can be solved without added cost, the initial concentration by the trough would reduce the lens area needed and produce a much shorter optical path in the light-guide for a given cell size, lowering the cost and improving the efficiency over current LGO systems. However light-guide optics rely on a low-refractive index gap to confine the light through TIR, resulting in the same hit to the efficiency as a two-part lens. The two-part lens appears to be simpler and has the same cell-width advantage.
Lens Alternative:	**Effects: Relaxes receiver-to-lens alignment**

Separate TIR tiles	**Reasoning:** The TIR cones need tens-of-micron alignment to the cells, while the lens-to-TIR alignment tolerance is much more relaxed – a few hundred microns along the trough's focus and up to millimeters across the troughs focus. Receivers with cone tiles could be place on the lens with 100 µm, versus 10 µm, tolerance. **Challenges:** The cone tiles would still need the precision placement on the receivers. Cone tiles would also create an extra optical interface which needed to be fairly bubble-free since bubbles would reflect 10% of their area worth of light and would distort the other 90% of their area. Since placing receivers on the lens eliminates the alignment of large parts, this is no longer needed unless over-molding cones on the lens back proves impractical.
Lens Alternative: Air-gap Lens	**Effects: Reduce absorption loss while allowing wider cells** **Reasoning:** The glass between the graceful curves of the front of the lens and the cones at the back absorbs light in addition to adding cost, so why not get rid of it with a lens with two separate layers? Since the alignment between the upper curve and the body of the lens is more relaxed than the alignment to the cells, and the glass is clear enough that a thin lens will not overheat with convective cooling, a two-part (hollow) lens is actually a workable approach. The main drawback is the light reflected back at each refractive index change, which will cut power by ~2.5% per air/glass interface with today's AR coatings. For a narrow lens the savings are simply not worthwhile. Better Anti-Reflection coatings would help air-gap lenses. The extra interfaces are protected, allowing a nano-textured coating, reducing the loss to only ~0.5% (Ulrike, et al., 2015).

	The TIR cones need tens-of-µm alignment to the cells, which would be easy if placed as separate per-receiver cone tiles. The lens-to-TIR alignment tolerance is more relaxed – a few hundred µm along the trough's focus and up to millimeters across the trough's focus.
Lens Alternative: Glass-cored cones	**Effects: Provides greater cone stability** **Reasoning:** Glass less expensive and even more stable than silicone. While traverse features are harder to roll-form accurately, over-molding removes the need for accuracy. The cones could be as close to the desired net shape as practical, with the remainder comprising over-molded silicone. There is no magic number –*none of a cone has to be glass* (and none will be initially). But glass would be even better (until the roll forming gets difficult), so somewhere around 1 GW glass-core cones *might* be introduced. 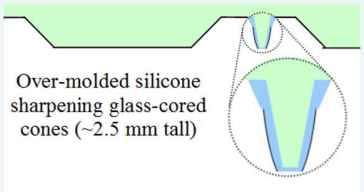 **Figure A.1:** *Glass-cored over-molded cone.*
Micro-Tracking Height (SAT) Lens-centered axis	**Effects:** Rotating about the top-center of the lens produces the highest concentration at low tilts, and also eliminates horizontal lens travel during micro-tracking. **Reasoning:** This is not used because when optimizing for latitudes 0º-37º, the concentration is limited by the spread of the trough's focus at high tilt, and the concentration can already exceed that desirable for today's cells. Lens-centered cooling also keeps the receiver from moving laterally during micro-tracking, so even the receivers at the ends of a module stay within the module. However, unless the micro-tracking

	drive arms are very short (making accuracy harder), the arms would not stay within a given module anyway, so horizontal travel is accommodated by hollow ends on the modules that bolt together.
Micro-Tracking Height (SAT) Cooling-centered Lens axis	**Effects:** Coolant tube axles directly from the manifold would simplify the hosing or the tubes by removing a bend. **Reasoning:** While the height of the center of the lens would change slightly as the receiver rotates; it would not be enough to keep the trough focal width even on the lens, and so cams or other height adjustment means would be needed and these are more complex than adjusting the axis of rotation.
Micro-Tracking Height (SAT) Micro-tracking height cam	**Effects:** A versatile way to control receiver height in the focus is to use mechanical cams in the micro-tracking linkage to automatically adjust the height according to the tilt. **Reasoning:** With coolant tubes as axles, it would have to lift the coolant pipes, too, and the offset axle does almost as well with fewer hassles. It is also possible to raise and lower a whole module, but this required more force (especially when wind, snow and freezing rain are considered.

Table A.2: *Table of TLC variations and effects.*

A.3 Alternate Aluminum-Fin Heatsinks

- ***<u>One-Piece All-Aluminum with Skived Fins:</u>***

A very similar one-piece all-aluminum heat-sink lid could also be made through skiving. Skiving does not have the extrusion-width limitation (this direction is the length in skiving), so the welding step would be avoided and heatsink could be truly one-piece. Decaluma (JINAN DEGA MACHINE CO, LTD.) makes a skiving machine that skives 1.2 mm fins up to 500 mm wide; there is no fundamental limit on skiving width, so multiple heatsinks could be cut from a skived block.

With a standard start/stop, fin-at-a-time skiving process, a skived-fin heatsink would not be as cost-effective in volume as an extruded-fin one would be unless the volume were high enough that one operator could manage multiple machines. Skiving thick fins also requires a ~5-mm-thick base, but this could be thinned to 2 mm or 3 mm to keep the assembly CTE in the target range; the extra 3 mm of aluminum is roughly the same as the amount saved by tapered fins (and is the shaved aluminum is worth recycling, at roughly half the raw material price). Skiving is a backup in case welding or brazing of extruded profiles fails.

- ***<u>All-Aluminum with Extruded fins on Extruded Lid:</u>***

The base and the fin-pairs can also be extruded separately, and fin-pairs welded or brazed to the base. The aluminum mass is similar to the one-piece heat sink above, and a six heat-sink block is still welded (or brazed), but the machining is different. The shearing of the fin pairs is highly automatable (and could punch out base that would block airflow), so shearing the 97 fin pairs is allocated 2 minutes at $60/hour, or $2. The welding is in thinner aluminum, which would be faster, but it consists of 40 mm welds rather than 1-meter welds, which slows it down, so the same 1260 mm/minute is allocated. Six heat sinks is 18.5 minutes of welding plus 0.5 minutes for change-over, or $54.70 of welding, or $9.12 per heat-sink (and 18 minutes is ample time to fixture the next assembly to weld).

No folding is needed, leaving just one cut to singulate a heat sink; corner trimming is still nice, so the total is $2 of post-welding processing. The total is thus $(17.10+9.12+2+2) = $30.24 = ~2.85¢/W. Post-welding polishing of the bottom might be needed for a tight thermal interface, so is allocated until a brazing process can be confirmed.

While this is more expensive when welding is used, a custom induction-brazing jig might make this lowest-cost processor of all. Multiple extruded lid-flange profiles would be placed in guides in the jig base and buttered with the braze compound. 97 full-length fin-pair extrusions would then be placed in the jig top, which would press them against the braze on the lid-flange profiles in the base. The jig top would have a magnetic-steel patch at the intersection of each fin-pair and flange, and coils in the base would then rapidly heat the patches of magnetic steel to braze temperature in a traveling wave. This would produce a lid-block from which 30 to 50 lids could be cut using parallel saws (like a bread slicer). In even higher volume, a custom sheer could singulate a heatsink lid with each stroke and perform the corner trimming and base punching at the same time, and fin

pairs could be stamped from a tapered profile. With all processes highly automatable, this could lower the cost per lid toward 2 ¢/W. An extruded base also allows varying the profile thickness, which could stiffen corners, thin non-critical zones, and better balance glass-to-aluminum CTE mismatch stresses, so this offers the potential of the best heatsink at the lowest cost.

A.4 CTE-matched-Heatsink Lid (backup)

Until the higher overall CTE from the all-aluminum heat-sink lid is validated (first through COMSOL analysis and then through thermal cycling of cells-on-substrate and heatsink-on-glass), a CTE-matched option is kept as a backup. The CTE-matched design comprises bonding folded or extruded fins to a strong, low-CTE module lid. There are two affordable choices for strong, corrosion-resistant materials with suitable thermal expansion and thermal conductivity: martensitic stainless steel, and titanium. Both of these are lower than optimal thermal conductivity, but are acceptable if painted bright white (e.g., titanium-dioxide-white) for the off-focus when tracking onto or off of sun).

- ***Titanium:***

Titanium is considered first because it is an outstanding CTE match to the lens at 8.6 ppm/K, making a hermetic seal simple, and it is adequately thermally conductive at 21.9 W/mK. The lens-tile back is a simple folded-sheet part, and titanium folds well so the 'machining' is just two slots for the terminal wires, and bolt holes for mounting (the end would be folded down which would produce mounting tabs from the extra width). These could be punched before folding, or, if laser cutting is used, be part of the cut pattern. The exposed sides would be painted white after bonding the fins.

- Titanium Cost:

With the great CTE-match to the lens glass, long flanges to balance the CTE stress of the lid's top are not necessary so the 'hot flange' can be ~12 mm shorter than in the all-aluminum heat sinks shown above (while still sealing well to the glass). The mounting flange can also be shortened at least 8 mm, so the total width of the titanium sheet is ~120 mm. This makes a 1 mm thick lens-tile back ~120 mm wide, or 206 cm^3 per module. Although titanium costs more than martensitic stainless steel, it is still affordable; with a density of 4.5 g/cm^3, the current price of $10,200/tonne (Metalary, Dec. 2017) is only 4.6 ¢/cm^3, or $9.50 per module or 0.91 ¢/W. While this is roughly

four times the cost of a thicker stainless-steel part, it still adds less than 1 ¢/W (and a single-module sheet for prototyping would be less than $100 from onlinemetals.com).

Aluminum can be bonded to titanium more easily that to steel because iron and aluminum form weak intermetallics. In fact, one way to bond aluminum to stainless steel is to coat the steel with titanium and then stick the aluminum to the titanium coating. So, if aluminum fins are brazed or welded to the housing, titanium is simpler, at least in the short run.

- ***Martensitic and Ferritic Stainless Steels:***

Martensitic and ferritic stainless steels have lower thermal expansion than the more common Austenitic stainless steels, and several of the low-cost versions are formable and weldable, making them mechanically suitable for the TLC module housing. While their corrosion resistance is not as high as marine-grade steel, they are resistant to water (and to methanol coolant) at all operating temperatures. Alloy 409 in typically the lowest-cost stainless steel, and alloy 430 is typically the next least expensive, but alloy 410 has a better CTE match and is still highly affordable at just 0.12 ¢/W more than the lowest-cost alloy 409.

Alloy 410 stainless steel is thus a good material for the module lid, being strong, formable, laser and induction weldable, inexpensive, and chemically resistant to key coolants as well as having acceptable thermal conductivity (~25 W/mK) and a CTE (9.9 ppm/K for 0ºC to 100ºC) between the 8.6 ppm/K of the lens and that of a rolled AlSiC heat spreader (11.5 ppm/K). In 2012 alloy 410 was 89¢/lb (AK Steel, 2012), and steel prices now are about the same as in August 2012. At a density of 7.8 gm/cm^3 that is ~1.5 ¢/cm^3.

Steel allows a full-width mounting flange (easier to bolt to) so a 128-mm width would be used. A 2 mm thickness would eliminate the need for gussets (even 1.5 mm might be sufficient), and 2 mm is 450 cm^3 of steel, or $6.72 or 0.64 ¢/W.

In the longer term, it is probably inexpensive to thermal-spray coat $$410 with titanium. This could be done as a long strip or even as a wide sheet before cutting the individual module covers. A strip 0.25 mm thick by 40 mm wide would be only 0.08 ¢/W of titanium, and would allow using steel that is thicker, stiffer, and stronger than a titanium back, while saving ~0.3¢/W.

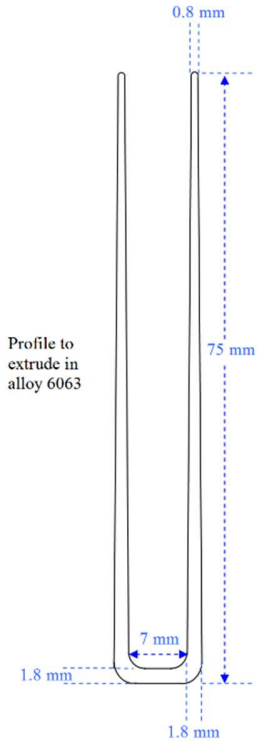

Figure A.2: *Fin-pair Profile.*

- ***Preventing Corrosion:***

Aluminum provides by far the highest heat conductivity per dollar of any durable solid material (graphite is too brittle, and sodium is too flammable), and its thermal conductivity is sufficient for the heat density and distance involved with cooling the focus of a single RP3 mirror. Aluminum and titanium have different enough galvanic potentials that a coating such as epoxy would be used over the Al/Ti bond region to seal out moisture. 410 Stainless is a better galvanic match to aluminum, but a sealant is still advisable.

With a sealed seam and white-painted cathodic metal, galvanic corrosion should be very limited, and thus harmless when spread over the huge anodic fin area. However additional protection can easily be provided (for example, for environments where salt is present) by providing sacrificial strips of zinc or magnesium (which are low cost as well as being less noble even than aluminum).

- ***Extruded Fin Pairs:***

There are numerous ways to make aluminum fins, of which extruded fin pairs would offer the advantages of tapered fins while using the least material. Fins would be extruded in long lengths

(5 to 8 meters), and a shear would then cut off 140-mm lengths. To keep the 'floor' or the 'U' from blocking airflow, it would be trimmed (except for the 40 mm to be bonded to the lid) as the fin-pairs are cut to length; these trims almost 5% of the aluminum (which gets recycled). The aluminum mass needed is similar to the $17.10 for the welded extrusion costed earlier.

- ***Fin Brazing to Titanium or Ti Coating on Steel:***

Aluminum can be brazed to titanium sheet with 17 µm (½ oz) copper foil – the foil forms a Cu/Al eutectic that wets both titanium and aluminum (Shaprio, 2017). ½ oz Cu foil is low enough cost to afford a continuous 40-mm-wide strip per module rather than a strip per fin; at 0.068 m^2 this uses 11 grams of foil. ½-oz foil is currently available at $9.92/kg (with raw Cu price at $6.72/kg) for a 10 MW quantity (and ~5% less at 1 GW), so this is only ~11¢ of foil per module or 0.01 ¢/W. Alternatively Al88Si12 is used to braze both titanium and aluminum to dissimilar metals, and thus can join the two together, but it is expected to be less convenient.

A module-length 'comb' clamp can have steel teeth (half as many as fins), each ~6 mm in width to fit between fins, a few cm tall for stiffness, and ~ 12 cm long to span the titanium cover. This can clamp all bottom-folds against the cover, with the copper foil sandwiched securely in between. This can then be run through an over in an inert atmosphere at a rate that gives the copper and aluminum time to form a eutectic melt. Since titanium does not burn in nitrogen until ~800ºC, a nitrogen atmosphere should be sufficiently inert (vacuum brazing produces even stronger joints but should not be needed for a non-aerospace application). The melt should form fast enough to make induction heating practical. Aluminum-to-copper brazing is not done in a furnace because the eutectic forms fast enough that unless the amount of copper is limited it will eat the aluminum before the furnace can cool. But what makes copper-to-aluminum brazing hard works in favor of brazing aluminum to titanium with copper foil by ensuring a fast process.

U-shaped fin pairs are extruded from aluminum (or stamped from sheet), and a lid is stamped from alloy 410 (possibly needing to be copper or titanium coated). Fin pairs are slid over the teeth of a full-heatsink jig (with the fins hanging down, and no precision needed), and wire-brushed with a fluxed brush. A fluxed copper strip is placed along the row of fin bases, the lid is clamped against the foil/fins. Either small induction heads are activated in series, or a small head is run along the module. 6 kW will heat a receiver-sized lid-region to braze temperature in ~1 second, so brazing totals less than a minute (and the fins rapidly cool the brazed assembly).

Alloy 410 is magnetic, so induction will heat it, and it will then heat the braze and then the fins, so the fins will be cooler than the braze as well as having a higher melting point1s. For titanium the clamp would be mostly non-magnetic, low-thermal conductivity stainless steel, with only a contact pad directly against the titanium being made of magnetic steel. Induction thus heats the part of the clamp that is against the titanium, when then heats the braze, so again the fins are below the braze temperature.

GH Induction Atmospheres would be a good place to start for a quote, but a rough estimate can be made. If placing a fin pair takes 5 seconds, then loading the pairs is just over 4 minute, or 5 minutes including rushing with the flux. Add a minute for placing the copper and closing the lid, one for the brazing, and one to unload, and the station is occupied for 8 minutes per heat-sink. This is low-skill labor, so $1/minute is pessimistic (but affordable. Shearing fins is low cost, so the total non-material cost should be ~$10 (and much of this could be automated at 1 GW).

The total cost should thus be roughly $24.25 for the aluminum, $9.50 for the titanium, and $10 for the processing, or $44, which is 4.2 ¢/W, and ~41 with steel, or 3.9 ¢/W. While a CTE-matched heatsink is only ~1 ¢/W higher in cost than an all-aluminum heatsink, a one-piece heatsink with high-thermal-conductivity flanges, and at a lower cost, is preferred.

- ***<u>Fin Stress:</u>***

The aluminum fins could have slit or folded bottoms to ensure no contiguous area larger than 1 mm x 1 mm bonded, but this should not be necessary. At aluminum's modulus of 69 GPa the CTE mismatch of ~15 ppm/K produces ~ 1 MPa/K of stress, so a ~550K temperature drop from brazing produces much more stress the yield strength of common aluminum alloys, and the aluminum will yield to relieve stress on initial cooling from brazing temperature. The brazing will anneal the aluminum, so it will be soft and will retain only ~30 MPa of residual stress).

The aluminum will rapidly strain-harden from any further temperature cycling and will also age-harden. The +/- 60K from winter night to summer on sun would only produce ~~60 MPa, which is less than the yield strength of even ¼-hard H12-temper aluminum (Aluminum, 2013), so the aluminum rapidly reaches the point where it does not plastically deform any further.

A.5 Heat Pipe with Fin Tube:

A heat pipe solves the issue of fin conductivity by using the latent heat of vaporization to distribute the heat to a condensation section where the heat is travels only a short distance through a sea of fins. However, a heat pipe needs to be hermetically sealed, and the distances are long enough to need gravity return (at least in the simplest fin arrangement) for the condensed coolant.

- Per-module Heat Pipe:

For a dual-axis tracker a whole-module heat pipe is practical. The fin tube per Watt is constant. The steel per Watt grows with the cube root of the trough width due to needing thicker steel, but this is a minor factor. The welding cost is driven by the module length, so it remains almost fixed during scaling, and thus on a per-Watt basis it decreases almost proportionately with the module width (out to several times the RP-3 mirror width).

A heat pipe transfers the heat so efficiently that the cooling scales well – just add more rows of fin tube. Although the slight slant needed to get gravity-assist for the return eventually pushes the fin tube to where it shades the mirror, enough for RP-3 fits the RP-3 mirror off-axis offset, and wider mirrors could easily be made with a proportionately wider offset.

A per-module heat-pipe would be difficult with SAT TLC because coolant could accumulate in some receivers while letting others starve for coolant and overheat.

- Per-receiver Heat Pipe:

For a dual-axis tracker a series of small heat pipes could also be used. With each heat pipe having a vertical length of fin tube, each heat pipe would be trivial to seal by welding or even epoxy. However, the individual heat pipes then have to be attached to a module back, negating the gain of simpler sealing. The natural fin orientation is also horizontal near noon so much more fin area is needed due to the poor air flow through horizontal fins when there is no wind, so the cost appears to be significantly higher.

For SAT TLC the high thermal conductivity would allow a sparse enough array (e.g., fins at several different heights) to get around the micro-tracking issue, but not the fin-orientation issue.

A.6 **Pumped Coolant with Fin Tube:**

Pumped coolant allows scaling to very wide troughs by forcing a liquid coolant across the back of the heat spreader, and thin films of flowing liquid transfer heat very effectively. Liquid can also distribute heat to multiple fin arrays that can use shorter fins to have higher heat transfer coefficients, thus reducing total fin area. Since gravity return is not needed, the fin tubes can be stacked below the module where they don't shade the trough. However pumped liquid cooling simply isn't necessary for DAT TLC.

For SAT TLC pumped liquid cooling is thought to be the simplest practical cooling, and long troughs naturally lend themselves to long lines of long modules that can have pump redundancy at low cost. Each receiver needs its own coolant feed and return, which can be through grommet seals or flexible tube. In either case minimizing the number of connections pushes toward larger receivers and wider troughs.

APPENDIX B
TLC ON SINGLE-AXIS TRACKERS

B.1 Single-axis-tracked TLC without Micro-tracked Receivers

This document section has not been updated to reflect the February, 2020 discovery of a way to micro-track optics within the TLC module without micro-tracking the receivers, and thus without micro-tracking the cooling. This discovery makes it possible to use low-cost turbulent flow cooling, instead of micro-channels or a flowing film of coolant. In particular, this works well for heat capture for uses such as desalination, where it not only provides heat, but reduces the cost and the complexity of the module (relative to the SAT TLC described below) and dramatically reduces tracker cost compared to DAT TLC.

B.2 Single-Axis vs. Dual-Axis TLC

While dual-axis-tracked (DAT) TLC has been the focus of this document, TLC can also work on a single-axis tracker (SAT). In SAT TLC, the trough, with the module attached, is tracked daily from east to west as CSP troughs are single-axis tracking is very low cost and very accurate (so low cost that even inefficient thin-film panels are often tracked to lower energy cost). Within the module receivers with small lens tiles are ganged together so that a single drive unit per module can rotate all the receivers at once to keep each lens pointing directly at the sun. This 'micro-tracking' moves a small mass a short distance within a protected module, and a single small stepper motor can micro-track many kW, so this costs much less than a two-axis tracker. However, the module has a cover glass to protect the moving parts, which reduces optical efficiency, and micro-tracking makes the two lowest-cost cooling options difficult. Together with the lower capacity factor from one-axis macro-tracking these offset most of the lower cost from one-axis tracking for roughly equal costs overall.

While two-axis trackers currently cost more than twice as much per area as even CPV-grade single-axis trackers, they allow simpler receivers with no micro-tracking, a lower-cost cooling option, lower assembly cost, and higher optical efficiency. While cost calculations currently show a

slightly higher cost per Watt once the tracker is included, the higher capacity factor from two-axis tracking roughly compensates for this.

TLC on a two-axis tracker has higher efficiency, low module cost, and lower development cost and risk than TLC on a single-axis tracker. **Why even consider single-axis-tracked TLC at all?**

- *Tracker Cost*

The main reason for considering a single axis tracked TLC version with internal micro-tracking is that two-axis trackers are expensive. CPV-grade two-axis tracking currently costs EU 72/m^2 (Trackers Feina), which is \$84.4/m^2, or \$80/m^2 (Morgan Solar), which at today's exchange rate are similar prices. At 38.7% efficiency, that is 20.7 ¢/W, or almost as much as the entire TLC DAT module (including mirrors, lenses, cell, substrates and cooling).

In contrast, CPV-grade single-axis Trackers Feina is only EU 31/m^2 for, and that includes engineering costs so two-axis tracking is ~\$50/m^2 more than CPV-grade single-axis tracking. At 40% efficiency that is a ~12 ¢/W savings in tracking cost.

However single-axis tracking has ~8 ¢/W in extra costs from micro-tracking and more complex housing, assembly and cooling, and loses 5% in efficiency from the cover glass, and 6% to 12% on capacity factor, so the complexity is greater, and the COE is higher, so the margins are smaller.

- *Future Two-Axis Tracker Costs*

As TLC prices fall with volume, the cost savings of a DAT module versus a SAT module also decrease, so dual-axis tracker prices need to fall to keep the overall cost advantage for DAT.

Can two-axis trackers cost fall roughly by half by 2025 and to 1/3 current cost by 2030? In 2011 two-axis CPV trackers cost \$170/m^2, so costs half fallen in half in the last six years, so it is not unreasonable that cost would fall significantly again.

CPV tracker data from the future is hard to find, but CSP Power towers use similar enough technology and studies have been done there. SunShot has a heliostat goal of <= \$75/m^2 by 2020 (DOE, 2014), with a "stretch goal" of \$60/m^2. Since bare-bones flat solar mirror costs are ~\$10/m^2, that leaves \$65/m^2 as a reasonable 2-axis tracker cost by 2020, and \$50/m^2 as likely by 2025. Heliostats also have more demanding optical tolerances than TLC, so TLC may cost slightly less and should not cost more. Thus, costs probably won't *quite* fall by 50% by the time TLC reaches

a gigawatt (2025 target) but will come close. Further cost reductions are harder due to approaching the cost of steel, so SAT TLC will probably gain relative to DAT at very large scales.

- ***Business Advantages:***

- There are more single-axis tracker companies than for dual-axis,
- Single-axis trackers are produced at a larger scale
- SAT TLC nicely fits the CSP infrastructure, for even better natural business partners.
- While either SAT or DAT TLC can produce heat by-product, it would remove a few cents/Watt of DAT's cost advantage from cooling

- ***Latitude:***

- SAT TLC improves in the tropics. Not only is the micro-tracking range reduced, allowing higher acceptance angle or higher concentration, but the impact of the second axis on capacity factor is reduced as well. (But SAT starts losing light at around 40º and it gets difficult to reach 1500X at around 50º).
- DAT TLC has an asymmetric acceptance angle to avoid a tropical tracking problem. The azimuth drive of the most common type of dual-axis tracker (altitude/azimuth) needs to spin fast in the tropics when the sun passes nearly directly overhead, and the asymmetric AA keeps the speed within the range of common trackers.

- ***Future Possible Advantages***

- AR Coatings are improving, and improved AR coating could greatly reduce DAT's 5% efficiency advantage from not needing a cover glass.
- On the other hand, something like higher concentration cells could strongly shift the balance towards the higher concentration affordable by the simpler optics of DAT.

- ***ATLC DAT vs SAT Comparison Table***

The balance is close enough and contains enough factors that a comparison table is needed to keep track of shifts in the balance:

Criterion	TLC for DAT	TLC for SAT	Notes
Cost 2nd 10 MW	20 ¢/W	29 ¢/W	Either easily beats silicon
LCOE 2nd 10 MW	Even lower	Very low	Either easily beats silicon
Margin 10 MW	Even higher	Very high	Can undercut silicon
Cost 1 GW	16 ¢/W	22 ¢/W	Either easily beats silicon
LCOE 1 GW	Even lower	Very low	Either easily beats silicon
Margin 1 GW	Even higher	Very high	Can fund growth even if silicon given away for free.
Tracker Cost	Fairly high	Low	Very strongly favors SAT
Tracker Accepted	Some (Heliostats)	High	SAT is for PV at utility scale, while dual-axis is largely CSP.
Typical tracker works in tropics	Yes	Yes	Most dual-axis is Alt-Azi, but tolerance sufficient for 90º sun
Module works at all latitudes	Yes	Yes, but not as simply	Single axis iffy beyond ~50º latitude, but can go vertical
Micro-tracking	None	Low	Strongly favors DAT
Capacity factor	High	Fairly high	Higher saves on inverter
Partial shading tolerance	Good	Near-Perfect	Can match strings by position on tracker on DAT
Efficiency	Very High	High	Should eliminate cover glass
Current Cell Cost	Low	Low	Slight edge to DAT: cells can be lower aspect ratio
Max Concentration	Very high	Fairly High	Hard to reach the cell limit with SAT.
Acceptance Angle	Very good	Not bad	DAT is much more generous
Heat Absorption Cost	Very Low	Very Low	DAT has more flexibility
Heat Absorption Simplicity	Fairly Low	Fairly High	DAT has no moving module parts
Heat Rejection	Very Low	Low	DAT has more flexibility
Primary Optics	Very low	Very low	DAT scales down better
Secondary Optics	Very low	Low	DAT can reduce by 50%
Receiver wiring	Extremely Low	Low	DAT has more flexibility
Isolation Cost	Low	Low	No significant difference
Module shell	Very Low	Low	Only around heat pipe on DAT
Receiver assembly	Simpler – no micro-tracking	Reasonable	Micro-tracking involves bigger parts, more alignment
Module Assembly	Simpler – no micro-tracking	Reasonable	DAT details have now been worked out.

Shipping cost	Even lower	Very low	Either easily beats silicon
Installation	Even simpler	Fairly easy	

***Table B.1**: Comparison table of DAT vs SAT TLC.*

Since TLC on SAT is still a viable alternative (especially with the new optical micro-tracking), the differences of SAT TLC modules from DAT TLC modules are covered in the remainder of this section. It is fairly likely that SAT could win in the tropics (where lower tilt improves gains acceptance angle, concentration and capacity factor) while DAT almost certainly wins in temperate latitudes (unless there is a use for the heat byproduct).

B.3 TLC Micro-Tracking

The biggest difference in a TLC module for single-axis trackers is the internal micro-tracking. TLC's micro-tracking rotates the receivers within the module so that their lenses always face the polar/equatorial tilt of the sun's position, and this requires the receivers to not collide when rotated, which drives the lens design, and requires the coolant connections to allow rotation and to operate over a 0º-180º range, which makes long-enough heat pipes difficult, as well as requiring drive mechanics and a cover glass to protect the moving parts.

- **Number of Lens Rows and Lens Focal Length:**
  - *Why Not a Single Lens per Receiver?*

A single simple lens of glass of an index of refraction of ~1.5 would have a focal length longer than its width, and hence outside of the tropics it would hit a neighboring lens during micro-tracking. While a two-part lens would eliminate the physical collision, the focus would still start to be intercepted by the neighboring lens before getting to the cells. (**This does not occur below ~15º latitude,** or on a trough slanted to within ~15º of at-latitude, so a system optimized for the equator or a near-latitude trough would have much more freedom on the lens.)

The first solution to this was a single compound lens, in which the foci from five lens arcs were combined by TIR facets to produce a short focal length, allowing full receiver rotation to any micro-tracking tilt. However, the lens needed higher-quality glass, it was not easily roll-formable as a

large sheet, and it needed a narrower rim angle for a given concentration, which would have either custom mirrors or lower concentration. Recent drops in RP3 mirror cost removed most of the cost advantage of the novel mirror shaping, which swung the balance to RP3 and lens tiles.

- Multi-row Lenses:

Multi-row lenses produce shorter focal lengths by simple scaling of the individual lens. While the rows of cells are then spread out across a wider region, the shallower lens has a diagonal shorter than its width, so the lenses clear their neighbors during micro-tracking.

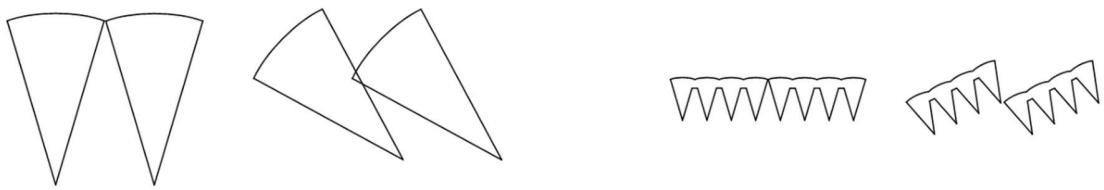

Figure B.1: *The need for multi-row lenses. Individual simple lenses collide (left), Multi-row lens tiles don't collide (right).*

While narrower lenses to fit in more rows increases the clearance *as a percentage of lens width*, this quickly reaches diminishing returns, after which the absolute clearance in millimeters begins to decrease even as the relative clearance continues to increase. For a given receiver length, the number of rows for maximum clearance varies with lens-width-to-focal length and the thickness of the cells and cell cooling below the focus.

The clearance does monotonically increase with the number of rows, but this quickly reaches diminishing returns, and for a given trough-width longer receivers start diminishing concentration at high tilts because the vertical extent of the receiver means that it cannot all be at the height of the focus of the trough. The absolute clearance increases linearly with the scale of the trough (and thus the receiver length allowable for a given concentration).

- Optimization of Lens Rows and Focal Length:

The lens optimization is an extremely complex trade-off, but it mostly comes down to concentration and cell size. As discussed above, too long a receiver requires a wider receiver at the maximum tilt (which is latitude dependent), and a wider receiver reduces concentration by

requiring more cell area. But a shorter receiver requires a thinner lens and thus a shorter focal length to be able to rotate without collision, and a shorter focal length requires either narrower lenses or lower concentration, and narrower lenses require either narrower cells or lower concentration.

Multiple rows produces an easy-to-align lens tile because it rests stably when flat on a flat surface such as the cells on the cold plate, ensure near-perfect lens-to-cell height and angular alignment, and also rests flat on the opposite face when aligning the receivers to the micro-tracking, ensuring near-perfect initial receiver angular alignment.

- Latitude 'Tilt' Effects

At higher latitudes the sun slants more relative to the trough, especially at noon in the winter. This stretches the width of the focus since the lenses are farther from area of the trough that is reflecting onto them. This effect increases with increasing focal length (decreasing rim angle) because the width of the trough (which does not change with tilt) becomes less relevant.

When the receiver tilts to match the sun, some parts of the receiver must be before or after center of the trough's focus. This effect decreases with decreasing trough rim angle because the light spreads less quickly. This effect is, however, proportional to the length of the receiver.

When a receiver's lens is horizontal (zero tilt) it sees the sun's image in the trough projected as a straight line. However, as a receiver tilts toward the sun's image in the trough, the sun's image is no longer straight below the receiver so the band of light that it 'sees' looks like a smile. This broadens the angular range of the image by an amount roughly equal to the sine of the tilt (the depth of the smile) times the cosine of the tilt (due to the smile being farther away from the receiver). This effect depends on the depth of the mirror's curve, which scales with the *square* of the rim angle.

The length and width of the receiver thus limit how far from the equator a given TLC system can still be and capture essentially all of the light. This is directly related to the concentration – A system good from the equator to 37º latitude should reach 1000X concentration with a reasonable tolerance budget, but by 50º latitude it becomes hard to exceed ~600X.

- Lens Defocused from Trough in at Low Tilt:

While optimizing a lens for low tilt is easier than for high tilt because the concentration is already higher, a lens optimized for high tilt would not work well if centered in the focus at low tilt. First, the lenses would see the light at a wider angle from being closer to the 'illuminated' band on the trough, so they would focus less well and some light would be lost on the second axis. And second, the light would fall only in the middle of the lens at higher concentration, where it would become focused too intensely for standard cells (even with some light being lost).

By moving the lens farther from the focus when at low tilt, the focus can be spread back out to 30X across the whole lens tile, and the lens curvature optimized for the broader angular range of the low-tilt focus. While this pays the concentration cost of both the broader angular range at low tilt and the broader focal width at high tilt, it prevents the overheating and loss of light.

Thus, micro-tracking the lens to high tilt not only rotates the lenses to face the sun, but also rotates them upwards into the center of the focus so that the lens tile is centered in the focus at noon in the winter when the sun is at its greatest tilt. (The refractive index of the lens lengthens the convergence of the trough's focus, so the ideal distance is a few millimeters closer than the lens tile being exactly centered.) This is discussed further under 'Micro-Tracking'.

- Number of Lens Rows Needed:

With the heat spreader and manifold thicknesses added to the minimum lens thickness, three lens rows on a receiver is the minimum number that can allow micro-tracking, and even that requires a trick to get room for sealing the perimeter and still has the cells too close to the edge for great cooling, and a lens focal length shorter than optimal. Going to four lens rows relaxes the constraints enough that it was thought to be optimal.

But as the advantages of small cells have been found to outweigh the disadvantages, the cells have gotten narrower, and this allows more lens rows to further relax the focal length (relative to the lens width) and to provide more room for the cooling manifold and the receiver sealing. **Thus, a five-row receiver appears optimal if an odd number is advantageous, and a six-row receiver if an even number is advantageous.**

- Optimizing the Receiver Width

Once the number of lens rows is picked, the optimal receiver length can be determined. Longer receivers reduce handling costs, and allow wider lenses and thus wider cells, but at even moderate

latitudes a soft limit is hit due to the high tilt of the receiver in winter. The trough's focus is already broadest at noon on the winter solstice, and the tilt of the receiver means that parts of the receiver must be either above or below the tightest focus and the longer the receiver the more this broadens the focus.

Part of the lens being in a broader region of the focus increases the length of the lens needed to capture it, thus increasing the length of the cell rows and thus lowering the overall concentration. One could simply forego the edges of the focus at the highest tilt, but this decrease energy capture (this trade-off is discussed under 'latitude effects' and under 'ideal lens height with tilt').

However, since the since the sun's image as magnified by the trough is narrower in angle at high tilts (due to being farther away along the trough), and since only part of the receiver is illuminated at high tilt (due to being partly shaded by a neighboring receiver), and since the illuminated receiver can be centered so that only half of even that illuminated area is above (and half below) the center of the focus), the effect is manageable up to fairly long receivers.

For example, for 60º tilt (the maximum at 37º latitude), only cos(60º)= L/2 of the receiver is illuminated and the vertical extent of that is only sin(60º)=0.866 times that, or 0.433L (Lengths), but with half above and half below the focus, the maximum distance out-of-focus is only 0.217L. Furthermore, the sun's image on the trough is reduced by 1.87x from its 53.1 degrees (a rim angle on each end), or to 28.4º, the focal broadening is only sin(28.4) = 0.476 times that, or 0.103L.

The *broadening* is actually larger at lower tilts, but the trough's focus *starts tighter* at lower tilts so for reasonable lengths the widest total is at the highest tilt. So, for reasonable receiver lengths each extra centimeter of length only adds just a trace over a millimeter to the receiver width needed, and since the focus is already ~60 mm, each extra centimeter of lens length costs only about 1.7% reduction in concentration.

The limit of current cells is about 1500x (a very soft limit - other designs often have hot spots that can exceed this by several times, but efficiency suffers), and TLC with an RP3 trough could reach this with tight tolerances at receiver widths of up to 100 mm if one pushes the cells smaller, and around 80 mm with the targeted cell size.

- The Inverter Effect:

With the module length picked at 4060 mm to match solar thermal troughs, the number of receivers per module depends on the receiver length. Since the receivers are in series, the voltage per module thus depends on the receiver length. Since large inverters come in 600V, 1000V, 1200V and 1500V (with 1500V expected to become prevalent at utility scale), having integral numbers of modules add good matches for (meaning a bit below) several of these voltages is good. At 1500V inverter should have an MP-tracking range of 000V to 1300V, so four 300V modules in series is near-ideal.

50 mm to 80 mm receiver lengths means 80 to 50 receivers, or 80 to 50 times the cell voltage. Ideally the cells' cumulative V_{OC} would fit within the inverter voltage, but most modern inverters forgive a higher V_{OC} as long as V_{MPP} is within range. The best 3J cells would have a V_{OC} of 3.15V at 1500x and V_{MPP} of ~2.85 at the target cell size, so optimizing for 3V is a good start. At 3V 50 mm receivers gives 240V, which fits tolerable in 600V and very well for 1000V 1300V and 1500V, so 50 mm receivers were originally used. But 50 mm even with five cell rows is ~400µ-wide cells, which is smaller than desirable. 60 mm with five rows works better, 500µ cells and 200V output, and 66.7 mm with six rows has cells a bit narrow but a convenient 60 receivers and 1 mm gap, while with five rows the cells are near ideal, and 180V is a reasonable match for each inverter. 80 mm is a great match for 6 rows, with 13.33 mm lenses and ~560µ-wide cells.

Looking at the upcoming 5J cells, the V_{OC} at 1500x would be close to 6V, so 50 receivers of 80 mm cells would produce a very convenient 300V with 6 rows. Since this is a convenient with for both, this will be used in further calculations. If an odd number of lens rows turns out to be desirable, 64-mm receiver length works well at 375V for both 1200V and 1500V inverters (the growth area), with a similar 12.8 mm lens width and 550µ-wide cells.

- Sealing the edge of the receiver

At the current size, a plain 6-row lens only has ~3.5 mm clearance beside the edge notch. This is enough to have the glass come down to the AlSiC for a seal, or for a metal sealing ring. On the back there is only ~3.2 mm beyond the cell; enough for the manifold to be beyond the cell, but not by much which slightly impairs cooling of the edge rows.

But this can be easily improved: It is the lens row on the edge of a receiver that controls how far in from the lens edge the edge cells are, so the middle rows are kept fully optimized while the edge rows are modified to relax the sealing tolerance and get better edge-cell cooling.

Using a lens curve trimmed off-center shifts the focus from and also lower the edge of the lens. As shown below, the edge lens curve is first slid away from the edge by up ~10% of the lens width, and down enough to avoid a jump in the lens surface. The lens is slightly closer, so the focal length is slightly shortened, but this is minor since a 10% shift is only a ~3% decrease in focal length for an optimum lens.

The new lens curve is then trimmed by that amount of the shift where it passes under the next lens and extended to the edge by that amount in the other direction, thus keeping the lens width constant. This also lowers the lens edge, further shortening the lens diagonal. The TIR cones and the bottom of the lens are then be rotated to point at the new lens curve. While the outer edges of the edge lenses are now steep and have slightly higher reflective losses, this is mostly offset by the flatter slopes at their inner edges, and for the polar edge of the lens tile this outer edge is shaded by a neighbor at moderate to high tilts anyway. The outer lens concentration is slightly decreased, but this is very minor for small shifts.

Figure B.2: *Lens Curve Shifted for Clearance.*

A 1.33 mm (10% shift lower the lens by 0.6 mm and the extension lowers the edge a further 0.8 mm. This 1.4 mm lowering gains an additional 0.49 mm on top of the 1.33 mm shift, for a total gain of 1.82 mm. The notch rotates by just over 3º, so it still has a draft angle of over 3º even on the worse side; the rotation claws back 0.07 mm of the gain for the lens (but not the manifold).

This boosts the available width from a skimpy 3.5 mm to a comfortable 5.25 mm for the lens, and from 3.2 to 5.05 mm for the manifold, allowing a good seal with almost unimpaired cooling. Allowing a more comfortable 1 mm gap, these are 4.5 mm for the lens and 4.3 for the manifold.

B.4 Micro-tracking Goals

The main goal of the micro-tracking is to tilt the receivers to match the tilt of the sun's position relative to an E/W line. However as discussed under 'Latitude Effects' in the optics section, the slant of the sun not only affects the direction of incoming light, but also the width of the trough's focus because the longer path from the mirrors to the receivers allows the sun's image to spread further, and even the angular range that any given are of lens sees. The micro-tracking therefore also moves the receivers into and out of the maximum focal concentration as it tilts the receivers.

- Preventing Excessive Concentration

The width of the focus of a parabolic trough depends on the sun's tilt relative to the trough's axis, and at higher latitudes this effect can be significant, especially for troughs with long focal lengths (at 37ºN latitude the focus is almost twice as wide at noon in December than it is in June).

Outside of the tropics this produces much higher maximum concentration in summer than at noon in the winter, and while this could be handled simply by having long enough lenses and enough area to handle the broadest focus, the highest concentration may exceed the limits of the cells unless to overall concentration is lowered, and even if the cells can handle the highest concentration, it is harder to have cells be near-optimally efficient across a broad concentration range.

To avoid excessive concentration, at low tilts the micro-tracking can lower the receivers (or the whole module) below the height of maximum concentration. This broadens the focus at low tilts so that it looks much like the broader focus at high tilts.

- Very High Tilts

At very high tilts much of the lens is shaded, so the ideal height changes significantly. Since the highest tilt is also the broadest focus from the trough, the illuminated part at the lower end of the receiver should be centered in the trough's focus. This then determines how wide the receiver lens tile must be.

- Amount to Raise and Lower Receiver

The amount to raise and lower the receiver depends not only on the size of the receiver, but also on the latitude and the trough rim angle. The initial design is calculated for latitude 37º and a single-RP-3-inner-mirror trough rim angle, as is shown visually in the following diagram.

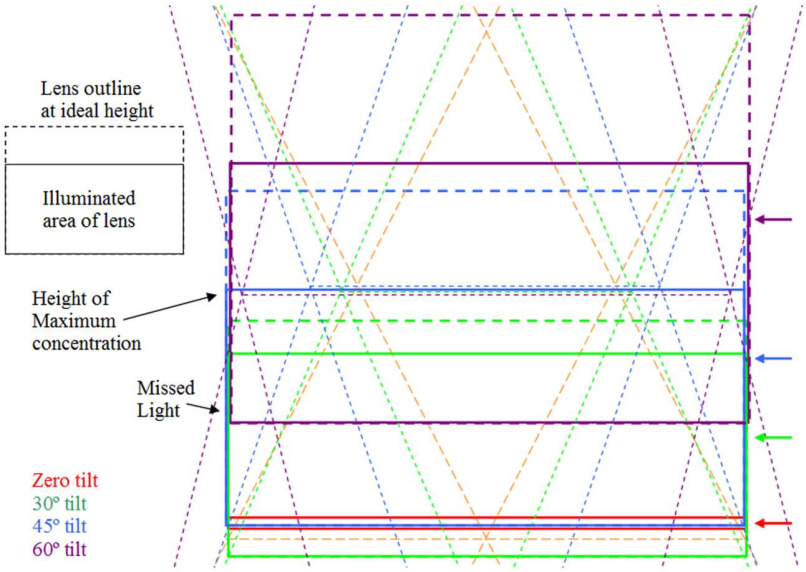

Figure B.3: *Ideal Lens Height with Tilt.*

As can be seen, an optimized lens for up to 37º latitude captures almost all of the focus even at noon on the winter solstice when the sun's tilt is highest (the solid purple rectangle almost gets the dashed purple lines throughout its height). The tiny bit of missed light is not significant – the dashed focal lines are where the focus from an edge-of-spec mirror installed at the limits of tolerance focuses the very edge of the solar disk, and if that tiny bit of light is missed only near noon near the winter solstice, the loss of energy per year is insignificant.

The heights of the center of the lens tile's face are shown by the colored arrows on the right (with an extra gold arrow added for 15º) can also be seen. As can be seen, the lens face starts out flat and low (red box, narrow because seen edge on), and the center rises as the lens tilts. At 60º (noon on the winter solstice at latitude 37º) the lens is almost face on (large purple dashed rectangle), but only the bottom part (solid purple rectangle) is illuminated (the top is shaded by the next lens in the module).

While the optimal rate of change is symmetric with direction, the maximum tilt is not symmetric, reaching ~60º at winter solstice noon and only ~30º at summer solstice dawn/dusk at latitude 37º.

- ***<u>Axis of rotation for micro-Tracking</u>***

  - *Height-Adjustment-Centered Micro-Tracking Axis*

If the micro-tracking axis is significantly above the receiver's lens, then micro-tracking a receiver by rotating it around that axis not only tilts the receiver to match the sun's tilt, but it also raises the lens away from the trough. If the lens is before the focus at zero tilt, this raises the center of the lens toward the focus, which is the direction required.

If the height of the micro-tracking axis is picked correctly, then the center of the illuminated region of the lens is at the center of the focus at maximum tilt, and the height of the lens at zero tilt spreads the light to the same width as at maximum tilt. The curves match fairly well, and by adjusting the height to improve the worst-case tilt (broadest focus), an optimum can be found for any latitude and lens width versus trough rim angle and focal length.

- Off-center Axis:

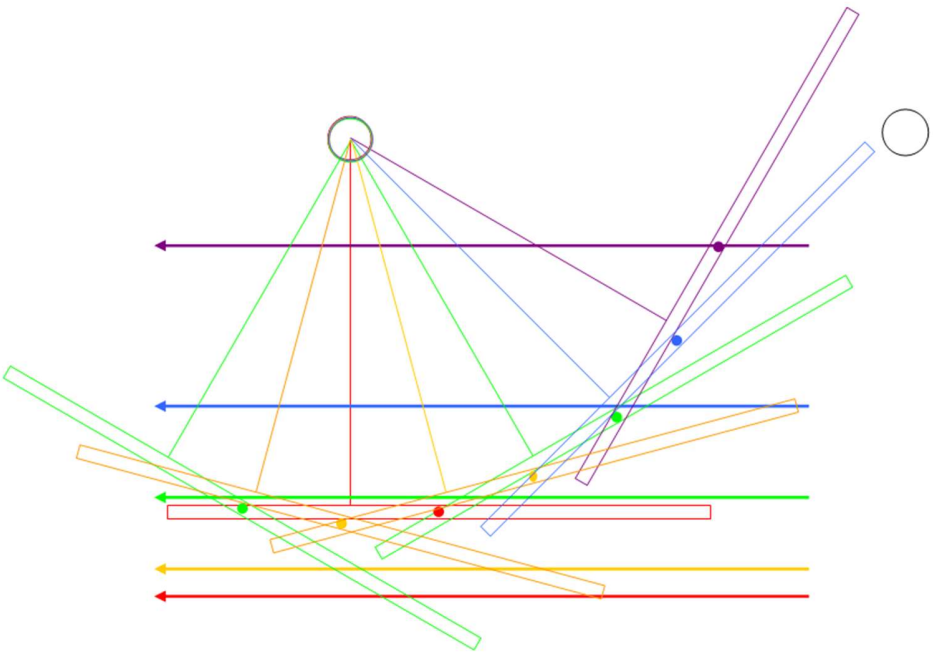

Figure B.4: *Off-Center Axis of Rotation.*

A further improvement can be made with an off-center axis of rotation. While the ideal 'lift' curve is symmetrical, there is major asymmetry in the lifting. At latitude 37º the maximum tilt toward the pole is only ~30º and it occurs at dawn and dusk when the sun is so low that there is little energy to capture, and the sun's path takes it due east/west before it is high in the sky. In contrast the maximum equatorial tilt is 60º and occurs at noon when the sun is much higher (providing light equivalent to 50% of equatorial noon), and the sun stays nearly maximally south for several hours. An off-center axis of rotation can thus favor the equatorial tilt over the poleward tilt for better overall performance.

A trough micro-tracking range optimized this way performs very well up to the latitude that produces the maximum tilt it was optimized for. It will still work fairly well at higher latitudes but will start missing a little light around noon around the winter solstice when the tilt exceeds its optimization. It will also work very well at more equatorial locations, and if the concentration already reaches the cell's optimum, then there is no need for an 'equatorial version'.

As can be seen, the centers of the receivers (colored dots) only crudely follow the ideal height. This approximation, which uses an axle 33.5 mm above the lens surface for the 50-mm long receiver for an RP-3-inner-mirror half-trough, is thought to be near optimal for a through-axle where one receiver's lens has to miss the axle of the next receiver and no cams for height adjustment are used. A better approximation can be obtained with a longer vertical arm for the receiver to pivot on if the axle is two short pieces that allow the receiver to pass between them (a much easier optimization). However, this approximation is good enough, as seen below:

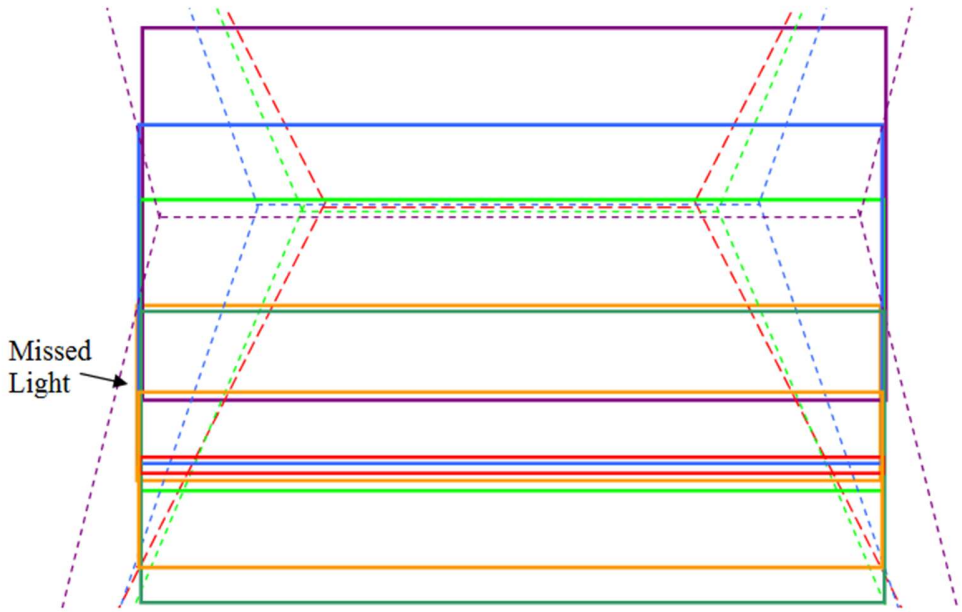

Figure B.5: *Off-center Axis Result.*

Although the ideal height of the center of the receiver was only crudely approximated, no additional focus is lost, even on the summer-solstice mornings, and the worst over-concentration is avoided. This allows even the simplest arrangement, which is to use the coolant inlet and outlet tubes at the axle, while keeping the larger coolant pipes static within the module.

- **Optimizing for different Latitudes**

TLC systems could certainly be optimized for different latitude ranges, but one should ask "Why bother?". Consider a system optimized for 37º (which includes almost all major solar markets and high-DNI regions) as illustrated above. If this TLC system is installed at higher latitudes, the 'loss' initially is only where the very edge of the sun hits a mirror at the edge of its specification mounted at the edge of its tolerance, as well as being for a short time of a day for a short period around the winter solstice. At 45º the system will start losing a few percent of the light at mid-day about six weeks either side of the winter solstice, reaching 12% on solstice noon, but the energy lost per year is less than 1%. And the only significant solar market higher than 45º is Germany, which is also the only significant solar market with sun quality too low to be good for TLC anyway. Thus, this is put off for the far future if needed at all

A system tailored for the tropics, or the same system could concentrate to ~40% higher concentration. But cells lose efficiency at higher concentration, and TLC for 37º can already exceed the maximum that today's cells can handle, so with current cells there is nothing to gain.

If this changes, it is trivial to use the same principles to optimize for any latitude. If, for example, Alberta becomes a huge market at 50ºN, a longer-lensed 1200x TLC system optimized for 45ºN would capture ~99% of the annual energy in the high-sun areas of Alberta, or if cells become available that are efficient at 2000x, then a tropical version could save a cent or two per Watt. But a system optimized for 37º fits all suitable major markets, so outside of this one section, this document looks almost entirely at a system optimized for the 0º-37º range.

- Specific Latitudes of Interest:
  - Latitude for U.S. DOE:

The U.S. D.O.E. uses Kansas City Missouri as an example of where they would like to see solar reach 3¢/kWh by 2030, and at 39.1ºN it is just outside the 37º optimal range. A quick check shows that to catch that light the focus would be 6% wider, and that thus the lens and AlSiC would be as well at a cost of 6% of 4¢/W, but only the bottom half of the cells would be needed in the extra area for 3% of 9¢/W, so the cost for covering a full 39º latitude would be an extra 0.51¢/W or 1.5%. But **not** trying to catch those last rays would only cost 0.15% of the power.

  - Latitude for Sherbrooke:

Sherbrooke is 45.4ºN, which would require a 35% wider lens and 17% more cell area to catch the whole sun (from the worst mirrors when maximally miss-tracked) at the suns most southerly at noon on the winter solstice, which would add 3¢/W, or 10% to the cost. But not bothering to optimize for 45º would cost only about 2% of the annual energy output.

  - Latitude 50º:

By latitude 50º (the sunniest parts of Alberta), 70% more lens and 35% more cell area would be needed, which would add 6¢/W. Adding 20% to the cost to catch the last rays when the sun is so low to the south is simply not worth it, so optimizing to 45º and skipping 8% of the energy for a few hours a day near the solstice (when there isn't much energy anyway), or ~1% over the course of a year would be reasonable. Even then a 37º system would be more cost-effective, losing only

~~4% of a year's energy while avoiding the extra cost. It probably isn't even worth a compromise that misses fewer rays unless winter power carries a high premium.

- ***Micro-tracking Mechanics***

  - *Axle:*

A receiver rotates about an axle. The lowest cost is to use the coolant feed tubes as an axle, in grommet seals in the main coolant pipes. Alternatively, the axle can be a steel rod for lower friction and the coolant can be fed in through silicone tubes. The molded BMC housing holds the glass-glass sandwich and connects it to the axle, the coolant and the micro-tracking tracing drive.

For the micro-tracked design, the receiver has a short pipe molded into the housing on each end; these serve as the coolant inlets and outlets for the receiver as well as the axle to rotate around, and a housing molded from BMC (a sturdy low-cost plastic used around car engines) holds the receiver, and coolant (basically car radiator fluid) flows in a thin film between the glass cover and the AlSiC. The AlSiC spreads heat from the narrow cells so this flowing film cools as well as a micro-channel cold plate, but at an order of magnitude less cost.

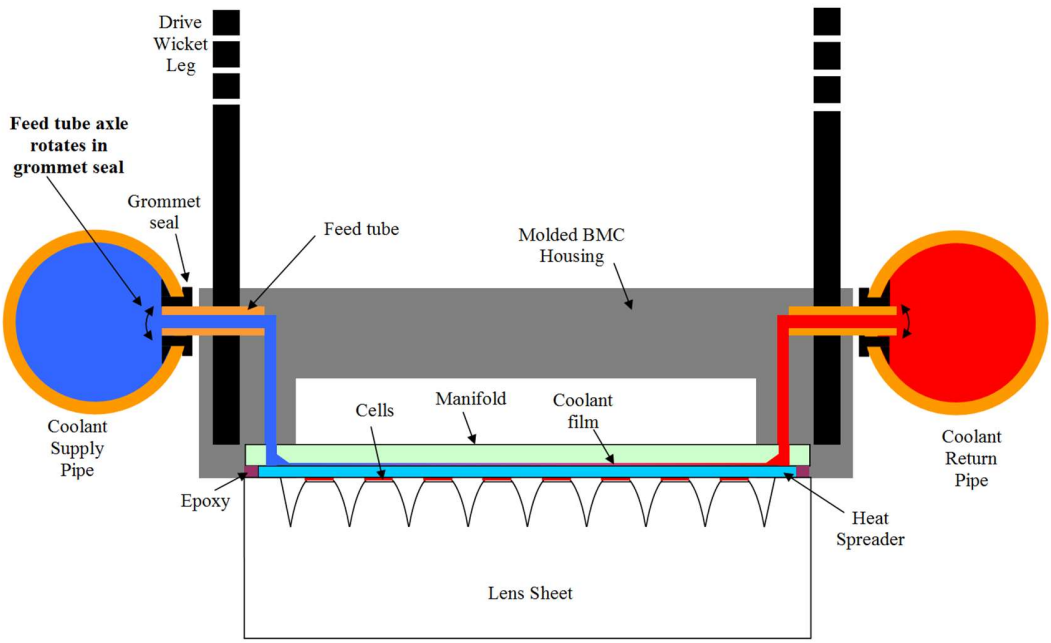

Figure B.6: *Receiver Cross Section with Micro-Tracking (~1.6x scale).*

- Micro-Tracking Wicket

The micro-tracking drive uses a wicket of bent steel rod, providing a pair of drive legs for high stability. The legs are long enough to extend from the axle to the body of the receiver, (although there is really no force after they pass the axle area). This is the axle offset minus the lens depth, or ~15 mm. (If silicone-tube feeders are used instead of cooling tubes in grommet seals, the housing can be two sections over-molded onto a steel axle).

Each leg also has to reach the height of the common drive rod that gangs the receivers together. The length of this section provides the micro-tracking accuracy, and while 100 mm is overkill for accuracy, it does prevent the receiver from hitting the drive rod up to the maximum ~68º tilt at ~45º latitude, allowing a TLC system optimized for equator-to-37º to be used as far north as southern Quebec. The total length of the wicket rod is thus ~310 mm.

- Higher latitudes:

Tracking at bit higher latitude requires a longer wicket and a taller housing; for example, to track to winter solstice noon at latitude 50º require another 32 mm. For a version for tracking at even higher latitudes, the wicket legs can be splayed several centimeters at the top and a drive rod on each side can allow the lens/manifold sandwich to pass between the drive rods and wickets. The limit then become the wicket top hitting the cooling pipe at over 80º tilt (but there isn't that much sun to go after at such high tilts anyway).

- Accuracy:

The torque to turn in each grommet seal is about 1/8 Nm. It doesn't really matter if the wicket legs all flex by the same amount (the micro-tracking could detect this through power output and adjust, but any difference in flex will show up as a micro-tracking error, so to be safe the torque is assumed to range from zero to twice the nominal value, or a range of ¼ Nm.

Each wicket leg is a 100 mm long rod with a force variance of 2.5N, and if the rod is 3/16" in diameter the deflection at the end of the rod is ~6.4 mils or ~164µ. This is 0.00164 time the length of the beam, or an angular deflection of 1.64 mrads, which is more accurate than most two-axis trackers and only 16% of the micro-tracking tolerance allowed in the TLC optics model.

- Wicket Cost:

The rod does not have to be precise because the initial receiver alignment will adjust for any discrepancies from a 'true' rod. Thus, simple drawn or hot-rolled rod will do, and this currently sells for $530/ton, or $585/tonne, or 58.5¢/kg. The mass is 0.566 kg/m, so it costs 33.1 cents per meter and 310 mm costs 10.3¢ or **0.2¢/W**.

Before the rod is bent, a 2-cm spherical BMC bead is slid onto the wire and to the middle. The bead has a second hole perpendicular to the hole through which the wicket rod passes; the drive rod will pass through this second hole later. A bead is ~4 cm^3 of BMC, or ~1.3¢ or ~**0.025¢/W**.

Bending steel rod is well-established process - a low-cost wicket-machine could position a bead, slid a rod, cut the rod, bend the rod and eject the wicket; at an easily achievable 6 seconds per wicket, this is 30 kW per hour, so each machine could do ~50 MW/year per shift. Machine cost is thus tiny, and even for one machine with a full-time attendant, the cost is only **0.1¢/W**. A person could manage several machines, so **0.05¢/W** is allocated for volume production.

- Receiver Housing

The Receiver housing will be injection-molded from Bulk-Molding Compound, or BMC, a sturdy, low-cost, coolant-compatible plastic used for car engine peripherals. It will have two channels molded in to carry coolant to the glass manifold, to which it can be attached with epoxy, and the housing will be over-molded onto the wicket legs and onto two short lengths of tube that serve as axle and also convey coolant to and from the channels.

- Housing Accuracy:

The receiver optics are sensitive to lens slant along the trough and slightly sensitive to slant across the trough, but both of these are locked in on initial alignment, so housing accuracy does not affect them. The receivers are quite insensitive to other alignments such as the collinearity of the axles or even if one main pipe is shifted relative to the other (thermal expansion), so the housing does not need high accuracy. Over-molding onto the wicket legs and axle tubes should supply very high accuracy for free, so no error budget is consumed by the housing. The BMC can be mirror-finished to reflect heat in the event that the tracker is misaligned.

- Housing Cost:

For prototyping the housing can be made from a 1" x 1" x 3" block of BMC and a drill press. BMC costs ~5¢ per cubic inch, so even that is only 15¢ of plastic. For production, injection molding can reduce the plastic used to about 18 cm^3 of plastic, or about 6¢. Molding cost should be low, so 8¢ is allocated in moderate volume and 4¢ in high volume. That makes cost 14¢, or **0.27¢/W** for introductory volume, and 10¢, or **0.2¢/W** at a gigawatt scale. The two 2-cm lengths of copper axle tube at 60¢/foot come to 8¢ or **0.16¢/W.**

B.5 Micro-tracked Receiver Cooling

With 50% of the optical energy of the intense focus being converted to heat, the cells need good cooling. There are two basic receiver cooling challenges: getting heat from the cells into coolant and getting the coolant into and out of the receivers. The cooling of the cells has gone through several different designs and the current design is thought to be near optimal, so the older designs are listed as obsolete.

The current TLC cell cooling system is very simple – the cells are narrow enough that a heat spreader can spread the heat to enough area that a thin film of water on the back of the heat spreader can rapidly absorb the heat. However, a lot of details go into that simplicity.

- ***Receiver Heat Absorption:***

- *Coolant Flow per Receiver*

There are several drivers of the coolant flow per receiver: cells (heat and temperature sensitivity), pumps, and pipe sizes, and these all need consideration. The film can then be adjusted to get the right pressure drop and cooling characteristics.

- Low-pressure / high-flow Pumps:

Gorman-Rupp (GRI) has low-cost reliable pumps (INTG-3) for low-pressure applications that pump just over 6 GPM, which is ~400 cm3/sec, or 1 cm3/sec per centimeter of a 4060-mm module. There is no point in NOT running the coolant through each block, even when the pump covers multiple blocks, so this is a reasonable starting point for coolant flow rates.

Because many flat-panel trackers are ~3.3 meters wide, two RP-3 inner troughs fit so even though the individual mirror is 1641 mm and the individual receiver is 55 mm, the coolant flow is evaluated for the 3.3-meter width. This also covers a scaled trough expected to be optimal for very high volume production.

- Matching a Standard Pipe:

6+ GPM is a good match for ¾" pipe, with a pressure drop in the pipes of only about 2 kPa (0.3 psi), so this is a good match for low-cost commodity plumbing. If the inlets and outlets are used the micro-tracking axle, then the pipe material should be fairly-well CTE matched to the drive rod (such as galvanized steel + steel), but otherwise the pipe can be any material although plastic near focused sunlight makes me nervous. Copper and steel are roughly the same price at $5/meter.

- Temperature Rise in Coolant;

At 3.3 meters wide, each centimeter of trough gets 3.3% of a square meter of sunlight, or 33 Watts, and ~10% is reflected infrared so even with 42% cells efficiency only about 50%, or 16.5 W/cm, is heat. Coolant absorbs ~4.1J/cm^{3}K, so a cm^3/sec of coolant per cm of trough absorbs 4J/cmKsec = 4W/cmK, so 16.5 W/cm means only a 4K rise in the coolant temperature.

The 4K rise costs only 4*0.06%abs = 0.24%abs = 0.6% of the cell's output, so this is a reasonable coolant flow rate (in addition to other thermal losses). A ~$100 pump can feed coolant to about 20 kW, so that's ~0.5¢/W, and doubling the number of pumps would only save 2K, or 0.3% of the output, and the value of that is only roughly half of the 0.5¢/W so it is simply not worth going for higher coolant flow to reduce the temperature further.

- The Flowing Film

The thinner the film, the better the heat transfer, so the film must be thin. But a thin film is in laminar flow, and in laminar flow pressure drop rises dramatically with thinness. Once glycol antifreeze is added the pressure drop becomes the limit for a pump serving more than one module unless the film is too thick to cool well, so a simple manifold becomes optimal.

- Simple Manifold:

Most manifolds have numerous branches to get very short flows, but just a few branches is sufficient to greatly reduced the pressure drop due to the number of branches dividing both the flow speed and the flow distance in the thin film. And is help not to go too low – if most of the

pressure drop is in the film, then the film flow tends to be quite uniform even if the manifold is stamped rather than machined.

A suitable material to CTE-match to the lens is glass, so the manifold is made from roll-formed glass. Ordinary soda-lime glass is fine because since this is not an optical element but a mechanical/fluidic element. The manifold only needs to be about 3 mm thick (the deepest channels are ~1 mm deep), so pricing from textured 3.2 mm solar cover glass is use (pessimistic); this costs \$4/m^2 or 2.2¢ for the size of the AlSiC plate, and that is only 0.04¢/W!

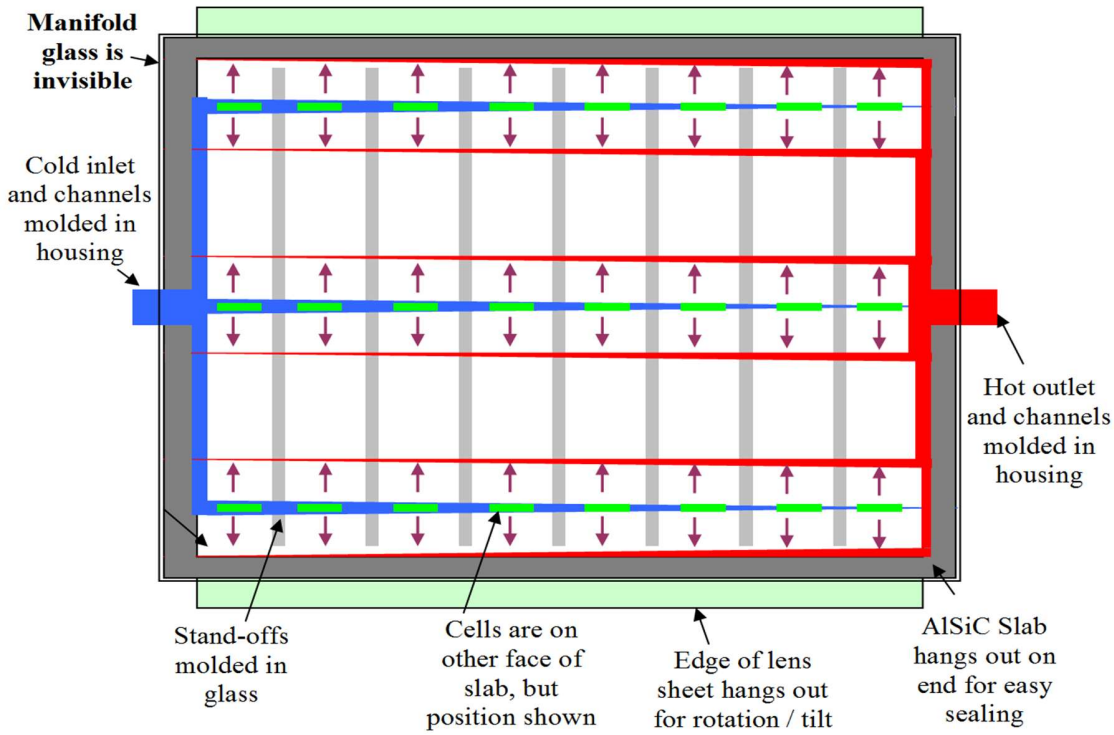

Figure B.7: *Coolant flow in receiver, only 3 rows shown.*

In the above illustration, the glass manifold is almost invisible and the coolant flowing within it and the surfaces below it can be seen. The cold coolant (blue) enters on one side and is divided into two tapered channels. The coolant flows in each direction from each of these channels to get to the nearest hot channels, which taper in the opposite direction (note that the cell array is currently 6 x 20 rather than the 3 x 8 array illustrated, but the principles are the same).

With the back of the AlSiC insulated by the LCP, coolant can squeeze between the glass and the LCP (which is highly chemically resistant). This forces the coolant to flow in a thin film from cold channel to hot channel (purple arrows). The glass manifold is epoxied to the LCP along every stand-off, and everything in this glass-AlSiC-glass sandwich is well CTE-matched so it is strong and extremely stable. And it will look cool!

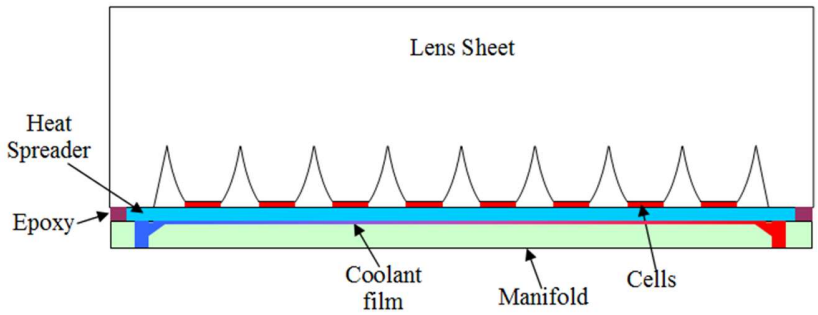

Figure B.8: *Manifold/Lens Glass-Glass Sandwich*

- Heat Transfer into Coolant:

The conduction of heat into a flowing film improves linearly with decreasing film thickness, and the 60μ film thickness set by the stand-offs is a moderately aggressive (thickness in recently-used micro-channel cold plates range from 50 μ to 160 μ). However, the heat spreader area is large enough to make up for this, even with only one transfer surface - there is ~10 cm^2 of exposed LCP near the cells and only ~40 W of heat to conduct, or only 4 W/cm^2 or .04 W/mm^2.

The thermal conductivity of the coolant is generally between 0.4 and 0.5 W/mK, depending on the glycol (antifreeze) concentration and the temperature (improving with increasing temperature), so 0.4 W/mK, or 0.0004 W/mmK, is assumed (pessimistic). The heat is conducted to the center of the channel on average, or only ~30 microns or 0.03 mm, so even at 0.0004 W/mmK, the 0.04 W/mm^2, and a conduction distance of only 0.03 mm, the temperature rise from conduction into the coolant is only 3K.

- Heat Conduction through the LCP Isolation:

If the AlSiC spreads the heat to an area of ~10 cm^2, then in normal operation there is ~7 W/cm^2 of heat to take through the LCP. 10 W/mK is 0.1 W/cmK, which is 60 degrees per centimeter; 100 μ is 1/100 cm, so that's only 0.7 degrees penalty for the LCP!

- Conduction in the Spreader:

However, to get down to 3K during transfer into the coolant, the heat has to spread to 10 cm^2, or 25 mm^2 for each cell. From a 2.505 mm x 0.546 mm cell, that takes spreading 1.8 mm in each lateral direction, and AlSiC is merely a very good thermal conductor and not a perfect one. This section will be replaced by a COMSOL FEA analysis, but for now a pessimistic bound is calculated by assuming that the heat flow averages halfway through the slab before spreading, spreads an average of halfway the distance, and then flows an average of half-way vertically.

This is pessimistic because it ignores the corner area, diagonal paths, and that more heat than average will flow on the shorter paths, but it is easy to calculate. The cell sits on a 1-mm-thick slab of with a thermal conductivity of 0.180 W/mmK, so the entire 0.75 W/mm^2 heat being conducted an average of 0.5 mm vertically before spreading costs 2.1K. The hot plug has a circumference of 6.1 mm or a circumferential area of 6.1 mm^2 or 4.46 times bigger so the heat per area is 4.46 times less or 0.85 degrees. The heat flow across the 25 mm^2 is then 0.03 w/mm^2 so flowing 0.5 mm on the average in 0.18 W/mmK AlSiC costs 0.083K.

The total thermal spreading in the AlSiC therefore costs less than 2.1+0.85+0.83 = ~3K. Adding in the 0.7K from the LCP and the 3K from the conduction into the coolant and the 4K rise in the coolant itself, and the total thermal penalty of the cold plate is ~10.7K. For starting from 0.75 W/mm^2 or 75 W/cm^2, this is actually very good performance at ~0.15K/W/cm^2.

- Pressure Drop

The very short distance in the thin film means very low pressure drop, and the total pressure drop in a receiver at operating temperature, even with glycol antifreeze geared for -40ºC, is only about ½ PSI. This will be further improved in the future by tapering the molded glass so that the coolant film is even thinner right under the cell where the heat transfer is most efficient and is slightly thicker away from the cell. The pressure drop in a whole module is only about 1 PSI so a $125

GRI INTG-3 pump (quote and data sheet from GRI) could service 8 modules with full flow, or 10 modules if the glycol is just for -20ºC.

- Cooling the Diode

The diode can take a much higher temperature, but if a lower-cost diode that is more sensitive is used then a short additional cold manifold finger can provide overkill cooling directly under the diode at no extra cost.

- Receiver Housing

The receiver housing is mentioned here because it carries coolant to and from the receiver. While BMC's CTE is about 8 ppm/K higher than the glass/AlSiC/glass sandwich, even with a +/-50ºC CTE temperature swing each end of the sandwich only moves ~12μ relative to the housing. The BMC flexes enough to accommodate this.

B.6 Micro-tracked Receiver Interconnections

Micro-tracking within the module adds complexity longer insulated wires to allow for movement, and connectors to allow receivers to be wired together late in the module assembly.

- ***Terminal Wire***

The SAT-receiver terminal wire is fine-stranded (65/30) 2 kV solar connection wire, with 12 gauge currently being the most cost-effective (12 ga wire is overkill, but in 2kV solar wire it costs almost exactly the same as 14 ga). Although the distance between connection points remains fixed at the receiver spacing of ~50 mm, about 100 mm of wire is needed to minimize flexing stress as the receivers micro-track (with the fine 0.25 mm strands this is pessimistic).

- ***Connector for the Terminal Wire***

A SAT receiver will have a connector cage soldered to a Ni/Au pad on the AlSiC. The connector will accept the incoming terminal wire from a neighboring receiver and connect it electrically to the AlSiC backplane so that when the terminal wire from one receiver is plugged into the next receiver during module assembly, the two will be electrically in series. The connector will also *mechanically* hold the receiver's own terminal wire by its insulation (depending upon the attachment process, a clip can also hold the far end of the bare section of the receiver's terminal wire as its terminal wire is attached).

B.7 Stripped end of Terminal Wire

The SAT receiver's terminal extends onto the top metal to supplement it with a low-resistance path. This extends the wire by ~50 mm; these costs more than bare wire because the terminal wire has high grade outdoor insulation that is stripped off and thrown away, but 12-gauge 2kV solar wire costs only about 1.4 cents/cm (Mouser catalog, 2016), or 7¢ for the extra 50 mm needed or 0.14¢/W. And this saves a 10¢ connector for the wire to the next receiver (because it is the same wire), and this bare wire can also be soldered (or epoxied) to the top of the diode as it passes over the diode on the AlSiC, eliminating the need for wire-bonds for the diode as well.

- ***Terminal Wire Cost***

The 12 ga wire costs 0.14 cents/mm, or 25.2 cents for 180 mm. For a 52-Watt receiver this is **~0.5¢/W**, *including* the stripped wire on the receiver.

- ***Resistance***

The higher current (~11A for the introductory product with the expected cells) flows through the 130 mm plus the equivalent of an extra ~15 mm on the receiver, so the low resistance of the 12-gauge wire is helpful. At 5.2 µOhms/mm, the 145-mm-equivalent wire is 750 µOhms so at 11.2A the voltage drop is 8.4 mV, which for the 46% cells is a 0.175% loss. Even with the current 3J cells the 17A is 12.8 mV which is only a ~0.45% loss. Adding the 5 mV from the top copper and the 4.8 mV from the bond wires gives a total of 18 mV, or 0.37% electrical loss (which is tiny compared to wiring losses in other HCPV systems). Even with current 3J cells it is only ~23 mV or a loss of ~0.8%.

- Volume-Production Receiver

An optimized SAT trough would probably be twice the mirror width and focal length, which still fits on a flat-panel tracker). The area (and thus current) would be doubled, and only the on-receiver wire length would change. With the same wire, the power loss would increase by ~0.2%, but lower diode, wire and handling costs would more than make up for this, and power transmission could

all connected to a common drive rod. A small stepper motor drives the rod back and forth on a daily basis, rotating all receivers to follow the sun on the second axis. Aligning the receivers initially is easy the linkages are originally free to slide along the drive rod, and the module is placed so still use 8-ga wire at $2.80/meter for wire runs to an inverter.

B.8 Module (4060 mm) that includes Micro-tracking

In the micro-tracked design, the receivers are ganged together for micro-tracking. Each receiver has a drive arm molded into the housing, and these are that all lenses sit flat on a flat surface, and the linkages are then locked to the drive rod with drops of adhesive. Each module will have its own drive motor and fine-tracking sensor; the power output of one or more receivers or of a module can be used as the sensor.

A 4060 mm module matches the length of standard thermal-trough receiver tubes, and the trough will use the inner segments of a standard RP-3 trough, allowing CSP infrastructure to be used. This will produce two 1.641-meter off-axis troughs on a tracker, so each 4.06-meter module will have 6.66 m^2 of aperture and have two rows of receivers (one for each half-trough), sharing some module infrastructure. The 46% cells expected at product introduction will produce modules of ~40% efficiency and 2.6 kW per module.

There will be fifty 80-mm-wide receivers in series, which for 46% cells at a V_{OC} of ~6V produces ~300V, which is suitable 600V, 1000V, 1200V or 1500V inverters by linking modules in series. Each receiver will produce around 50W in full sunlight.

Reusing CSP infrastructure helps greatly on production cost and especially on capital cost, at low volume. All parts are simple and low cost, and at an initial volume of a few tens of MW, the production cost is expected to be around ~30 ¢/W for ~40% efficient modules. While the cheapest silicon panels (made at thousands of times the volume) will by then be approaching that cost, the higher efficiency and durability of TLC add great value, so TLC will be competitive even at low-volume production for product introduction.

- ***Cover Glass***

A cover glass should not be necessary since the modules are basically hermetically sealed glass/glass modules, and the housing could be folded to protect the micro-tracking from dirt and weather. But a cover glass looks right on a 50-year product, and it makes the metal work simpler, so a cover glass is included.

The cover glass is the same as the lens glass, and thus is the same as a high-quality flat panel cover but 25 times smaller per aperture area and thus 40x smaller per Watt. Module covers cost only a few ¢/W, so this costs less than 0.1¢/W even from the best solar glass AR-coated on both faces.

However, while the economic cost is small, even with an AR-coating the glass still reflects ~5% of the light. This raises the cost by 1.4¢/W as product introduction, and it lowers the value by ~3¢/W due to decreasing the overall efficiency. However, if it makes the difference between a 30-year life and a 50-year life, that is worth over 20 ¢/W at product introduction.

It should be noted that if glasshouse solar like GlassPoint proves economical, then one could dispense with both the housing and the cover glass.

- ***Module Shell and Support***

  - *Module Shell*

A module will be housed in a long sheet-metal shell. While the glass lenses of the receivers will be very durable, a cover glass is added where light from the trough enters the module to target a 50-year product life. Glass and metal at the focus of a parabolic trough is still the *only* solar technology field-proven with over 100 MW of utility-scale solar on sun for over 25 years, so a very long product lifetime should be achievable.

The shell is not hermetically sealed (any water vapor that gets inside and condenses will be baked off as soon as it goes on sun) – but it will keep dust and snow/ice off of the micro-tracking mechanics, allowing low-cost light-weight mechanisms to be used. The shell also supports the coolant pipes, and they in turn support the receivers and the micro-tracking linkage.

The shell flexes sufficiently to simply be pulled outward at the bottom enough to go over the coolant pipes, and then springs back. A bead of adhesive is then applied; a previous project got 36 meters of bead from a $6 tube (Dow Corning Silicone 995), so that's $1.50 per receiver module or **0.06¢/W**, and the shell is then pulled outward a bit less and the cover glass is placed. The sides are then held against to cover glass while the adhesive sets.

To accommodate the receiver travel during micro-tracking, if the shell is closed-ended the receivers have to be sized to leave about 2/3 of a receiver in spare space. However, it is more efficient to use open-ended shells where one shell attaches to the next, with the modules at the trough ends capped. To prevent leakage where the cover glasses abut, a bead of clear silicone on one cover-glass edge is compressed when the shells are bolted together.

The cost of the shell is very low since it is made from roll-formed galvanized sheet. Even if a solid 20-gauge steel and 20 cm wide and 15 cm tall with 2.5 cm folded under on each side, that is

only ~2.2 m2 of steel per module. At 8.1 kg/m2, that's ~18 kg of steel sheet, or about $12 at current prices (steel catalog on web) or an insignificant 0.25¢/W.

- ***Module Electrical Output***

While the receivers within a module are interconnected during module manufacturing, the connections from module to module are made in the field. The low current allows standard solar-module connectors to be used to connect the modules in series. With roughly an order of magnitude fewer connectors per Watt than flat panels, standard 1500V solar connectors between modules are essentially free (in terms of both parts and installation labor) – at $3.40/pair, they are 0.13¢/W for introductory modules (Mouser catalog, 2017, quantity 1000), and ~0.06¢ a watt for optimized modules, and judging by price trend, ~0.05¢/W for large-scale production.

- Module at Introduction

With an RP-3 inner segment trough focusing on each side of the module, a module has two rows of receivers. With 80 mm receivers and 0.75 mm of clearance between them, a 4060 mm row of receivers has 50 receivers on each side, or 100 receivers in total.

  - Power with Cells Expected at Introduction

46% cells are expected to be available by product introduction in 2021, and their higher voltage would produce ~7A per row, or 14A per module. 80 receivers would add to ~375V, making four modules in series a perfect match for a 1500V inverter. Each module will be 5.3kW, so each string can feed 20 kW of inverter (strings look like strings of flat panels, and strings can be in parallel for bigger inverters). Just as with flat panels, inverters can be over-provisioned – for 30% overprovisioning one simply uses ~60 strings per MW of inverter, or ten 100-meter troughs.

  - Current cells

Even if all cell progress stopped, the module would produce ~150V at 32 Amps with today's 3J cells. This would use four blocks in series for a 600V inverter, six in series for 1000V, or eight in series for 1200V, or ten in series for 1500V. At 4.8 kW per block, each string is 18.4 kW to 48 kW depending on the inverter.

- Optimized TLC Trough

The optimum width trough (~2x) with the scaled receivers (2x the width to match the trough width) will be ~100W at about 22A with the cells expected at product introduction. 50 of these will fit for block with 1 mm clearance, for 5.3 kW at ~300V V_{OC} and V_{MPP} of ~250V. This makes strings of five a great match for a 1500V inverter, and ten 100-meter trough rows a generous ~30% over provisioning per megawatt of inverter.

- ***<u>Coolant Distribution</u>***

The receivers are fed coolant from a copper pipe and drain into a similar pipe. These pipes connect to a bundle of fin tube that can add to the strength and rigidity of the module as well as providing low-cost heat rejection. The coper pipe sells retail (chain hardware store website catalog) for $3.50/meter for ½" pipe for introductory receivers, and $5/meter for the ¾" pipe for an optimized trough, and two are needed per module, so that's $28 for 2.65 kW or 1.06¢/W (pessimistic since cost would be lower in volume). At a gigawatt the optimized trough would cut this to $40 for 5.3 kW, and a 10% volume discount is assumed for 0.68¢/W.

- Coolant Connections

The receivers' coolant distribution must be connected to the module's overall coolant distribution system in a manner that allows the receivers to rotate during micro-tracking. There are several ways to do this including the following:

  - Feed Tubes in Grommet Seals:

A grommet seal is a low-cost doughnut-shaped seal that provide a fluid-tight seal around allow a smooth tube. For example, some car radiators use coolant headers with rows of grommet seals for both input and output and have smaller finned coolant pipes in the seals to carry coolant from input header to output header. Some grommet seals allow the tube to rotate within the seal, for example, my cappuccino machine has a pivoting "foaming" arm that rotates in a grommet seal. Grommet seals are essentially free at <0.1¢/W (Alibaba – no U.S. pricing found online).

An O-ring is also a form of a grommet seal; in a sink with a pivoting spout the spout typically is in an O-ring or other grommet seal. Grommet seals are typically made from an elastomer such as rubber to compensate for any imperfections in the tube and for effects like thermal expansion, and the tubes are typically made of metal for durability.

A TLC receiver can have its coolant inlet and outlet tubes form its axis of rotation and have these tubes in grommet seals the module's coolant distribution pipes. It is simple and low-cost to drill rows of identical holes in a pipe and to insert grommet seals (as is done in some car radiators). A jig can hold a whole module-worth of receivers at the correct spacing, and the coolant distribution pipes can then be slid onto the aligned tubes. The tubes can be molded as part of the cold plate housing or can be metal (preferably copper to match the cold plates) inserted into the housing. The tubes can have tapered ends to make insertion easier.

The grommet seal in my cappuccino machine requires less than one inch-pound of torque, which is about 1/8 Nm. A set of 50 receivers with 2 seals each would thus require about 12 Nm of torque. If the micro-tracking arms are 10 cm long, that means a force of 120N or 25 lbs.

- Twisting feeders

The receiver's coolant distribution can also be connected to the module's coolant distribution through rotational-axis tubes that twist rather than rotating within a seal. However, such tubes are less preferred than grommet seals because they tend to flex and provide less-precise alignment.

- Bending Tube

It is not necessary to use the coolant feed tubes as the mechanical axle. For example, the mechanical axle can be a rigid metal rod at the axis height, and the feed tubes can be highly flexible tubes that bend to accommodate the receiver's rotation. Small-diameter silicone tube is sufficiently flexible, and it only costs ten cents per meter (un-negotiated quote from insistent Chinese supplier) in quantity making it a preferred embodiment.

The length needed is only a few centimeters more than a loop for flexing, and an 8 cm diameter loop is sufficient, so 35 cm is allocated. Two such looped tubes are needed per receiver if tubes are used, or 70 cm or 7 cents or 0.14¢/W. A compression band (like that used for PEX tubing, but smaller and lighter) would cost ~0.01¢/W each for four, so 0.18¢/W is sufficient for the tube and connections if tube is used.

- Twisting Feeder Boots

The coolant tubes could also be over-molded with an elastomer to produce a 'boot' for each feed tube; these could then be clamped with compression bands around the feed tubes, with enough free portion to twist the ~90º needed during micro-tracking. This could even be done as a backup for

O-rings or grommet seals, providing double protection. However, given that very high-grade grommet-seal materials are affordable, this should not be needed.

- Coolant

Water is the best coolant, with antifreeze added as needed for the installation site's climate. The coolant is electrically isolated, a triazole anti-corrosion agent will provide further insurance against corrosion-produces particulates, and an anti-fowling agent will prevent bacterial growth. Such mixtures are available in industrial quantity with either propylene glycol (non-toxic for prototyping when spills are inevitable) or ethylene glycol (standard automotive antifreeze).

  - Coolant Leaks

Although grommet seals and silicone tubes should both be leak-free, with several hundred connections per module leaks may be a limiting factor as seals age. Propylene glycol evaporates ~0.1% as fast as water and even ethylene glycol and benzo-triazole evaporate ~0.01% as fast as water, so for tiny leaks the coolant would simply bake out of the non-hermetic module. Another possibility it to use an alcohol antifreeze; this evaporates faster than water and hence a larger leak can be tolerated without significantly impacting module performance. Evaporation thus lets a small coolant leak be dealt with as routine maintenance rather than immediate maintenance.

- ***Heat Rejection***

In the micro-tracked version, the coolant for a whole 4-meter module runs in a cold pipe that feeds all receivers in parallel, and the receivers drain into a common hot pipe. Regular plumbing pipes each get a series of holes, and each hole has a grommet seal; these seals hold the feed tubes of the receivers. Spent coolant passes through a set of fin tubes (industrial radiator) under the **next** module. A pump ever four modules gives a ~100-meter trough six pumps driving a coolant loop, with coolant passing through receivers, fin tube, receivers, fin-tube, etc., down the trough and back. This provides excellent cooling, and redundancy in case a pump fails, at low cost.

If low-grade heat is desired as a by-product, then the fin tube is replaced with larger-diameter cold and hot trunks with solar-thermal connections at each end of a trough row (and without the cost of glass receiver tubes or vacuum seals).

- Fin Tube

TLC can either supply low grade heat or can reject it to the air. When rejecting the heat, each radiator section has eight fin tubes, each 5/8"-on-1" with 1/16" wall and 5 FPI. The area of a fin is 3.19 in^2, so with 5 FPI and two faces per fin, eight 12-foot tubes provide 1.78 times the mirror area for a twin-mirror trough. With more energy converted, this is similar heat per area to reject as a well-ventilated flat panel that has two faces to reject from, and the tandem cells are less sensitive to heat, so TLC retains a thermal advantage. And fin tube is affordable – at <$2/ft in 10 MW quantity (Tex-Fin quote, 2017) it is only 3.52 ¢/W. It probably decreases only slightly at 1 GW, but that is already highly affordable (and that is an un-negotiated price, too). Also, 8-fin-per-inch rolled-aluminum fin tube with a 25 mm tube OD and 50 mm fin OD was available from India at $0.80/ft in 10 MW quantity (old 2007 price), and a 5 FPI version would provide comparable fin area at ~5 ¢/W. This has the higher CTE and electrochemical potential of aluminum, but that should not present any problems, however its cost would have to come down.

- Viscosity Effect in Fin Tube

While in the heat absorption turbulence is undesirable because it would increase the pressure drop in the fine channels, in the fin tube turbulence is desirable because it increases the heat transfer and it is laminar flow that is undesirable due to poor heat transfer. With eight 7/8" ID fin tubes in parallel, the flow goes laminar when the viscosity of the coolant exceeds about 2 cPs.

For 45% glycol this temperature is about 35ºC (DOW), so basically a system protected to -30ºC by 45% glycol will warm until the coolant approaches +35ºC, at which point it will cool well. In contrast a system with anti-freeze protection to -15ºC (~30% glycol) will only warm to ~+20ºC before the cooling is fully operational. While a system could be optimized for a colder climate to take advantage of a lower cell temperature, in most major solar markets an HCPV system will typically have coolant at above 30ºC anyway and tandem cells are not very temperature sensitive, so this is not expected to be worthwhile until very high volume. Alternatively, ethanol makes a good low-viscosity anti-freeze for cold climates, and it costs less than glycol as well, but can be corrosive so the fin tube might have to have a stainless-steel tube, which would raise its cost.

- Controlling Lens Temperature

The 'viscosity effect' may also prove to be an advantage. Fresnel/box CPV has problems with the lens refractive index, and thus the focal length, changing with temperature, causing impaired

focusing on cold weather. TLC should be much less sensitive to this since the thermo-optic coefficient of glass is typically less than 10 ppm/K (various web search hits) and silicone is typically around 350 ppm/K (Norris, et al.), but it is good to have an extra trick, and the pump speed could also easily be varied for further control of lens temperature. This could also come in handy in prototyping or early production where the front lens might be silicone-on-glass instead of roll-formed into the glass, or if injection-molded PMMA lenses are used initially (PMMA has a thermo-optic coefficient of ~130 ppm/K (Sultanova, et al., 2013)).

- Coolant Loop

Coolant alternates from absorbing heat in the receivers in one module to rejecting heat under the next module, with the other half of the receivers cooled on the return loop.

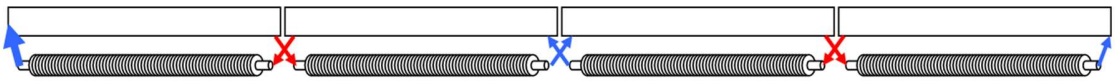

Figure B.9: *Coolant Loop.*

The above figure illustrates the principle but is not even close to scale. Each rejection section has multiple fin tubes in parallel, each fin tube is roughly 70 times longer than it is wide, and a typical trough (~100 meters long) would have roughly two dozen modules in a loop.

- Cooling Pump Redundancy

Putting many modules in a loop allows pump redundancy at low cost. Six low-cost INTG-3 pumps can cool an entire 100-meter-long-trough of modules, which reduces the cost of pumps compared to each module being its own loop, and also provides sufficient redundancy to keep working even if a pump or two fails (since the pumps are good for tens of thousands of hours, this allows pump failures to be taken care of as scheduled maintenance rather than requiring emergency repairs).

A $125 pump can thus service 22kW of receivers. Each pump also has its own 50W solar panel to drive it at a cost of $50, so the pumping costs $175/22kW or 0.8¢/W at 10 MW. Pump prices come down another 10% at 1000 pumps (22 MW) for 0.74¢/W. A further 10% reduction to 0.67¢/W is assumed for gigawatt scale; this is pessimistic because prices will almost certainly drop >10% at 50x the volume.

- Solar Panel Shading:

A corner case on the solar panels for the pump (and even the micro-tracking occurs when the sun is low to the east or west, and the neighboring trough shades these auxiliary-power panels while half of the trough is still illuminated. At this point the heat load can be almost half of full load, while the sun intensity on the auxiliary panels drops to 1/20 of full intensity.

The micro-tracking only needs 3W at a <0.1% duty cycle, or a few mW on the average, and this can be met with a small capacitor that stores a few milliJoules, and the V_{MPP} stays relatively constant so the capacitor charges even under low light.

The pumps, however, need continuous power or bad things happen quite fast. While 5% of power would allow 20% of pumping and temperatures would be less than 15K above normal, the resistance of the pump does not change, so without an extra regulator trying to keep current up (reducing voltage), the aux panels also would drop so far from their maximum power points that they would produce less than 1% of normal power and the pumps probably wouldn't even turn.

The easiest solution would be a regulator to keep the panel near the MPP, and larger solar panels. To keep within the chosen pump's spec would require 14% ($9^2V/24^2V$), and a regulator is not perfectly efficient so allowing 20% for a safety margin means quadrupling the panels, which would cost another 0.69¢/W, plus another ~0.06 ¢/W for the regulator. 0.75¢/W (and falling over time) is affordable but annoying, and it can be decreased. The pump power needed drops slightly faster than the square of the flow, so this would almost keep the temperature rise in the coolant at normal in the worst corner case. But with less light on the cells, the other thermal costs would be lower anyway, so the coolant could run at about 20% of the normal rate and the cells would be at the full-sun temperature. With 6 pumps in a loop only 2 of the pumps have the 'booster' panels, reducing the cost to 0.25¢/W while still providing redundancy (since even with one pump the cells will only warm by a few degrees).

An alternative answer is to provision two of the pumps with a half-sized panel on each edge of the trough so that one is always unshaded. This takes no extra watts of panel, but takes extra wiring run along the back of the trough and up through the supports. Given that panel costs fall over time, wiring is probably more expensive due to labor, but it is only a couple of dollars of wiring so it is worth considering in case the labor can be made simple.

It should be noted that this issue becomes slightly more significant with the single off-axis trough because the panel gets shaded when the mirrors are ~70% illuminated.

- Plumbing Connectors

The ¾" pipe is connected by flexible stainless steel ¾" water-heater connectors. These cost ~$5.40 each (retail, chain store web catalog 2016), and four are needed per module so that's ~$21.6 for 5.3 kW or 0.4¢/W. This is pessimistic because volume purchasing can beat retail.

- ***Volume and Cost of Coolant***

The two ¾" pipes hold 2.7 liters and four fin tubes with 7/8" ID hold 5.7 liters for 8.4 liters per module, or ~9 liters including in the receivers. Concentrated coolant (~100% glycol) is $13/gallon for a single 55-gallon drum (supplier web site, 2016), so a 45% mix (good to ~-30ºC), costs ~$1.60/liter or $14.40 per module or 0.27¢/W. A 10% volume discount is assumed for gigawatt-scale production. (The initial design uses a ½" pipe pair per half-trough, which is basically the same volume). This glycol cost allocated is pessimistic because most solar sites, especially for CPV, do not require freeze protection to -30ºC, especially since even -15ºC coolant wouldn't freeze solid at -30ºC but just would be too slushy to allow operation.

- ***Micro-tracking***

While part of micro-tracking is a receiver cost, micro-tracking also has module components. The receivers are all linked to a common drive rod. This is kept from buckling by a receiver wicket every 50 mm, so it can be quite modest in diameter and 3/16" rod is used since it is used for the wickets. Stainless steel 304 has a CTE very similar to that of copper pipe, and the 2.3 kg for 4060 mm costs $5.20 at current prices (web catalog, 2017). This supports a half-module or 2.65 kW, so the cost is ~**0.2¢/W**, and it supports twice the Watts in an optimized trough for **0.1¢/W.**

- Thermal Expansion Accuracy Limit

The steel rod will expand in length similarly to the coolant pipes because they will be at roughly the same temperature, being inside the same, white-painted housing (whose absorptivity can be picked to equalize temperature), but a difference of 10K, which produces 170 ppm expansion, is

allocated. The 0.7 ppm/K CTE difference times the +/- 40K operational temperature range amounts to another 28 ppm, so a total of 200 ppm is allocated.

The module will be micro-tracked from its middle, so each end is 2030 mm from that point and the total of 200 ppm is thus 0.4 mm. On a 100 mm lever arm this is 0.4%, or 4 mrads, and is the biggest source of micro-tracking error. This could easily be reduced with a longer micro-tracking arm, but there is no need with the current cells which are limited to ~1500X.

If tighter tracking tolerance is ever needed (for example, for 2000x cells if these become available in the future), the rod could be lower expansion and could be heated slightly by running a current through it to almost eliminate the micro-tracking error (however this is not expected to be used).

- Micro-Tracking Initial Alignment

The initial in-factory receiver alignment is simple. Receivers are slid onto the common drive rod through the holes in their micro-tracking beads, which fit loosely enough to slide freely. Even spacing is then ensured with a spacer that has a guide every receiver-width – the guides simply slide into the ~1-mm gaps between the receivers.

The receivers rest on a flat surface and are then clamped to it to ensure that the receivers are level, and thus that the lens faces are coplanar. There will be no appreciable force on the free-sliding beads. Each bead is then locked in place on the rod with a droplet of epoxy or other adhesive (the even spacing makes this easy to automate), and the epoxy cured (typically with UV, but other means such as time, heat and/or external catalysts can also be used).

This locks in the relative accuracy of all of the receivers. With four central 13.3-mm lens rows the leveling points on a lens are 50 mm apart, so 10μ of inaccuracy would only be 0.2 mrads, and lens accuracies are typically better than that. The beads still turn without excessive force on their wicket rods, but even this only takes enough clearance for the BMC's slightly higher thermal expansion; at 4 ppm/K difference and 100K range this is 0.4 μ/mm on the ~5 mm rod, or only a few microns which has no appreciable impact on the accuracy.

- Micro-Tracking Drive

A variety of stepper motors are available with lead screws. A Thompson ML17 with a 25-0031 lead screw provides sufficient force for a whole module and also provides a complete overkill 4μ step size, so the motor does not consume any of the tracking tolerance budget. It can 'only' drive

at ~10 mm/second, but that is still over 1000 times faster than needed, and even at that speed it consumes less than 3W!

I don't have a price from Thompson, but similar ML17 lead-screw steppers sell on Amazon for $50 *retail* for 3D printers, so $50 wholesale is allocated for a Thompson stepper, and this is probably pessimistic at high volume. $30 is allocated for a controller, a 5W solar 'panel', and a sheet-metal 'wart' on top of the module cover and a steel-rod pivot for the stepper. For the initial product that's 3¢/W, but the stepper is sized to drive a larger module; $80 thus covers a full 5.3 kW module and thus costs 1.5¢/W in large-scale production.

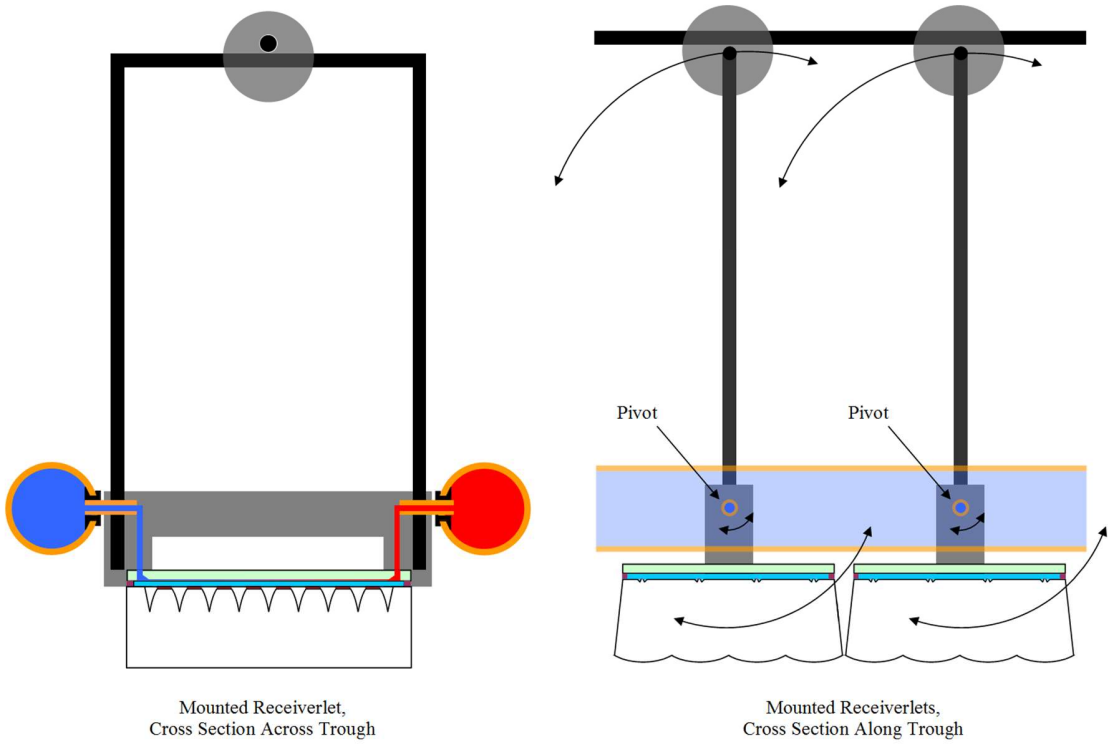

Figure B.10: Receivers on Rod.

- Junction Box and Connectors

The PVRD FOA lists a junction box at $6.00 and sealing it at $1.60, so initially that's $7.60 for 2.65 kW = 0.29 ¢/W, and 0.14¢/W with a 5.3 kW module for custom mirrors.

- Module Supports

The off-axis trough allows simple supports since a vertical support even a few inches wide will not block any sunlight. Supports can also be guyed at the top or even in the middle if needed to prevent buckling. This allows standard profiles like 2" channel, pipe or angle to be used; 6 feet per 5.3 kW dual-mirror module is ~$10, or 0.2¢/W.

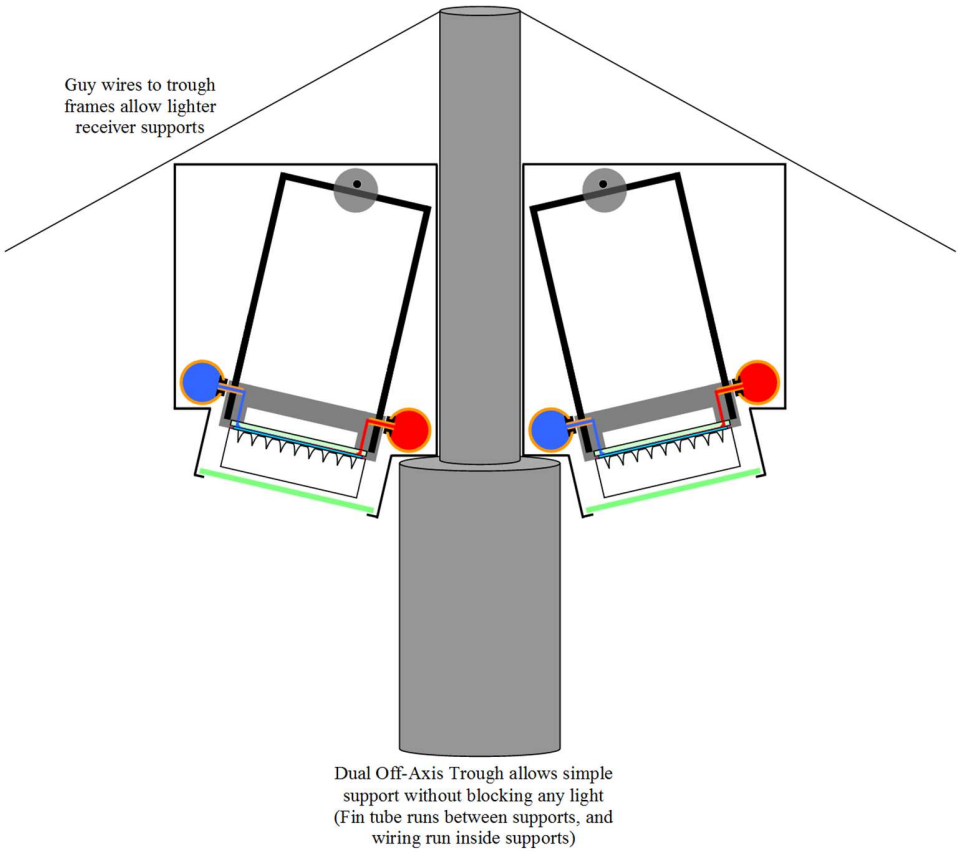

Figure B.11*: Dual-off-axis Slanted Modules.*

- Receiver slant Housing Shape

The slant of the receivers, and thus the fronts of the modules, should match the average slant of the half-trough mirror to minimize Petzval and other off-axis effects. The housing should also support the cold and hot pipes to keep them from bending, and flat surfaces are best for these (tiny glass-wool strips would be used to prevent contact and minimize galvanic corrosion, although with the module being baked dry this should not be an issue). Beyond this the housing shape is not critical and a reasonable shape was drawn for illustration.

B.9 Micro-tracked Module Assembly Walk-through

- **Cooling :**
  - Copper pipe is drilled
  - Copper pipe has connectors soldered on
  - Grommet seals are places in holes
- **Housing :**
  - Housing is roll-formed folded from sheet metal
  - Housing is cut to length
  - Mating ends are stamped onto housing
- **Receivers :**
  - Receivers are placed in jig
  - Jig is clamped to align receivers by flat lens surface
  - Copper pipes are aligned to axle tube
  - Copper pipes are pressed onto tube ends
  - Drive Rod is threaded through wicket beads
  - Adhesive is dabbed onto each bead to lock it to drive rod
  - Drive rod is pushed back and forth to ensure all seals seated
  - Pressure test for leaks
  - Each connector gets a squirt of conductive epoxy (if connector needs it)
  - Receiver terminal wire from previous receiver is inserted
  - Module terminals are attached
  - Epoxy is heated or UV cured to set
  - Receivers are unclamped
- **Housing**
  - Receiver rack is placed into housing jig
  - Housing placed over rack of receivers
  - Bottom supports retract to make room for cover glass
  - Silicone bead is dispensed
  - Cover glass is placed
  - Sides are clamped onto cover glass to lock in alignment
  - Clamped module is set aside for curing
  - Silicone end gasket bead is placed on cover glass end
  - When silicone is cured, drive actuator is attached along with stepper
  - Gasketed wart is placed and bolted on
  - Air filter over terminal wire
  - Junction box is placed and bolted on
  - Temporary end caps are placed for shipping

B.10 Module Assembly Cost

The module assembly involves stringing the receivers onto the drive rod, connecting the receiver wires, spacing the receivers, installing the cooling pipes on the axles, clamping the receivers flat,

the locking the receivers to the drive rod, mounting the pipes in the cover, installing the cover glass, and installing the drive and the sheet-metal wart that protects the micro-tracking drive.

While stringing the receivers on the drive rod would be easy to automate, connecting the terminal wire to the next receiver is probably not yet automatable and should take 15 to 20 minutes for 50 receivers.

Spacing the receivers is easily automatable with a 'comb' of plastic wedges, as is clamping the receivers in place once spaced. With slight tapered feed-tube ends, putting a coolant pipe onto the tubes should be automatable (and would take only a couple of minutes if done by hand, and adhering the beads to the drive rod is easily automatable. The cover will probably be placed by hand with a jig, the adhesive-dispense automated, and the cover glass placed with a jig.

The total is about ½ hour which at $40/hour factory time is ~$20 for 2.65 kW or ~0.75¢/W. Even without further automation, this would be cut in half to 0.38¢/W with an optimized trough of twice the width.

B.11 Single-Axis Tracking Requirement

The relaxed requirements for the first axis concentration bring the stiffness and accuracy needed to within the capabilities of low-cost single-axis tracker for flat panels. A single-axis tracker manufacturer who has also built two-axis CPV trackers has provided costing for doing so (Trackers Feina Quote, 2017), and even with the engineering cost included, for a single 10MW order the added cost per square meter is more than made up for by the higher efficiency, resulting in a lower tracker cost per Watt.

While the initial quote could be brought down through a competitive purchasing process and volume, its cost already good enough that the TLC tracking cost per Watt is much less than the tracking cost per Watt for any silicon panels, even after upgrading the tracker for CPV and de-rated TLC for not using diffuse light (which in a typical U.S. location brings TLC's equivalent efficiency down to ~30%. Even with current cells, TLC retains the tracking cost advantage until low-cost silicon panels reach 22% (versus typically 16% today).

- Tracker Cost Reduction:

No tracker cost reductions relative to the CPV-grade single-axis tracker quoted are currently in the cost model, but this is unreasonably pessimistic:

- The price used is from the first supplier's initial quote, un-negotiated
- The price used includes the non-recurring engineering
- Higher volume would reduce the cost differential
- The fine-tracking sensor could be using the trough's electrical output as a 'sensor'.

9. LIST OF REFERENCES

AK Steel. (2012). Specialty Stainless Sheet & Strip Stainless Price Schedule. *AK Steel web site*.

Allan. (1997). THERMAL CONDUCTIVITY OF CEMENTITIOUS GROUTS FOR GEOTHERMAL HEAT PUMPS. *Report for U.S. Department of Energy*.

Aluminum, U. (2013). Chemical Composition and Properties of Aluminum Alloys. *United Aluminum web site*.

Amada Miyachi America, Inc. (2016). Laser Welding Fundamentals. *Welding Guide*, 35-36.

Angel, et al. (2014). Advanced Manufacture of Reflectors. *DE-EE0005796*, 8.

ARPA-E. (2015). Micro-scale Optimized Solar-cell Arrays with Integrated Concentration (MOSAIC). *ARPA-E Funding Announcements*.

Assembleon. (2021). AX-501 SMT Placement specifications. *Web, accessed 2021-04-07*.

Atlas Steels. (2013). Aluminum Alloy Data Sheet 6063. *Atlas Steels web site*.

AZUR SPACE. (2020). CPV Triple Junction Solar Cell - Type 3C44C (5.5*5.5mm²). *downloaded from AZUR SPACE website*.

Bar-Cohen, et al. (2003). Design of Optimum Plate-Fin Natural Convective Heat Sinks. *Transactions of the ASME, Vol. 125*, 208-216.

Becky Laland. (2009). Kokomo company a glass act since 1888. *Cincinati Enquirer* , 55.

Bensch, W. (2011). Handling of semiconductor products: recommendations for handling, processing and operating conditions of CPV multi-junction solar cells. *downloaded from AZUR SPACE website July 2020*.

Beukema, et al. (2019). Moldable optical silicone elastomers spark creativity in LED lighting. *Laser Focus World*.

Blais. (2022). *VIEILLISSEMENT ACCÉLÉRÉ DE L'OPTIQUE SECONDAIRE D'UN SYSTÈME PHOTOVOLTAÏQUE À HAUTE CONCENTRATION.*

Borgerson, B. (2009). Life Expectancy of an Injection Mold. *MATRIX Plastic Products*.

Borrelli. (1999). *Microoptics Technology: Fabrication and Applications of Lens Arrays and Devices.* Corning, NY: Marcel Dekker, Inc.

Brakels. (2016). SunPower Solar Panels: One of the best panels lots of money can buy. *SolarQuotes Blog*.

Brown. (1984). The History of Power Transmission of the Rectenna,. *IEEE Transactions on Microwave Theory and Techniques, Vol. MTT-32, No. 9*.

Burroughs, e. a. (2010). A New Approach For A Low Cost CPV Module Design Utilizing Micro-Transfer Printing Technology. *AIP Conference Proceedings 1277, 163 (2010); doi: http://dx.doi.org/10.1063/1.3509179*, 163-166.

Campeau, et al. (2013). SunPower Module Degradation Rate. *SunPower Corporation*.

Carron. (2016). Startup makes residential solar panels twice as efficient. *Phys.org*.

Chiu, et al. (2014). Direct Semiconductor Bonded 5J Cell For Space And Terrestrial Applications. *EEE Journal of Photovoltaics, 4(1)*, 493-497.

Chylak, et al. (2006). Ultra-low-loop Wire Bonds. *Solid State Technology*.

Cooper. (2014). HIGH-CONCENTRATION SOLAR TROUGH COLLECTORS AND THEIR APPLICATION TO CONCENTRATING PHOTOVOLTAICS. *PhD Thesis*.

Corino. (2016). Die SunOyster – solare KWK. *SunOyster web site*.

CREE. (2017). CREE XLamp LEDs Solder Joint Reliability Study. *CREE Support Document*.

Davis. (1993). Machining of Aluminum and Aluminum Alloys. *ASM Specialty Handbook*, 341.

DeGroot. (undated). Silicones for Photonics. *Inorganic Polymers*.

Deign. (2017). Rapid-cleaning robots set to cut solar energy losses, labor costs. *NewEnergyUpdate*.

DeMilo et al. (2007). Thermally induced stresses resulting from coefficient of thermal expansion differentials between an LED sub-mount material and various mounting substrates. *Proceedings of SPIE, 6486*.

DOE. (2014). DOE-FOA-0001186 "APOLLO". *DOE*, 8.

DOE. (2016). Photovoltaics Research and Development (PVRD) 2. *FOA 1654*.

DOE. (2017). Photovoltaics (PV) Innovation Roadmap. *DE-FOA-0001764*.

DOE. (2018). Funding Opportunity Announcement (FOA) Number: DE-FOA-0001840. *EERE Funding Opportunity Announcements*.

DOW. (n.d.). *DOWTHER Engineering and Operating Guide for Inhibited Ethylene Glycol-based Heat Transfer Fluids.* DOW.

Elgin. (Undated). ÅngströmLink AL-3252 Datasheet. *Fiber Optic Center*.

Ellis. (2019). Faster, Cheaper, Simpler. *SMT007*, 82-84.

Elrod, S. (2008). III-V Solar Concentrators:. *Lecture at Stnaford*.

Espinet-Gonzalex, et al. (2012). Triple-junction solar cell performance under Fresnel-based concentrators taking into account chromatic aberration and off-axis operation. *AIP Conference Proceedings 1477*, 81-84.

Farris. (2021). Building A Better, Brighter LED Headlamp with Top-Side Alignment Process (TAP). *SMT007 Magazine*.

Fatemi, et al. (2013). Qualification and Production of Emcore ZTJ Solar Panels for Space Missions. *Photovoltaic Specialists Conference (PVSC), IEEE 39th*.

Feldman and Margolis. (2020). Q4 2019/Q1 2020 Solar Industry Update. *NREL/PR-6A20-77010*.

Fidaner and Wiemer. (2013). Power Output of Multijunction Solar Cell Wafers. *CS MANTECH Conference*, 439-442.

Fidaner and Wiemer. (2013). Power Output of Multijunction Solar Cell Wafers. *CS MANTECH Conference*, 439-442.

Frenzel at al. (2016). On the current and future availability of gallium. *Resources Policy*, Volume 47, Pages 38–50.

Furman, et al. (2010). A HIGH CONCENTRATION PHOTOVOLTAIC MODULE UTILIZING MICRO-TRANSFER. *Photovoltaic Specialists Conference (PVSC), 2010 35th IEEE*, 477-480.

GH Induction Atmospheres. (2011). The Brazing Guide. *GH Induction Atmospheres Turnkey Heating Solutions*.

Gonzalez, et al. (2009). Reliability analysis of temperature step-stress tests on III-V high concentrator solar cells. *Microelectronics Reliability*, 673-680.

Goodrich, et al. (2013). A wafer-based monocrystalline silicon photovoltaics roadmap: Utilizing known technology improvement opportunities for further reductions in manufacturing costs. *Solar Energy Materials & SolarCells*, 110-135.

Green, et al. (2016). Solar cell efficiency tables (version 49). *Progress in Photovoltaics*.

Gritz et al. (2009). Antenna-coupled-into-rectifier infrared sensor elements and infrared sensors. *US Patent 7,679,057*.

Guo, et al. (2017). Long-range hot-carrier transport in hybrid perovskites visualized by ultrafast microscopy. *Science*, 59.62.

H. Lv, W. H. (2018). An improved model to predict thermal runaway in concentrator III–V multi-junction solar cells. *International Journal of Low-Carbon Technologies, Vol. 13, Issue 4, https://doi.org/10.1093/ijlct/cty043*, 432–437.

Harman. (1992). Wire bonding - towards 6-sigma yield and fine pitch. *IEEE Transactions on Components, Hybrids, and Manufacturing Technology*.

Harter, et al. (2017). The Effect of Area Shape and Area Ratio on Solder Paste Printing Performance. *SMT Magazine*, 44-58.

Hayashi, et al. (2015). High-efficiency thin and compact concentrator photovoltaics with micro-solar cells directly attached to a lens array. *Optics Express*.

He, et al. (1998). Reliability in Large Area Solder Joint Assemblies and Effects of Thermal Expansion Mismatch and Die Size. *International Journal of Microcircuits and Electronic Packaging*, 297-305.
Hernon Manufacturing Inc. (2018). Precision Jet Dispensig Valve. *Hernon web site* .
Hopmann and Röbig. (2016). Injection molding of high precision optics for LED applications made of liquid silicone rubber. *AIP Conf. Proc. 1713, 040003-1–040003-5; doi: 10.1063/1.4942268*.
Horowitz, et al. (2015). A Bottom-up Cost Analysis of a High Concentration PV Module. *NREL*, 15.
Huang, et al. (2017). *High performance surface-modified TiO2/silicone.* Scientific Reports | 7: 5951 | DOI:10.1038/s41598-017-05166-7.
Huawei. (2017). How String Inverters Are Changing Solar Management on the Grid. *Greentech Media White Paper*.
J. Wheeldon, A. W. (2011). Efficiency Measurements and Simulations of GaInP/InGaAs/Ge Quantum Dot Enhanced Solar Cells at up to 1000-Suns Under Flash and Continuous Concentration. *CPV-7*.
Jacob and Nitz. (2017). The Role of Draft Facets in Temperature Effects in Silicone-on-Glass Fresnel Lens Applications. *CPV-13*.
Jordan, et al. (2016). Compendium of photovoltaic degradation rates. *Progress in Photovolataics*, 978-989.
Kontges, et al. (2014). Review of Failures of Photovoltaic Modules. *External final report IEA-PVPS*.
Kramer. (2006). Mineral COmmodity Summary 2006: Gallium. *U.S. Geological Survey*.
Kurtz. (2009). A Bright Future for CPV. *NREL*, 14.
Kurtz, et al. (2013). Photovoltaic Module Qualification Plus Testing. *NREL*.
Kutscher, et al. (2010). Line-Focus Solar Power Plant Cost Reduction Plan. *NREL*, 15.
Las Marias. (2017). Achieving the Perfect Solder Joint: The Many Perspectives on Soldering. *SMT Magazine*, 16-26.
Lechivic et al. (2009). Solder Joint Reliability. *Materials Science and Technology*, 1-8.
Levine. (2016). Wire Bonding. *Electronic Device Failure Analysis, volume 18 #1*, 22-28.
Limata. (2019). *Limata X1000 series data sheet (accessed 2021-09-03).* Retrieved from imata.com.
Löckenhoff, et al. (2010). Water Cooled TJ Dense Array Modules for Parabolic Dishes. *CPV-6.* Freiburg: AIP Conference Proceedings 1277, 43 (2010); https://doi.org/10.1063/1.3509229.
Luce and Cohen. (2010). THE PATH TO VOLUME PRODUCTION FOR CPV OPTICS. *IEEE PVSC*.
M. Steiner, G. S. (2013). Modeling the thermal runaway effect in CPV modules. *AIP Conference Proceedings 1556, 230 (2013); https://doi.org/10.1063/1.4822238*.
M. Steiner, G. Siefer, and A. W. Bett. (2013). *Modeling the thermalrunaway effect in CPV modules.* AIP Conference Proceedings.
M. Steiner, P. K. (2015). CPV module design optimization for advanced multi-junction solar cell concepts. *AIP Conference Proceedings 1679, 100005 (2015); https://doi.org/10.1063/1.4931552*.
Mauve and Mayoh. (2018). The Printed Circuit Designer's Guide to Thermal Management with Insulated Metal Substrates. *I-007 eBooks*, 15.
Maxfield. (2020). Laser Imaging PCB Solder Masks is a Game Changer. *Electronic Engineering Journal*.
Mehos, et al. (2016). On the Path to SunShot: Advancing ConcentratingSolar Power Technology, Performance, and Dispatchability. *NREL*.
Meiser. (2014). Analysis of Parabolic Trough Concentrator Mirror Shape Accuracy in Laboratory and Collector. *Doctoral Thesis*, 49.
Miller and Stephens. (1979). Solar Energy Collection System. *U.S. Patent 4149521*.
Mints. (2017). Slippery Slopes — The Trolley Problem, Solar Module Price Choices and Climate Change. *Renewable Energy World*.
Munsell. (2016). Solar Module Prices Reached 57 Cents per Watt in 2015, Will Continue to Fall Through 2020. *GTM Research, March*.
N. Caillou. (2021). DEVELOPPEMENT D'UN MODULE OPTIQUE A DOUBLE CONCENTRATION POUR DU PHOTOVOLTAÏQUE A CONCENTRATION. *Masters' thesis*.
N. Hayashi, M. T. (2017). Thin Concentrator Photovoltaic Module with Micro-Solar Cells which are Mounted by Self-Align Method Using Surface Tension of Melted Solder. *CPV-13, AIP Conf. Proc. 1881, 080005-1–080005-6; doi: 10.1063/1.5001443*.

Norman. (2017). HCT-PV First Functional Prototype. *SIPS Proposal 1654-1632*.
Norman et al. (2018). Trough-Lens-Cone Optics with Microcell Arrays: High Efficiency at Low Cost. *AIP Conference Proceedings, Volume 2012, Issue 1, id.090006*.
Norman, et al. (2010). Systems for Cost-Effective Concentration and Utilization of Solar Energy. *USPTO 20200263709*.
Norman, et al. (2019). *Trade-offs and Optimizations in Trough-Lens-Cone Optics for High Efficiency at Very Low Cost.* CPV-15.
Norman, et al. (2020). Massively-parallel Microcell Arrays and Thermal Runaway in Unilluminated Cells. *AIP Conference Proceedings 2298, 040001 (2020); https://doi.org/10.1063/5.0032175*.
Norman, et al. (2021). Molding Arrays of Tertiary Optical Elements for Microcell. *CPV-17*.
Norman, et al. PVSC. (2019). Reducing CPV Materials Cost Through Multistage Concentration. *IEEE 46th Photovoltaic Specialists Conference (PVSC)*.
Norris, et al. (n.d.). Silicone Materials for Optical Applications. *Dow Corning Web SIte*.
NREL. (2009). [3] NREL, "Concentrating Solar Power Projects: Solar Electric Generating Station I), NREL, (2009). *NREL Report*.
Nurmi. (2005). Reliabity of SnAgCu Solder Joints Under Thermo-Mechanical Stresses. *PhD Thesis*.
Olny. (2014). ESTIMATING SHEET METAL FABRICATED PARTS. *ETM Manufacturing Cost Guide*, 30.
Patrick. (2015). PCB Stackup Cost Information. *AerospacePal*.
Pedretti-Rodi, et al. (2015). Trough Collector withConcentrator Arrangement. *USPTO 20150354856*.
Persson. (n.d.). Machining guide lines of Al / S iC. *MMC-Assess Thematic Network*.
Persson. (n.d.). Machining guidel ines of Al / S iC. *MMC-Assess Thematic Network*.
Pihl, et al. (2012). Material constraints for concentrating solar thermal power. *Energy 44*, 944-954.
Powel, et al. (2015). The capital intensity of photovoltaics manufacturing: barrier to scale and opportunity for innovation. *The Royal Society of Chemistry*.
Pruzan, et al. (1991). Design of high-performance sintered-wick heat pipes. *Pergamon Press Plc.*
Qin, et al. (2015). Advances in wire bonding to lower the cost of interconnect. *Chip Scale Review, Volume 19 #6*, 26-29.
R. Herrero, M. V. (2015). Understanding Causes and Effects of Non-Uniform Light Distributions on Multi-Junction Solar Cells: Procedures for Estimating Efficiency Losses. *AIP Conference Proceedings 1679, 050006 (2015); https://doi.org/10.1063/1.4931527*.
Rickover, H. (1953). Paper Reactors, Real reactors. *Letter, quoted in Hearings Before the Joint Committee on Atomic Energy*, 1702.
Rioglass. (2017). Quote received from Rioglass for 10 MW of tempered-glass RP3 trough mirrors. *Private e-mail*.
S. Li, L. S. (2012). A Composite Hardness Stamp in 184 PDMS for Nanostructures Transfer in High Fidelity. *12th IEEE International Conference on Nanotechnology (IEEE-NANO)*.
Sargent and Lundy. (2003). Assessment of Parabolic Trough and Power Tower Solar Technology Cost and Performance Forecasts. *NREL*.
Schachinger, M. (2020). EU spot market module prices by technology. *pv-magazine.com*.
Schneider. (2016). A Quantitative Comparison of Central Inverters and String Inverters in Utility-Scale Solar Systems in North America. *Schneider Electric*.
Shaprio. (2017). Filler metals for brazing titanium to aluminum at 550-580ºC. *Titanium-brazing-com website*.
Sharma, et al. (2015). A carbon nanotube optical rectenna. *Nature Nanotechnology*.
Shemer. (2017). CSP heliostat innovations set to deliver big cost reductions. *NewEnergyUpdate: CSP*.
Shin-Etsu. (2016). Characteristic Properties for SIlicone Rubber. *Shin-Etsu Silicone*.
Shirazy, e. a. (2013). Capillary and wetting properties of copper metal foams in the presence of evaporation and sintered walls. *International Journal of Heat and Mass Transfer*, 282–291.

Shirazy, et al. (2010). A Parametric Investigation of Operating Limits in Heat Pipes Using Novel Metal Foams as Wicks. *ASME 2010 8th International Conference on Nanochannels, Microchannels, and Minichannels*.

South Bay Technology. (undated). Characterizing the Model 850 Wire Saw. *South Bay Technology web site*.

Spectrolab. (2011). C4MJ Metamorphic Fourth Generation CPV Technology. *Datasheet on web site*.

Sultanova, et al. (2013). Characteristics of Optical Polymers in the Design of Polymer and Hybrid Optical Systems. *Bulg. J. Phys. 40*, 258-264.

SUNLAB. (2014). Current Directions for III-V Multi-junction Solar Cell Development. *ELG6373*, 12.

SunPower. (2017). E-SERIES COMMERCIAL SOLAR PANELS. *SunPower web site*.

Trabish. (2013). GlassPoint Brings 7 MW of Enclosed Trough CSP On-Line in Oman. *Greentech Media*.

TRUMPF. (2016). Laser welding for sheet metal parts . *Machining News*.

TRUMPF. (n.d.). TruLaser Weld . *TRUMPF Brochure*.

Ulrike, et al. (2015). Gradient index antireflection coatings on glass containing plasma-etched organic layers. *Optical Materials Express*, 1259-1265.

Victoria, et al. (2014). Temperature effects on two-stage optics made of silicone. *AIP Conference Proceedings 1616*.

Vitro. (2020). Starphire® Technical Product Data. *Datasheet*.

Wadia, et al. (2009). Materials Availability Expands the Opportunity for Large-Scale Photovoltaics Deployment. *Environ. Sci. Technol. 43*, 2072–2077.

Wallace. (2017). Precision glass molding boosts gamut of optics. *Laser Focus World, May*, 31-36.

Wang, et al. (2009). An Assessment of Immersion Silver Surface Finish. *Journal of ELECTRONIC MATERIALS, Vol. 38, No. 6*.

Wells. (1981). Electromagnetic Radiation Collector System. *U.S.Patent 4281640*.

Wheelwright, et al. (2014). Freeform lens design to achieve 1000X solar concentration with a parabolic trough reflector. *Proceedings of SPIE - The International Society for Optical Engineering (Vol. 9293)*.

Wiesenfarth, P. B. (2017). *CURRENT STATUS OF CONCENTRATOR PHOTOVOLTAIC (CPV) TECHNOLOGY 1.3.* NREL/Frauhofer.

Wong, et al. (2011). Visualization and Evaporation Resistance Measurement forGroove-Wicked Flat-Plate Heat Pipes: Various Working Fluids and Powder-Groove Evaporator. *10th Intenational Heat Pipe Symposium*.

Woodhouse, et al. (2016). The Role of Advancements in SolarPhotovoltaic Efficiency, Reliability, and Costs. *ON THE PATH TO SUNSHOT*, 30.

Y.S. Liao, J. H. (2004). A Study to Achieve a Fine Surface Finish in Wire-EDM. *Journal of Materials Processing Technology*.

Yeomans. (1994). Radiant energy collecting apparatus. *U.S. 5592932*.

Yeomans. (1997). Radiant Energy Collecting Apparatus. *U.S. Patent 5592932*.

Zogbi. (2016). *Passive Electronic Components: State of the industry.* TTI Market Eye.

Zuehlke and Gebhardt. (2016). Dicing silicon carbide power devices using thermal. *Chip Scale Review*, 39-42.