PERISCOPE
PERIapsis Subsurface
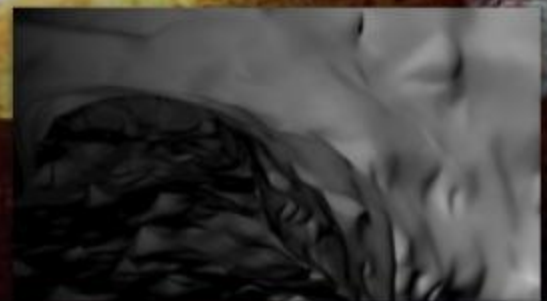# Cave Optical Explorer

NIAC Phase 2 Final Report
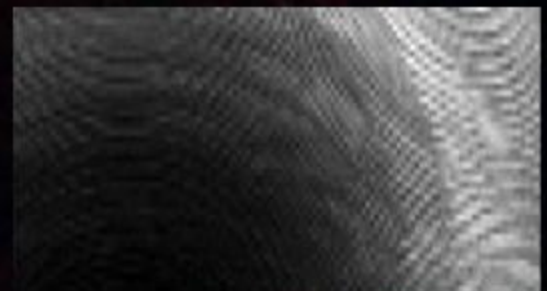Nosanov Consulting, LLC
PI: Jeff Nosanov

Co-Is:
Andreas Velten, University of Wisconsin
Karl Mitchell, Jet Propulsion Laboratory

In situ LIDAR scan of cave interior

Beyond-line-of-sight reconstruction of cave interior from remote position

PERISCOPE Phase 2 Final Report
2015-2018

PI: Jeff Nosanov, Nosanov Consulting
Co-Is: Karl Mitchell, JPL, Nitin Arora JPL (departed JPL prior to project conclusion)
Andreas Velten (University of Wisconsin, Madison)

This project could not have been completed without the hard work of many others. Please see the Special Thanks section for a list.

Note on cover graphic: the inset image at upper right is a photograph of a lunar skylight. The two greyscale graphics at the lower right represent our system performance. The upper one is a LIDAR scan of a cave skylight at El Malpais National Monument, New Mexico. The lower one is a reconstruction of that cave using the portable system developed under this project.

Introduction

This final report is the culmination of an idea that began six years ago. The question was whether a tabletop-scale, beyond line-of-sight imaging technique (developed by Co-I Andreas Velten at the MIT media lab) could be used at a 10 km-scale orbital altitude around the moon to image the interior geometry of skylight and other subsurface cave structures. The past six years have seen team members change institutions, have children, and experience other significant life changes. The world has also changed in significant ways.

The first Phase 1 NIAC proposal was rejected in 2013, the second version accepted in 2014, and the Phase 2 proposal was accepted in 2015. I would like to thank Josh Schoolcraft at JPL for helping mature the concept prior to the initial unsuccessful submission, which led the following year to the successful Phase 1 submission. As PI it has been a privilege and honor to lead this idea through these changes, and I speak for the team's immense gratitude and joy in submitting this final report detailing the work on the PERIapsis Subsurface Cave OPtical Explorer (PERISCOPE) concept.

Figure 1, right upper: current observation of lunar skylight geometry is limited to just beyond the "rim" of the opening, even under ideal illumination conditions. Figure right lower: active laser illumination of skylight floors enables multiple-reflection beyond-line-of-sight imaging of interior structures.

We are extremely grateful to have had the opportunity to explore this concept. This final report will be organized roughly around the proposed tasks from the initial proposal. The report will try to explain where the work deviated from the work we planned, and where it went. We will conclude with a realistic appraisal of this mission concept from an implementation point of view, and some ideas of how it may fit into the current discussion of both SMD and HEOMD operations.
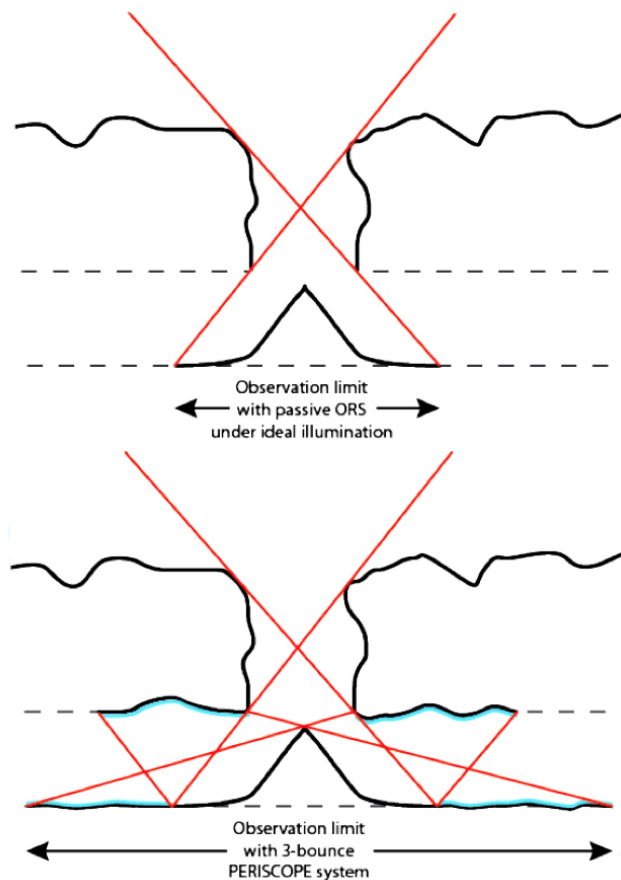

Figure 1: Beyond line-of-sight system can "see" farther into cave structures than orbital remote sensing using ambient lighting.

This final report is dedicated to the children who came into our worlds during the Phase 2 period:

Ilian Raahi Velten

Rose Celeste Ashley

Science

The Science at project inception (2012)

Although their presence was hypothesized over 130 years ago [Nasmyth & Carpenter, 1874], the discovery of putative lava tubes (caves) on the Moon and Mars is only a recent phenomenon, and even then the interpretation as lava tubes is not 100% certain. At inception of the PERISCOPE project, observations of skylights above putative caves on other worlds had only happened a few years prior, with Martian skylights reported first on the basis of Mars Odyssey's THEMIS (thermal emission spectrometer), in 2007 [Cushing et al., 2007], followed in 2009 by similar and larger skylights on the Moon by the SELENE (Kaguya) camera system [Haruyama et al., 2009]. Following their discovery and initial analysis, 100s more skylights, "putative caves" and "atypical pit craters" were found, in both reassessment of existing data and some additional planned observations. Arguments for underlying caves were initially circumstantial, being typically found within basaltic lava flows in contexts similar to lava tubes on Earth. However, follow-up imaging using oblique viewing angles and coordinated lighting geometries confirmed subsurface voids up to 20 meters beneath overhanging rock [Robinson et al., 2012].

While martian voids may have an astrobiological component to their attraction [Boston et al., 2001], lunar speleology is motivated more by what subsurface voids represent to 1) basic lunar science, and 2) lunar engineering. Planetary science will benefit significantly from direct exposure to crustal rocks from deep within lunar volcanic and impact melt deposits. Significantly for long-term human exploration or habitation, protection from many surface hazards (vacuum conditions, micrometeorites, dust, solar wind, extreme temperature cycling, and radiation) is achieved only a few meters beneath the surface. To simplify base construction and reduce

engineering costs, such structures may require only a minimum of "retrofitting" to become useful as habitations or caching supply depots.

State of Science at Phase II beginning

The PERISCOPE study focuses primarily on lunar caves, due to the potential for being imaged in orbital scenarios. In the intervening years, from 2012-2015, scientists developed further rationales and interest in the scientific value of lunar caves. It does not appear that they are likely to be sinks for water-ice [P. Hayne, pers. Comm.] due to the relatively warm temperatures (~-20 degrees celsius) in the caves leading to geologically-rapid migration of unbound water due to sublimation, and inevitable loss through any skylights. However, the skylights themselves reveal apparent complex layering [Robinson et al., 2012] which may speak to a more complex multi-stage evolution of mare flood basalts than previously considered, and so their examination may provide even more insight into the lunar mare, which in turn provide a primary record of early solar system crustal formal and evolution processes. Further extrapolation of these insights can be found within the exoplanet community of researchers, who find the information useful for calibrating star formation and planetary evolution models. In addition, catalogues of lunar [e.g. Wagner and Robinson, 2014] and martian [e.g. Cushing and Okubo, 2015] skylights, "caves" or "atypical pit craters" have been developed, with numbers for both bodies now in the low hundreds thanks to additional high resolution surveys and revisiting the existing image databases.

Independent developments since Phase II beginning

Since Phase II studies commenced, gravity data from the GRAIL spacecraft [Chappaz et al., 2016], and radar sounding from the SELENE (Kaguya) Lunar Radar Sounder (LRS) [Kaku et al., 2017], have both tentatively confirmed the existence of an anomaly consistent with a lava tube extending for many tens of kilometers from one of the many pit craters. The growing interest in lunar caves has led to preliminary studies into mission architectures that could descend into them [Haruyama et al., 2015; Kerber et al., 2018], primarily to study the layering in the skylight walls and unravel lunar mare evolution, but also to perform initial reconnaissance into the interior, and explore the lava tube hypothesis (admittedly, viable competing hypotheses are somewhat thin on the ground).

Of perhaps more immediate concern, in the context of a growing interest in establishment of a possible future lunar base, is the lack of useful data for planning and navigating a traversal into these systems, as well as confirmation that the accessible extent is indeed useful for HEOMD purposes. An initial reconnaissance of the hundreds of lunar skylights would allow considerable refinement towards the most desirable target for development of such infrastructure.

Science and exploration traceability

We wrote the Phase 2 proposal assuming that development of a Science Traceability Matrix (STM) would be a reasonable goal, and proposed to do just that. However, as we began our Phase 2 study and thoroughly evaluated the literature we realized that the state-of-the-art knowledge of the inside of the skylights was extremely limited. So limited in fact that we can only confidently and clearly define a single specific science objective with associated measurement requirements for such an STM:

"Verify that lunar pits are collapses into lava tubes."

The objective would be to demonstrate lava tube-like structure/morphology, with a measurement requirement of 5-m structural resolution at N-skylight diameters, with N >5-10.

A secondary objective to measure potential benches of lava tubes on walls would require resolutions of 0.1-0.5 m, and was considered "daisy-chained" from the above objective, and of unclear science benefit even if discovered, beyond putting some constraints on flow models, and so was deprioritized.

Other use cases for PERISCOPE have emerged during the same period that fall outside of traditional science objectives, into the domain of providing critical support/reconnaissance for other activities. These involve characterizing cave structure for:

1. providing protection for space weather and temperature extremes, and
2. potential exploration by cave-specific robotic mobility platforms,
3. assessing astronaut accessibility.

The first of these would be met using the measurement requirements stated above for lava tube verification. The second and third are discussed below.

A single PERISCOPE mission could advance these science, human accessibility, and robotic accessibility goals across dozens or hundreds of caves. We conclude that this is the principle value offered by a potential PERISCOPE mission as a subsurface reconnaissance mission to be conducted prior to further surface or base operations.

Reconnaissance for robotic and human explorers

We have contacted the PIs of two efforts to develop platforms to enable robotic exploration of caves on the Moon. We consider these to be generally representative of all robotic platform, as one is a high-TRL ready solution with limited cave mobility, and the other is a low-TRL solution that offers extreme cave mobility that can traverse a wide range of different surfaces.

The first, Axel (Nesnas et al., 2008), is a wheeled robot considered sufficient technologically-ready (TRL 5-6) that is in the process of being proposed for implementation in the near-future under Discovery 2019 (PI: L. Kerber). It has two sub-implementations, including the traditional Axel (a two wheeled rover with a tether that connects it to a mothercraft) and DuAxel (formed by having two Axels dock to either side of a central module, which offers a mode in which a single Axel could undock and drive away while remaining tethered to the other one which is anchored on a surface). Both implementations can work on any slope up to 90 degrees, given rappelling capability, an up to 30 degrees unaided by a tether (experiments/published), and tethers of up to ~500 m in length are achievable (I. Nesnas, pers. comm.). However, Axel is limited in its ability to traverse rocks of more than a wheel radius in height, although experimental solutions exist for traversing 1.8x wheel radius, such a platform would be more complex to support both flat and extremely rock terrains. As a result, to accommodate reasonable wheel sizes, data is required to allow <=1 m 3D resolution along potential traverse from a landing site, into a cave interior.

The second solution, LEMUR (Parness et al., 2017), is a limbed robot that has been undergoing field testing at terrestrial caves at El Malpais National Monument under PSTAR funding, and is considered to be at a less mature level (TRL 4-5), and deemed not yet ready for proposal for flight missions. It has a spider-like design with gripping "feet" on the end of multiple limbs with various degrees of freedom, designed for climbing on any secure surface, and is largely indifferent to roughness scales of >~ 2 cm, as long as there exists fine-scale millimetric roughness (often measured as 'kurtosis') on secure wall and ceiling surfaces allows climbing under the most challenging of circumstances (A. Parness, pers. comm.). Under circumstances where kurtosis is insufficient, its ability to "walk" is comparable to Axel, and so requirements are similar, but there is additional benefit to producing photometric data, particularly of brighter (hence rougher) surfaces, which allows some constraint of kurtosis.

Finally, requirements for human explorers are more complex. In Earth, cavers are able to traverse almost any feasible surface or roughness using very limited tools. However, in a lunar environment, the requirement to wear a protective spacesuit is limiting, especially given that basaltic rocks can be rough and sharp, thus providing a hazard. We propose that deriving recon requirements at this stage is premature, and that further research in astronaut cave mobility is desirable. Furthermore, considerable testing and refinement of spacesuits may be necessary to enable safe astronaut mobility within such an environment, unless the cave in question is found to be particularly benign.

As a result, this non-science reconnaissance objective would be to provide path planning for a robotic or human explorer, with a measurement requirement of <= 1 m structural resolution at N-skylight diameters is sufficient, with N >5-10. The objective would be enhanced by derivation of BRDFs (bidirectional reflectance distribution function) in order to constrain roughness models, and higher-resolution (<= 10 cm) structural resolution to highlight climbing challenges.

Scattering link budget

We had also proposed to provide real-world inputs for hypothetical and observed lunar cave, skylight and surface scattering properties, based on field observations and literature review. These would have included full 3D laser scanning mapping of an Earth analog skylight, working with a LiDAR technician, cataloguing of published lunar optical scattering and structural properties, and provision of terrestrial basalts with different weathering states. Literature review during the study showed that we know very little about optical scattering of candidate materials, and so some degree of speculation is necessary.

Most of the moon's surface undergoes extensive space weathering and is covered with dust, and so its surface optical properties are not appropriate for a permanently shadowed cave. Terrestrial basalts may not be appropriate, as they are rapidly exposed to an eroding and oxidizing (largely due to water) environment. The best analogs are probably unweathered terrestrial basalts, which are chemically similar to those erupted on the Moon, but may be extremely dark, with as low as ~3% optical albedo [Evans et al., 1982]. Furthermore, freshly-emplaced lava often has enhanced specular characteristics with corresponding lesser reflectivity at other angles, especially those close to backscatter, which is of particular relevance to PERISCOPE geometries; This effect may be mitigated somewhat by the tendency of "drip" structures on the ceilings of caves. As a result of these considerations, we propose that any PERISCOPE implementation for studying lunar caves will have to be engineered based on extremely conservative assumptions of the Bidirectional Reflectance Distribution Function (BRDF) of the surface. For want of better data, we assume a reflective surface equivalent to a lambertian (hemispherically isotropic) scatterer with >1% albedo for cave interiors and >8% for skylight floors (based on lunar mare), which for the anticipated 3-bounce system gives an effective reflectivity multiplication factor of 6.4 x 10^-5. These assumptions should allow a link budget to be constructed that is sufficient to guarantee mapping. However, due to the large laser powers and potential dynamic ranges that this necessitates, we propose a logical next step should be to further refine these assumptions by implementing a terrestrial analog field campaign, including collection of fresh lava samples for analysis using an optical goniometer, and testing of PERISCOPE over a recently-emplaced basaltic lava tube. This could be a proposed activity in a 2019 NIAC Phase 3 proposal.

Complementarity of Radar

We explored value-added and complementary activities and instruments that could be part of a PERISCOPE caves-focused lunar mission, e.g. use of PERISCOPE as a topographic mapper, detection of non-basaltic materials and resources in caves, utility of PERISCOPE in polar studies, addition of radar sounding, etc. Radar sounding is specifically of value to mapping distribution of cave systems (see A-Team study, below, objective 1b), but is typically extremely challenging in environments where caves are formed, due to "clutter". However, the Moon might provide a compelling radar sounding target, due to the lack of liquid water which makes terrestrial radar sounding more challenging.

Radar sounding through basaltic rock is typically limited to 10-20 wavelengths of depth penetration [Y. Gim, pers. comm.], and the ability to resolve scatterers (cave floors and ceilings)

is problematic at sub-wavelength scales. As a result, caves inferred from skylights observed on the moon, with void scales of 10s up to ~100 m, and widths and depths of ~100 m), benefit from operating at frequencies consistent with wavelengths of a few to 10s of meters. While this makes challenging achieving our target spatial resolutions of 5-, 1- and 0.1-m discussed above, the ability of radar sounding to map potentially extensive lava tube systems would be valuable for determining their extents and orientations, and showing lava tubes where no skylights exist. Therefore we propose that PERISCOPE and a radar sounding solution could work well together in tandem.

The only previous radar sounder to fly to the Moon was the 4-6 MHz (50-75 m) Kaguya Lunar Radar Sounder (LRS). However, the spatial resolution was coarse compared with requirements to detect caves, and even though one study proposed detection [Kaku et al., 2017] of one large cave, its results were marginal at best, if not questionable.

One previously flown radar sounder, SHARAD, has flown to Mars onboard Mars Reconnaissance Orbiter (MRO). This functions at a much more appropriate 15-25 MHz (12-20 m), making it ideal for the larger observed skylights, and serving as a model for our assessment of radar complementarity. If flown to the Moon, SHARAD would be able to resolve mid-large caves, given its ~12-40 m, although this would be insufficient to give the 3D resolutions discussed above. For the MRO 400-km orbit, it was capable of achieving horizontal surface resolution of around 3-6 km (cross-track) and 0.3-1 km (along-track) [R. Seu et al., 2007]. While it is possible to enhance the cross-track resolution in post-processing, using super-resolution techniques where ground-tracks are close together [Raguso et al., 2017], it has not been reported to date to have resolved any of the proposed Martian tubes, possibly in part be due to their smaller scales (44% of lunar, based on gravitational scaling). However, the GRAIL mission demonstrated the feasibility of, and collected gravity data to make easier to implement, elliptical orbits as low at ~15-km on our airless moon. If implemented, we might anticipate horizontal resolutions improved by more than an order of magnitude, to 112-225 m (cross-track) and 11-38 m (along-track), and potentially improved even further using superresolution techniques, which should be relatively trivial using the PERISCOPE orbital specifics (adjacent tracks of <100 m separation) discussed in the phase I report. As a result, we propose that cave detection and planimetric mapping is viable with radar sounding on the Moon implemented using the same mission design requirements as PERISCOPE.

Note that a system with an expanded frequency range up to maybe 15-100 MHz would be preferred [M. Mastrogiuseppe, pers. comm.], allowing smaller caves to be resolved, and giving a greater bandwidth for improved processing.

We can simultaneously use PERISCOPE to validate radar techniques applied to cave detection, due to the improve spatial resolution. Furthermore, use of PERISCOPE as a basic LIDAR will enable surface topography to be refined, improving radar sounding modeling. In addition to the subsurface voids identified from orbit on the Moon, over 1,000 volcanic lava tube skylight candidates have been detected on Mars using a variety of imaging platforms from orbit over the past 12 years [starting with Cushing et al., 2007]. A high level of interest surrounds these

discoveries within both the science and human exploration communities. Of particular significance for astrobiologists is the potential for near-surface underground environments to support indigenous life on Mars as protective habitats. Thus the application of PERISCOPE-based technology in other solar system settings could find strong future support [Cushing et al., 2007].

JPL A-Team Study, 2016 Results

The PERISCOPE team conducted an A-Team study via the JPL Innovation Foundry on October 18th and 20th, 2016. The goal of the study was to assess the feasibility of the PERISCOPE mission concept and identify major challenges. The threshold goal was to mature the instrument concept to the point where it could be taken to a Team X instrument study, and the target was to mature the mission concept including an estimate of feasibility, cost and concept of operations.

Specific questions included:

- Does the optical design work?
- What are the issues?
- What is the best power system?
- Is there something better than the baseline?
- Specifically, will supercapacitors work better than batteries?
- Are the pointing scenarios feasible?
- Would dynamic slewing be preferable to the baseline fixed pointing?
- Can the instrument satisfy the requirements from the science? Identify a feasible minimal lunar mission design concept, within realistic program parameters.
- Stretch objective: Explore additional concepts, including as part of larger Moon and Mars mission architectures.
- Stretch objective: Is there added value to having other instruments on the platform, for example radar sounding?

Table 1: Details on location and known geometry of several example lunar skylights.

Mission	Location	Lat	Lon	Size (m)	Floor Depth (m)
T	Lacus Mortis	44.962	25.61	140x110	80
B	Central Mare Fecunditatis	-0.917	48.66	130x110	30
T	Mare Tranquilitatis	8.335	33.222	100x88	105
B	Mare Ingenii	-35.948	166.053	100x68	45-65
T	Marius Hills	14.091	303.23	58x49	40
W	>>200 more	>-61	<+45	<~100x100 (Median ~20x20)	<4-50

We identified the following questions as major mission drivers, with associated measurement goals:

1. Science Questions

    a. Are Lunar skylights the result of lava tube collapse?
        i. Detect lava tubes
    b. Do tube orientations support a specific model of impact melt or mare formation?
        i. Constrain orientations of N>6 tubes per basin (to ~10°)
    c. What is the cooling history of mare floods/flows?
        i. Measure benches on tube walls (0.5m resolution)
    d. Do subsurface voids preserve volatile materials indicative of past impact or outgassing activity?
        i. Detect "bright patches" of significant (~5x+) backscatter contrast

2. Exploration Questions

    a. Do significant permanently shaded volumes exist underground on the Moon suitable for human exploration or bases?
        i. Possibly target 3-m resolution 3D model beyond illuminated zone
        ii. Provide structural data suitable for assessment of robotic access/navigation challenges/constraints (breakdown rubble) 0.5-m 3D model, both rubble pile and into lava tube (Lemur-inspired)

Problem: Laser Power

Among major technical challenges identified by team is the requirement of a 1kw laser. A 1kW laser has never been flown in space, and is risky and expensive. The closest equivalent instrument is the ICESat 2 Lidar, which has a 1m aperture but only a 20W laser and cost 200M (for a flagship-class Earth observation mission). If you could go Class D, there is precedent for a Space Station mounted Lidar of comparable laser power, with a 0.6m aperture, that was 50M (CATS).

The recommendation that emerged from the team was to reduce the average laser power by decreasing the pulse frequency, and accept a greater challenge in the mission pointing and ConOps by doing a fast "nod" maneuver over the skylight.

Problem: Surface Knowledge

In order for the instrument concept to work, the visible portion of the cave floor must have known shape, to within the accuracy of the instrument. For 1m or 0.5 accuracy, this exceeds the

resolution of existing data. This can be achieved by partly using the instrument in "Lidar mode." We De-focus the laser, and use a separate detector (or filter and primary detector) to collect first-bounce photons. An operations approach is to possibly use one observation pass just for Lidar. This Lidar data would be of higher resolution and accuracy than any existing Lunar data sets (conceivably to 0.1m).

Technical Questions: Power

The main tradeoff is between a battery-only design or a battery/supercapacitor hybrid. We determined that a Li-ion, battery-only power system should be sufficient for spacecraft and instrument power. Supercapacitors are more efficient and therefore produce less heat; they are also thermally more forgiving; and they can be cycled more. However, this would require a mass penalty; supercapacitor/battery hybrid design would probably about double the mass of the power subsystem. For a very high power laser, the need to be careful not to damage optics from heat is present. For cooling, usually lasers use ammonia. For such a short burst, our experiment might be able to get away with using paraffin wax.

Technical Questions: Optics

Mirror size: How big can the mirror be? 1m is common; 2m would be beyond what has flown on an instrument and so is unlikely to be acceptable for a mission that is already using an entirely new imaging modality. Better to use a 1-meter mirror.

Question: Can you use a deployable mirror? A fixed mirror is generally lighter and lower-risk. Recommendation: If you can fit your mirror in the launch fairing, use non-deployable mirror.

Regarding mirror quality: lambda/4 would be sufficient. Other recommendations include Filter: Eliminate Earthshine; 0.1nm (full-width, half-power), 20% loss in filter Laser: Single wavelength, 532nm, need to lock laser frequency.

Technical Questions: Pointing (Feasibility of Nod Maneuver)

Can we get enough torque to spin fast enough? (~5 deg/s) If you use the biggest reaction wheels in the catalog, and you don't have any appendages off of the spin axis, you could be in the ballpark of feasibility. We must keep "nod" axis moment of inertia under 10 kg*m^2 Could use control moment gyros, like on Space Station.

An unorthodox approach was discussed: slowly spinning up the spacecraft like a wheel, end-over-end, and timing it just right to hit the skylight as it passes. Pointing experts discouraged this; it is very unconventional and would be extremely hard to hit the skylight with dead

reckoning and orbit determination. Could use "Terrain Relative Navigation" to determine the time to start the roll, but this would add cost.

Two instrument concepts (low-power "ICESat analog" and high-power) were developed with the following characteristics:

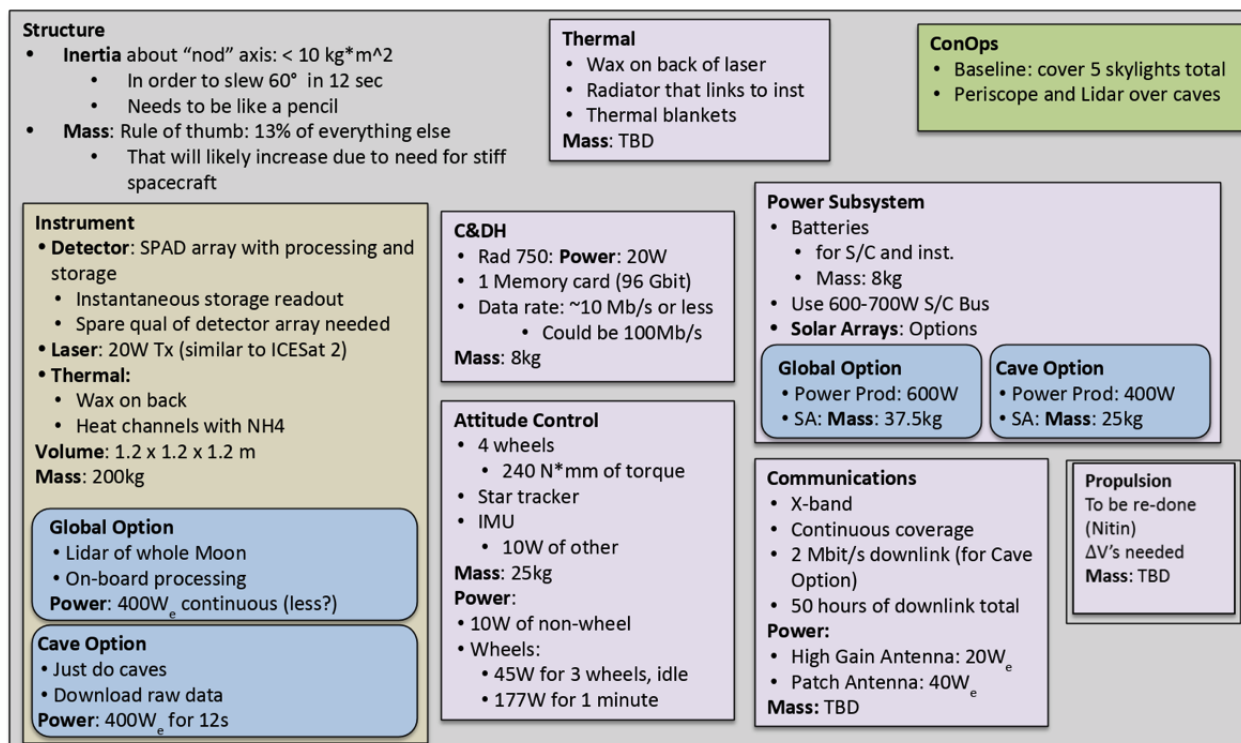*Table 2: Low-power instrument concept*

	Low-Power (ICESat Analog)
ConOps	Pass 1: Lidar • Defocus laser Pass 2, 4, 5: "Nodding" • Start obs @ 60° (30° from vertical) • 5% nod rate • Variable roll rate-- ≈ sin wave Pass 3: Lidar w/ nodding
Observation Time	12 s
Laser	Min 2.5 W @ 10 kHz (May require 10W, 20W preferred) • Can get a 10-W laser off the shelf, but need a shorter pulse length ns --> 20 ps • Can be defocused for lidar
Optics	1 m aperture
Detectors	Separate detector or lidar

Pointing	• Reference: ICESat-2: 13.3-arcsec control, 3 sigma
	• Cross-track control: 1.0 mrad (0.05 deg) (206 arcsec)–total roll-up of boresight error
	   • Need to calibrate boresight: 68 arcsec? Comes out of 206
	• Along-track control: 18 m from OD (*fancy* OD)–1 h before --> 0.5 mrad (350 arcsec) reqt total roll-up of boresight error
	• Using dead reckoning with OD, gets very hard
	   • Terrain Relative Naigation (TRN) is one solution
	   • --> 1 m? --> 1.5 mrad reqt
	   • Lower risk, but expensive and not flown yet
	• Required slew rate is ~5°/s; but star-trackers cut out at 3°/s
	   • Initial estimate: order of magnitude higher torque than any in database --> could use CMGs (see "Mission Concept" sheet for more info)
	• "Rolling Wheel" option – not how things are typically done, scary

Communications/ Ground System	• For OD: One pass per orbit
	   • One-hour passes? Or just keep it on; probably want continuous
	   • If using TRN, might not need every orbit
	• Comm: 10-W tx, 20 We <?>
	• X-band–25 orbits of data --> 50 h of downlink (D/L)
	• ~2 Mbit/s D/L upper limit

Power	Instrument: 400 W when operating

Table 3: High-power instrument concept

	High-Power
ConOps	Pass 1: Lidar • Defocus laser Pass 2-5: Nadir-pointing
Observation Time	0.03 s (50 m skylight)
Laser	1 kW @ 1 MHz • Can be defocused for lidar
Optics	1 m aperture
Detectors	Separate detector or lidar
Pointing	• Similar accuracy constraints? • Similar motivation for TRN? • No problems with slew rate
Communications/ Ground System	(Same)
Power	(Same)

The overall mission concept includes the following systems:

Table 4: A-Team mission concept summary (2016)

Structure
- **Inertia** about "nod" axis: < 10 kg*m^2
  - In order to slew 60° in 12 sec
  - Needs to be like a pencil
- **Mass**: Rule of thumb: 13% of everything else
  - That will likely increase due to need for stiff spacecraft

Thermal
- Wax on back of laser
- Radiator that links to inst
- Thermal blankets
Mass: TBD

ConOps
- Baseline: cover 5 skylights total
- Periscope and Lidar over caves

Instrument
- **Detector**: SPAD array with processing and storage
  - Instantaneous storage readout
  - Spare qual of detector array needed
- **Laser**: 20W Tx (similar to ICESat 2)
- **Thermal:**
  - Wax on back
  - Heat channels with NH4
Volume: 1.2 x 1.2 x 1.2 m
Mass: 200kg

Global Option
- Lidar of whole Moon
- On-board processing
Power: 400W$_e$ continuous (less?)

Cave Option
- Just do caves
- Download raw data
Power: 400W$_e$ for 12s

C&DH
- Rad 750: **Power**: 20W
- 1 Memory card (96 Gbit)
- Data rate: ~10 Mb/s or less
  - Could be 100Mb/s
Mass: 8kg

Attitude Control
- 4 wheels
  - 240 N*mm of torque
- Star tracker
- IMU
  - 10W of other
Mass: 25kg
Power:
- 10W of non-wheel
- Wheels:
  - 45W for 3 wheels, idle
  - 177W for 1 minute

Power Subsystem
- Batteries
  - for S/C and inst.
  - Mass: 8kg
- Use 600-700W S/C Bus
- **Solar Arrays**: Options

Global Option
- Power Prod: 600W
- SA: **Mass**: 37.5kg

Cave Option
- Power Prod: 400W
- SA: **Mass**: 25kg

Communications
- X-band
- Continuous coverage
- 2 Mbit/s downlink (for Cave Option)
- 50 hours of downlink total
Power:
- High Gain Antenna: 20W$_e$
- Patch Antenna: 40W$_e$
Mass: TBD

Propulsion
To be re-done (Nitin)
ΔV's needed
Mass: TBD

A-Team Study Conclusions

No critical physics-based flaws of the design were identified. High laser power (1-5kW) was identified as a major challenge, as it is two orders of magnitude beyond what has flown. A solution is to

1. reduce laser power (to 10-20W),
2. reduce laser pulse repetition rates (1 MHz → 10 kHz)
3. Adopt a "nod" maneuver over target to extend observation time and get needed pulse number (i.e. photon count).

Closest analogy: ICESat-2, flies in a few years (from time of this study), >$100M. If it can fly as Class D, some precedent in CATS.

Path forward involves significant testing in lab and terrestrial environments. Lab tests with small scale structures are described below in this report, and the team plans a field visit in May 2018 to a real cave. Next steps will include flight testing using helicopter, aircraft, and drones.

Team X Activity

We had proposed to conduct a Team-X activity to cost and detail a PERISCOPE-based mission in full detail. This was abandoned, because internal and external variables identified during A-Team study were too massive for detailed Team-X costing models. Instead, the effort focused more on minimizing these variables.

Geometry Reconstruction Method

We have developed a method that describes a non-line-of-sight (NLOS) imaging experiment using a virtual camera projected onto the relay surface that images using virtual waves that are waves of intensity fluctuations propagating on an optical carrier wave. In the following we derive the properties of these phasor field waves and show that they propagate according to the same law as electromagnetic waves. This means that the NLOS imaging problem can be re-formulated as a line-of-sight (LOS) problem. Any possible NLOS camera has an exact mathematical analogous imaging system in optical, microwave, and ultrasound imaging. We can thus solve all non-line-of-sight image reconstruction problems using readily available algorithms for image reconstruction and processing in any of the line of sight imaging fields.

Phasor Field Model

An example of NLOS optical experimental setup is shown in Figure 3. The transmitter is a laser with nanosecond to picosecond long pulses. For the receiver, we consider an ultra-fast camera,

such as a single photon avalanche diode (SPAD) [references 1,2]. The setup also comprises a relay wall, with a diffusely reflecting, ideally Lambertian surface. The light pulse generated by the laser travels to the relay wall arriving at point p and scatters from the wall in all directions; part of the photons reach the target, and a fraction of these travel back to the wall. The ultra-fast camera is focused on q and measures the light flux reflected at q as a function of time. The acquired data is a 7-dimensional space, because it is a function of the 3D coordinates of p and q, as well as a function of the time, t. The problem of reconstructing a 2D or 3D image of the hidden object is an inverse light transport problem.
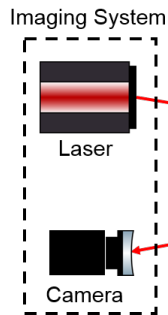
Figure 2: Traditional imaging system
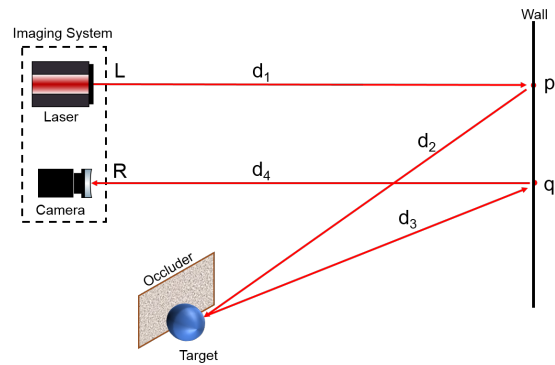
Figure 3: A generalized Time of Flight, Non-Line-of-Sight (NLOS) imaging system

The light transport theory models the propagation of light through a scene (see, for example, [3]) and allows us to infer about the scene by analyzing the data captured by the camera [4].

Previous approaches have used ray optics and attempted to model the light propagation through the scene as a linear operator that can be inverted with a variety of well-studied inverse methods [5-15]. If posed in this way, the reconstruction problem is only approximately linear for very simplistic scenes. The ray optics model also poorly captures the underlying physical light transport processes. Nonlinear inverse methods for more complex scenes have been proposed, but the added level of complexity makes their application challenging. Model complexity and an inaccurate modeling of real light transport also make it challenging to conduct more fundamental studies. There are unanswered questions regarding null-spaces, attainable resolutions and contrast, how to deal with multiple reflections in the hidden scene, the role of the bi-directional reflectance distribution function (BRDF) of the surfaces in the scene, and relation between reconstructions and scene complexity.
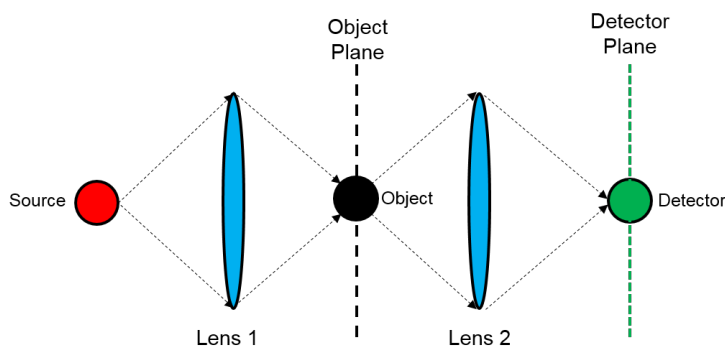
Figure 4: Unfolded schematic of the wave path in a line-of-sight (LOS) imaging system. For simplicity the light is transmitted through the object rather than reflected from it. A real imaging system may operate in transmission (e.g. a microscope), but normally uses reflections.

In this following section, we introduce a new signal processing framework that ties NLOS TOF light transport to conventional physical light transport by introducing a new quantity, the Phasor Field (or P-Field) to describe an NLOS imaging system, such as one shown in Figure 4, as a conventional imaging system that is projected onto the relay wall. We show that the data collected by an NLOS imaging system is fundamentally equivalent to the data collected by a conventional LOS imaging system, such as a camera, ultrasound, or microwave imaging system that is placed at the relay wall and is observing the hidden scene directly. The parameters of this proxy imaging system, including illumination strength, sensitivity, aperture shape and size, and wavelength are computed from the properties of the NLOS imaging system used. This approach, therefore, provides an intuitive understanding of the methodology and capabilities of NLOS imaging. Our method models the entirety of the physical light transport process as a linear, time invariant system. The motivations for this approach are as follows:

- Make predictions about the best possible performance of a given NLOS imaging system.
- Translate insights for efficient imaging system design and methods from conventional imaging systems.
- Address fundamental questions about the NLOS reconstruction process by re-formulating them as questions about conventional imaging processes, which have been extensively studied.
- Make use of the insight that NLOS TOF light transport, as described by our model, is fundamentally a linear, time invariant process to analyze the NLOS reconstruction process and develop optimal reconstruction methods.

The core idea of our contribution relies on defining the P-Field. This quantity is the complex envelope of optical irradiance, for which we define an amplitude and a phase, analogous to the complex envelope of the E-Field. A phasor representation of radiance was discussed in [16], where the authors propose a framework to analyze the light transport in correlation-based TOF ranging. We introduce a different phasor representation that is linked to physical light transport in two different ways: (1) A P-Field wave propagates very similarly to an E-Field wave, and (2) within our assumptions the P-Field propagator models the physical light transport process. This allows us to link the limitations of our reconstructions to fundamental physical limitations.

Non-line-of-sight imaging with virtual Phasor field cameras

The wave equation for time-harmonic E-Fields or magnetic fields, known as the Helmholtz Equation, describes the propagation of energy in the form of an electromagnetic wave in any linear, isotropic and homogeneous medium. The Helmholtz Equation is satisfied by all orthogonal scalar components of the E-Field. The Green's Function-based solution of the wave equation for the E-Field is given by

$$E(x,y) = K'_E \int E_0(x',y') \frac{e^{jK|r|}}{|r|} dx'dy'.$$

(2)

If

$$\mathcal{E} = E_0(x',y') \frac{e^{jK|r|}}{|r|},$$

(3)

represents an E-Field wavelet contribution from each location (x',y'), then (2) is expressed as

$$E(x,y) = K'_E \int \mathcal{E} dx'dy'.$$

(4)

The quantity |r| is the absolute distance between single location (x',y') in A and (x,y) in Sigma, $E_0(x',x')$ is the amplitude of the E-Field at location (x',y'), K is the wavenumber of the E-Field and $K_{E'}$ is the E-Field coefficient of proportionality. This is also known as the Huygens-Fresnel principle [17,18] and explains the transfer of fields from any generic location (x',y') inside plane A to an individual location (x,y) on another plane Sigma separated by a distance z from A. In other words, the Huygens-Fresnel integral comprehensively describes the physical process of light transport rendering, the propagation of an E-Field distribution or a wavefront from a light source through a scene to a camera, inverse rendering, and image formation (which is the propagation of a wavefront from the aperture of a camera back into the scene). The Green's function propagator $e^{jK|r|}/r$ is most-commonly known as the normalized Huygens E-Field wavelet.

In this paper, we model a temporally modulated light source and camera as emitters and detectors for a virtual wave of P-Field fluctuations of optical carrier irradiance; i.e.,

$$P_0 e^{i\omega_P t} = \left(\int_T |E|^2 e^{i\omega_E t} dt \right) e^{i\omega_P t}.$$

(5)

In (5), the Electric-Field (E) is integrated over at least one full cycle of its oscillation period. The integral in the parentheses describes the light intensity or photon number as it would be seen by a fast detector. The fluctuations of the intensity are described as an oscillation with frequency ω_p. Any intensity profile can be described as a superposition of phasor field waves – intensity waves with amplitude P_0. We show that the propagation of the complex P-Field amplitude is analogous to E-Field propagation and is described by a propagator P which is analogous to the Huygens-Fresnel propagator E for E-Fields and the P-Field distribution P(x,y) in Sigma is given by

$$P(x,y) = K_P \left| \iint_A P_0(x',y') \frac{e^{j\beta|r|}}{|r|} dx'dy' \right|.$$

(6)

The quantity

$$\mathcal{P} = P_0(x',y') \frac{e^{j\beta|r|}}{|r|}$$

(7)

is a P-Field wavelet contribution from location (x',y') in A with an amplitude P_0(x',y'), K_P is the P-Field proportionality coefficient and beta is the associated P-Field wavenumber. B is expressed in terms of the P-Field wavelength λ_P, the corresponding P-Field frequency ω_P and the refractive index of the medium of propagation n between A and Sigma as

$$\beta = \frac{2\pi}{\lambda_P} = \frac{n\omega_P}{c}.$$

(8)

Equation (6) is sufficient for many NLOS imaging scenarios where light reflects off diffuse surfaces, phase information in the E-Field component is lost, and the resulting intensity variations from the E-Field contributions are small resulting in carrier irradiance that is approximately uniform in space and time. In such cases, where the planes considered in the model are reflective or transmissive diffusers, the P-Field propagator can be used on its own.

If the scene involves non-Lambertian surfaces or partial coherence effects such as speckle, or if coherent detectors are used, the carrier propagation makes a significant contribution to the overall result and has to be modeled. In this case we show that if the coherence length of the carrier is shorter than the P-Field wavelength, the full propagator for the modulated light wave can be written as the product of unmodulated carrier propagator and the P-Field propagator. This, in turn, allows us to express the magnitude |I(x,y)| of the irradiance distribution I(x,y) in Sigma to be expressed as

$$|I(x,y)| \propto \left| \iint_A \mathcal{P}dx'dy' \right| \times \left| \iint_A \mathcal{E}dx'dy' \right|^2.$$

(9)

In the remainder of this section, we derive the properties stated above and show some examples of the applications of this model. We leave the exploration of all the implications of the model and its use in experimental NLOS reconstructions to future publications. We present a brief overview of imaging with the Huygens-Fresnel integral and its evolution into much simpler

Fresnel and Fraunhofer integrals under certain approximations. Using this formulation, an NLOS imaging system can be described as a virtual LOS imaging system on the relay wall and any existing LOS imaging method can be implemented for this virtual imaging system.

Application Examples

In a traditional line-of-sight (LOS) scenario, as the one shown in Figure 2, we consider an imaging system, composed of a light source and a camera and a point-like target, which we consider our object under investigation. The source emits a monochromatic wave that travels to the object, passing through a lens, whose goal is to refocus the electromagnetic wavefront onto the object (we assume that the object is located on the focal spot of the lens). After the light interacts with the object, it travels to the camera, once again passing through a lens which focuses the diffuse light into the receiver.

The Phasor Field virtual wave approach can be used to describe any time-based lensless imaging system, such as the ones presented in [20]. Our primary interest is however in its use for virtual camera projection in NLOS imaging as illustrated in Figure 3. The objective is to capture an image of the scene via a diffuse reflection at a relay wall. The imaging system can project temporally modulated intensities on the relay wall and detect incoming intensities that strike the wall. The wall thus becomes the aperture of a holographic P-Field projection and detection system. After detection of the P-Field at all points on the aperture any conceivable imaging system can be realized through digital post processing. A simple imaging lens, for example, applies a position dependent phase delay to the signal followed by a summation of the fields on the camera pixels and a time integral over the absolute value to reconstruct a 2D image of a scene.

Through a P-Field imaging approach a reflective structure such as a wall can be treated as a 'non-rough' P-Field aperture, similar to a P-Field lens, where despite the loss of spatial coherence of the E-Field optical carrier upon reflection, the P-Field contributions are preserved, enabling us to treat NLOS imaging as conventional LOS imaging, but in the realm of P-Fields instead of E-Fields of the optical carrier. As the effect of the surface roughness is minimal to the P-Field phase as compared to the phase of the optical field, we can effectively treat any rough surface in a multi-bounce NLOS imager as a P-Field aperture of negligible roughness. Consequently, this allows us to describe an NLOS imaging system as a LOS P-Field imaging system.

$$P(R) = \iint_Q P(q)e^{j\phi_l(q)+j(\frac{R-q}{ct})}dq,$$

(10)

where $P(R)$ is the P-Field of the on the image plane, R, $P(q)$ is the P-Field at the relay wall at point q, Q is the set of relay wall points, and ϕ_l is the phase shift applied by the computational lens which in a real lens is achieved by traveling through different amounts of glass. The details of this reconstruction operation and the design and demonstration of a virtual camera projection

system are beyond the scope of this work. Here we present some insights about the properties of this virtual imaging system that can be derived from the phasor field formalism.

Light source and detector design

To project a light source onto the relay wall a collimated laser is usually used to create a small spot that can be treated as a P-Field point emitter. The spacing between these point emitters can be derived from phased array theory. Phased antenna arrays require an antenna spacing of $\lambda/2$ to be able to project and detect wavefronts of arbitrary shape without artifacts. Consequently the spacing of laser positions on the wall should be about half the P-Field wavelength λ_P. The situation for the virtual detector positions is similar. However, an ideal imaging system would also have to maximize the total area of relay wall that is sampled, i.e. maximize the fill factor of the virtual camera. The detector array is thus composed of closely spaced area detectors of side length $\lambda/2$.

While the local density of camera and projector spots should be $\lambda/2$, parts of the relay wall regions can be left out without affecting the image quality. This is analogous to blocking part of the aperture of a camera. As long as it does not decrease the aperture diameter such blocking only lowers the signal strength but has very little effect on the image.

Resolution

It is important to realize that image quality is fundamentally only limited by the signal-to-noise Ratio (SNR). Resolution of the final image does not have to equal the resolution of the imaging system itself. Common methods of 'super resolution' work by changing the definition of resolution or by changing the properties of the scene. The resolution of an imaging system is typically defined as the size of the systems airy disc which is a measure of the highest spatial frequency that can be faithfully resolved by the imaging system. The distance from the central maximum to the first minimum of the airy disc in the far-field is given by the Rayleigh resolution limit

$$p_E = 1.22 \frac{\lambda_E r_E}{d_E},$$

(11)

where p_E is the size of the resolved patch, r_E is the distance between the patch and the aperture, λ_E is the light wavelength, and d_E is the diameter of the aperture. Applying the P-Field principle we can thus estimate the resolution of a virtual camera as

$$p_P = 1.22 \frac{\lambda_P r_P}{d_P}$$

(12)

where p_P is the size of the resolved patch, r_P is the distance between the patch and the aperture, λ_P is the virtual wavelength and d_P is the diameter of the aperture. The Rayleigh limit can be derived using the Rayleigh-Sommerfeld diffraction integral. To formally derive it for P-Fields, one simply has to replace the quantities in that derivation by the corresponding P-Field quantities.

Reconstruction examples

Figure 5 shows an example of an NLOS reconstruction performed with the phasor field method. The photograph on the top right shows a corridor and the main image shows a top down view of the reconstruction. We are able to reconstruct the corridor with approximately 10 cm accuracy which is in agreement with our resolution estimates.
A further description of the reconstruction method is provided in a manuscript that is currently under review (https://arxiv.org/abs/1810.07535).

A current limitation of the experimental setup is due to the low number of data points we can collect with our single pixel sensor. In simulations we can use array sensors and simulate a dense sampling of the relay wall yielding significantly better results. Figure 8 shows the reconstruction of a 2m by 2m by 2m large office scene. In this case we sample a grid of points with 1 cm spacing on the relay wall. We are currently developing array detectors as part of a different program.
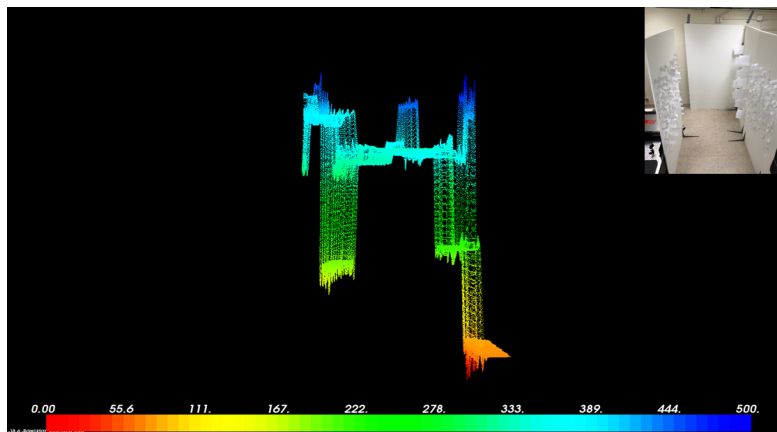
Figure 5: The photograph on the top right shows a corridor and the main image shows a top down view of the reconstruction.
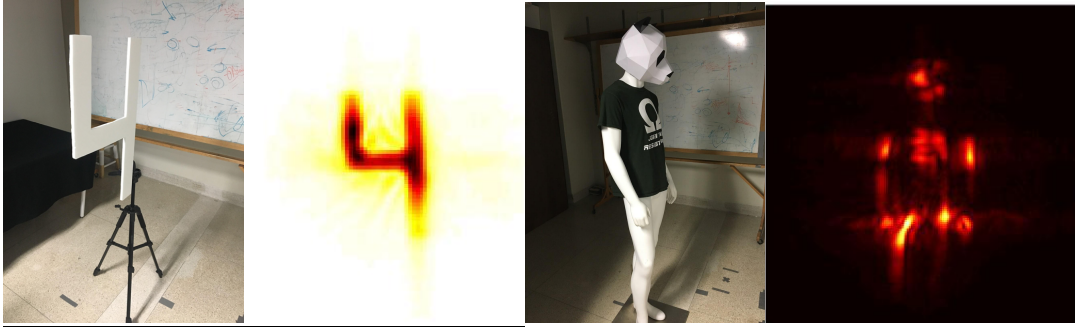
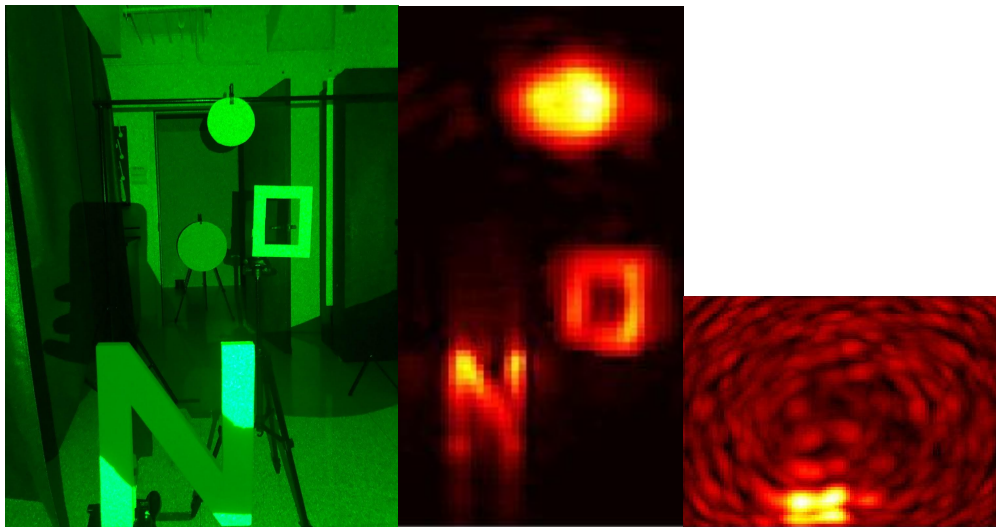Figure 6: Further examples of reconstructions of physical scenes.

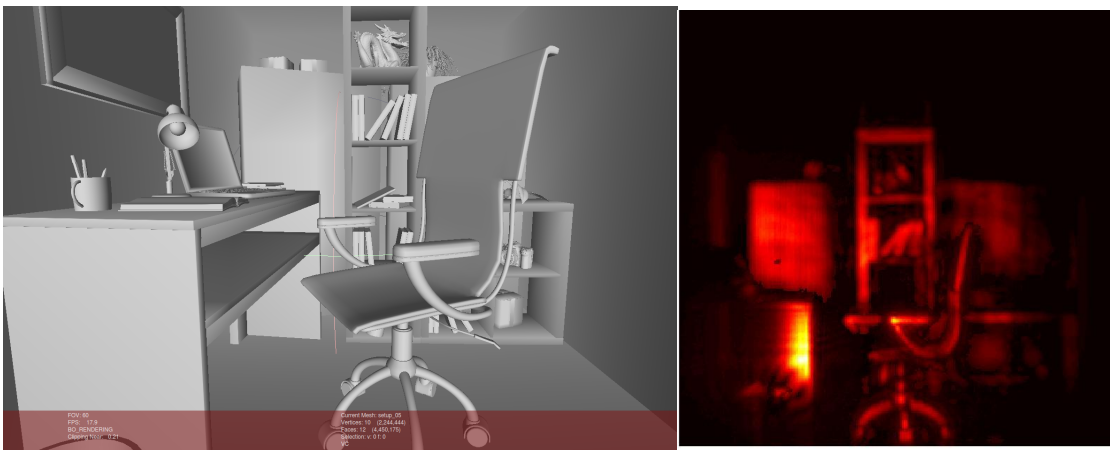Figure 7: Reconstruction of cardboard letters "N" and "O" using the phasor field formalism.

Figure 8: Simulated reconstruction of an indoor scene. The scene model is shown on the left from a 3d graphics program, and the 2d projection of the reconstruction is on the right.

Portable, self-calibrating imaging system

We have developed a portable optical system for NLOS. The novelties are as follows:

- To select a suitable relay wall (Fig. 3), we utilize the data recovered by a lidar unit and a color CMOS camera;
- To recover the exact position of the laser and SPAD on the relay wall (cf. points p and q in Fig. 3), we triangulate the information from a stereo camera setup, composed of two black-and-white CMOS cameras.

System Description
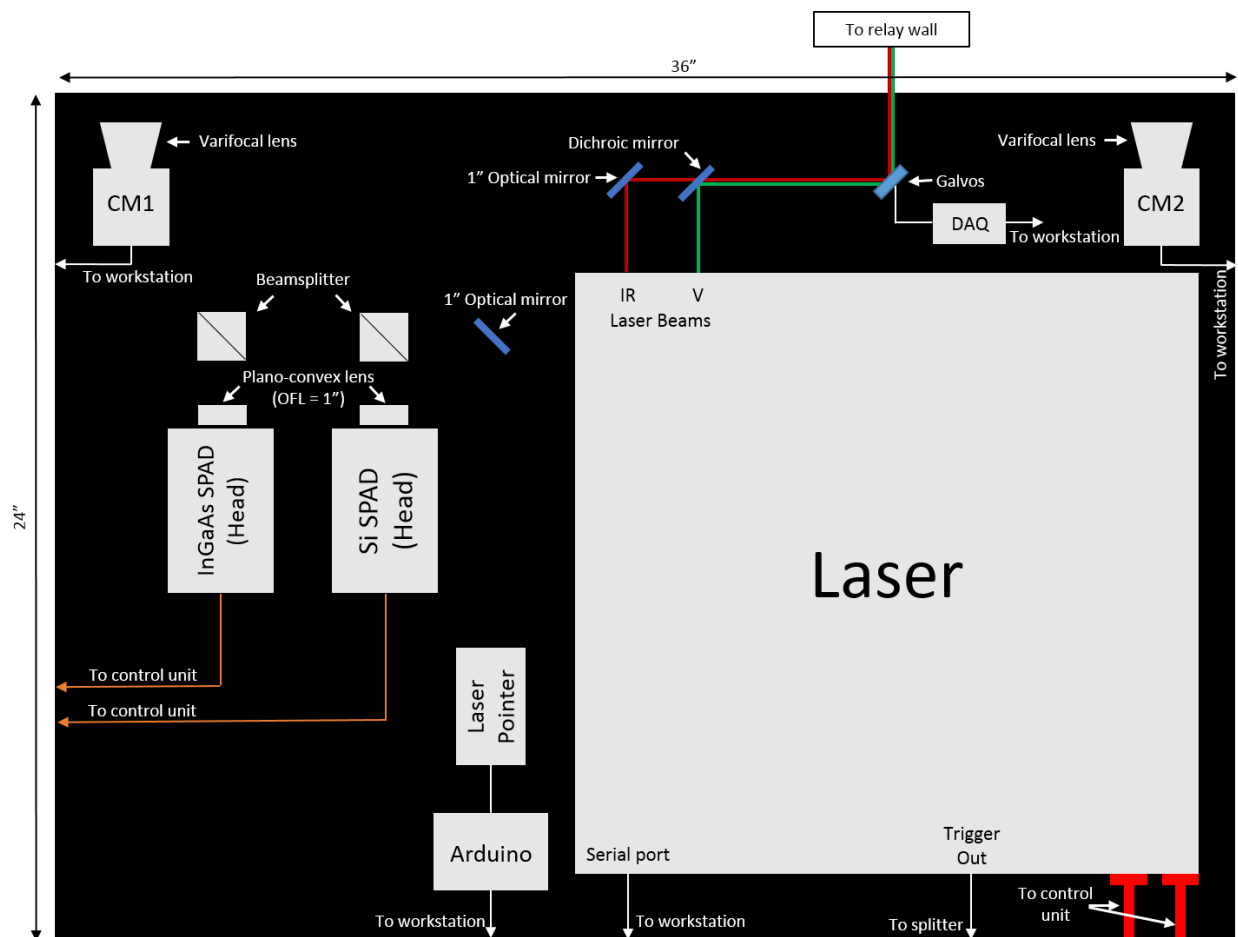
Figure 9: Schematic of optical system

Recall that the goal of the system is to recover a 3D image of a target hidden behind an occluder, as shown in Fig. 3, by illuminating different points on a relay surface (wall), considered an ideal Lambertian scatterer. In Fig. 9, we show the proposed portable NLOS imaging system, which uses ultra-fast optical equipment. It is mounted on top of a 36" x 24" optical breadboard,

clamped to a movable cart. We can subdivide the NLOS imaging system into two subsystems: the first one is responsible for the visible scene inference, while the second one is for the hidden scene inference. The former is composed of a lidar and a color CMOS camera, whereas the latter comprises a picosecond pulsed and tunable laser, two SPAD (ToF cameras), a programmable picosecond delayer, a time correlated single photon counting (TCSPC) unit, two galvanometer actuated mirrors (galvos) and two monochrome CMOS cameras. The equipment is connected to a workstation that allows us to modify the devices' parameters, acquire data and automate the data acquisition process.

Visible Scene Inference

In order to identify a suitable surface that acts as our relay wall, the portable NLOS system is equipped with a lidar scanning system and a color CMOS camera. In a laboratory setup, it is always possible to create an artificial and suitable surface. This can in fact be achieved using different type of materials or special diffuse coating paints (see [b16]). On the other hand, in an outdoor scenario, the lidar and CMOS camera approach allows to analyze all possible relay walls and select the best for the specific purpose. It also helps in determining whether certain surface areas of the relay wall should be avoided, because of their morphological characteristics (i.e: avoid holes, specular areas, and so on).

The CMOS camera is a Sentech STC-MBA5MUSB3 [b30], which is composed of a 1/2.5" CMOS sensor with a resolution of 5 Megapixels, connected to our workstation through a 3.0 USB connection.

Hidden Scene Inference

To infer about the hidden scene, we exploit ultra-fast optics equipment. As a light source, we utilize a Onefive Katana 10 HP fiber laser [b31], whose operating wavelengths are λ_{IR} = 1064 nm and λ_V = 532 nm. The laser emission pulses are 35 +/- 15 ps long and the laser's pulse repetition frequency (PRF) is tunable in the 0.05 - 10 MHz range. The beams are overlaid on top of each other using a combination of a 1" optical mirror and a longpass dichroic mirror with a 650 nm cut off wavelength. The Katana HP also produces a trigger output, a pulse of 2.5 V amplitude and the same PRF of the laser. This signal is available at Trigger Out, with a nominal delay of 24 ns (w.r.t. the laser beam pulse); the nominal jitter is <2 ps. As will be explained later, the laser's trigger output is employed as a synchronizing signal.

Let us now consider Fig. 3. Once the laser illuminates a generic point p on the relay wall, the photons scatter in all directions. Part of these photons reach the target(s) and, assuming no inter-reflection or sub-surface scattering on the target(s), a fraction of these photons return the relay surface on point q. We assume that the SPAD, our ToF camera, is focused on q, a point on the relay wall. More specifically, we focus the SPAD on a known 1 cm^2 area (whose center is q) by placing in front of the ToF camera a 1" diameter lens with a 25.4 mm focal length. We

assume that the detector's field of view (FOV) throughout a data acquisition time does not change.

Since the light source is capable of emitting pulses at two different wavelengths, the portable NLOS system is equipped with two different SPADs, developed by SPADLab (Milan, Italy). In order to detect photons at λ_V, we use a Silicon (Si) SPAD [27] and its specs are summarized in Tab. 5; the Si SPAD looks at a point q_1 on the relay wall. To decrease the background noise arriving at the detectors, we employ bandpass filters. In front of the Si SPAD, we place a bandpass filter with a peak transmission at λ_V and full width at half maximum (FWHM) bandwidth of 10 nm. For λ_{IR}, we exploit an Indium-Gallium-Arsenide / Indium Phosphite (InGaAs/InP) SPAD [b26, b32], whose most important characteristics are provided in Tab. 1; we assume that this SPAD is focused on a point q_2 on the relay wall. In front of the InGaAs SPAD, we place a bandpass filter with a peak transmission at λ_{IR} and full width at half maximum (FWHM) bandwidth of 10 nm. Note from Fig. 9 that there are two beamsplitters, located in front of the SPADs.

Table 5: Silicon SPAD Specifications

	Si	InGaAs
Sensor diam., $(\varnothing)$ (μm)	50	25
Photon Eff. (%)	24	15
Time Jitter (ps)	50	150
Hold-off Time (s)	$50 - 500$ n	$1 - 3000$ μ
Dark Count Rate (cps)	20	< 10k
Afterpulsing prob. (%)	< 1	29

When a photon arrives at the SPAD's detector head, namely where the active chip is located, it triggers an internal electron avalanche. This has two effects on the SPAD: 1) it becomes `blind' to subsequent photons for a specific amount of time, called hold-off time, and 2) after an internal delay, a NIM Logic signal is generated. This signal has a low logic level of 0 V and a high logical of -800 mV. Moreover, the NIM signal is available at the Photon Out port of the SPAD's control unit (cf. Fig. 9).

To register the time-of-arrival (TOA) of the photons, we use a PicoHarp 300 TCSPC unit [33], which produces a histogram of the photon counts as a function of time. This device has two inputs and are required to be in the [0, -800] mV range (the ideal suggested range is [-200, -400] mV [b34]). Fig. 10 summarizes how the laser, SPAD and TCSPC are interconnected. Note the presence of a delayer, which is explained in the following sections. We connect the Katana's Trigger Out to the TCSPC's Channel 0, with an appropriate inverter and attenuator (for our purposes, we use a PicoQuant "SIA 400 - Inverter and attenuator module"); for safety, we also add another 10 dB attenuator before connecting it to the TCSPC. The InGaAs or Si SPAD's Photon Out is attached to the TCSPC's Channel 1. Independently of the SPAD, this connection also includes a 10 dB attenuator. Since this TCSPC unit only accepts two inputs, we manually attach the Si or InGaAs SPAD Photon Out to the TCSPC's Channel 1, depending on the data that we want to acquire.

The TCSPC's histogram bin size can be selected between 4, 8 and 16 ps; for our purposes, we use 4 ps, the unit's maximum resolution.
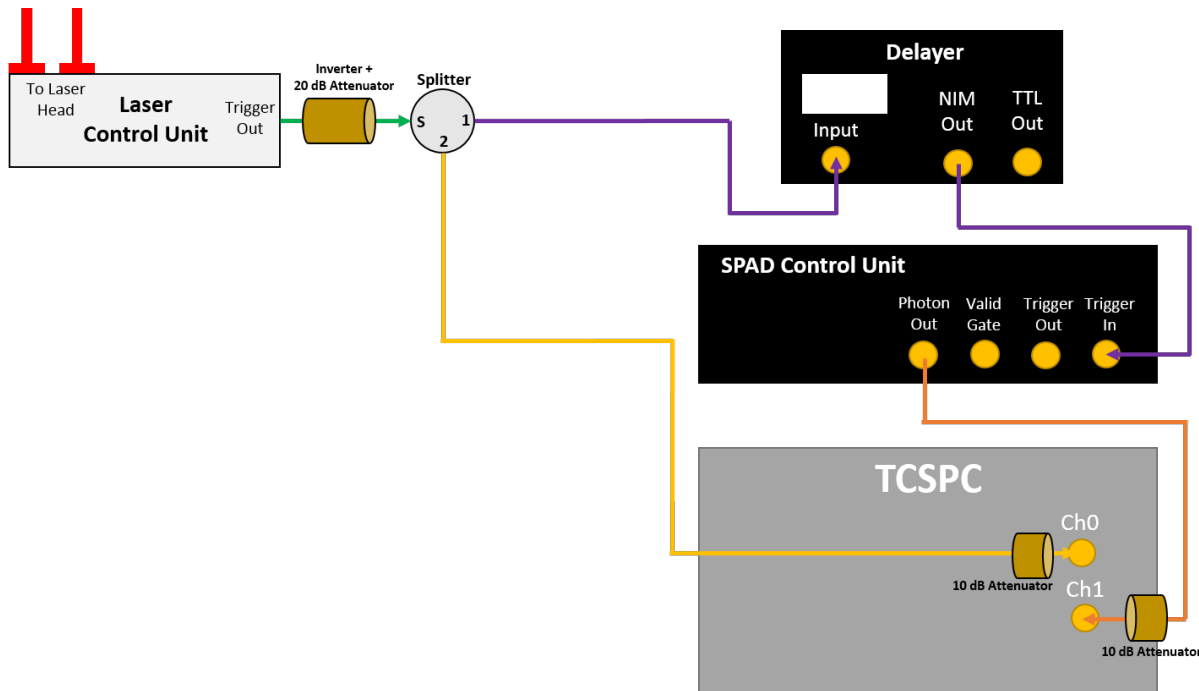
Figure 10: System Electronic Wiring

System Setup

The previous section provided an overview on the ultra-fast optical equipment and how it is used to acquire data. In the following subsections, we explain some advanced features that are fundamental to acquire a meaningful dataset.

Laser operations

The Katana's crystal can be operated at different temperatures. To get an efficient output of λ_V, the temperature needs to be set in the [145.75,146.40] degree C range, depending on the selected PRF. For our measurements, we have used PRF = 10 MHz and we set T_V=145.80 degree C. Using these parameters, both laser beams are available at their respective output channel (cf. Fig. 9). Furthermore, the laser's pump current can be adjusted in the [0, 10] A range. At the maximum current, the laser can reach a power of P_{IR} = 2 W and P_V = 1 W for λ_{IR} and λ_V, respectively.

The Katana laser parameters can be adjusted using an RS232 (serial) connection between the laser and the workstation and a series of terminal commands. To execute the required commands, we use CoolTerm, a serial port terminal freely available. The specific connection parameters (baud rate, data bits, parity check, etc...) and commands are included in the Katana's operations manual.

Gating the SPAD

Our SPADs can be used in two modes, denominated free running mode and gated mode. In the former, the device is always kept on; namely, photons are detected when they arrive in the chip's active area, unless the SPAD is currently in its hold-off time window. This mode is suitable for testing purposes, as well as for situations where the signal of interest is small. The latter mode is called gated mode and fundamentally allows the SPAD to time-gate out photons. The gated mode is recommended when the SPAD needs to be selectively turned off, particularly useful when strong light signals may arrive near the photon's TOA of the hidden scene. The gating technique can be interpreted as a way to decrease the dark count rate (DCR), afterpulsing probability, and consequently increasing the overall Signal-to-Noise (SNR) ratio. It can be controlled through the software that is provided by the developer, since the SPAD's control unit is connected to the workstation through a USB connection.

To successfully use the gating feature of the SPAD, the laser's Trigger Out is connected to the TCSPC's Channel 0 as well as to the SPAD's Trigger In, through the use of a signal-splitter. The SPAD's Photon Out is still connected to the TCSPC's Channel 1, as described above. The laser emits a laser pulse and, after a delay, it also sends out an electrical pulse from the Trigger Out. When this signal is recognized by the SPAD, with a delay of τ_T, the gating window turns on for a period of time, T_{ON}. Photons that arrive at the SPAD's active area, while the window is on, generate an avalanche. After a delay of τ_A, the NIM signal is available at the SPAD's Photon Out. In Fig. 11, we provide the visual explanation of what we have just discussed. The exact values of τ_T and τ_A depend on the specific length of T_{ON} and whether the synchronizing signal is external (i.e.: coming from the laser) or internal (i.e.: using a synchronizing signal internal to the SPAD). Note that photons arriving while the gating window is off do not trigger an avalanche signal.

Recall from the previous discussion that we want to use the gating feature to filter out the unwanted signals and focus on the signals coming from the targets. The SPAD's software allows the user to control the gate width T_{ON} and its repetition rate (cf. Tab. 6), which should be set to the laser's PRF. However, the software cannot shift the gating window over time. To do so, we use a picosecond delayer, specifically a picosecond delayer, capable of time-delaying the signal by [50, 5000] ps [35]. This device allows us to shift the gating window over time. From a hardware point of view, we connect the laser's Trigger Out to the delayer's Input and the NIM Output to the TCSPC's Channel 0. With some back of the envelope calculations, as well as using the free-running mode and a dummy Lambertian target, it is possible to identify the TOA region of the hidden scene and therefore select the correct values to be inserted into the delayer. If we need a delay greater than 50 ns, we can connect the laser's Trigger Out to the delayer's Input using a longer SMA cable (a 1 m long cable roughly delays the signal by 5 ns); of course, this method can be used as long as the cable attenuation does not significantly destroy the signal.
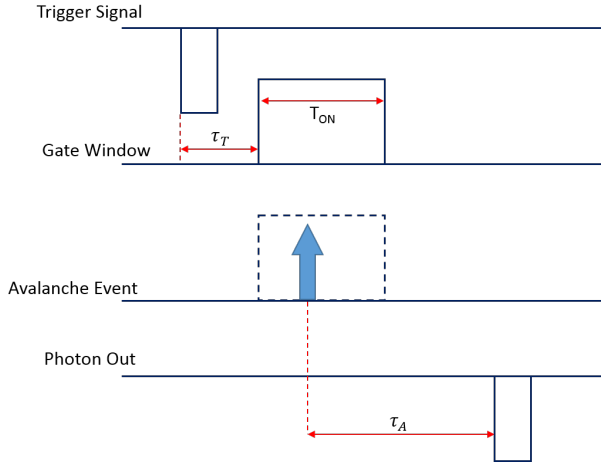
Figure 11: Trigger timing

	Si	**InGaAs**
Width, T_{ON} (ns)	$2 - 500$	$10 - 500$
Rise Time (ps)	$100 - 200$	$700 - 1500$
Rep. Rate (MHz)	$0 - 80$	$0 - 133$

Table 6: Gate details

Scanning Multiple Laser Positions

The reconstruction algorithm relies on a dataset acquired using multiple laser positions on the wall. To steer the beam towards different spots on the wall, we use a galvo system and a data acquisition device (DAQ), specifically a National Instruments USB-6001 DAQ [36]. The galvos are two mirrors placed orthogonally from the other, allowing us to move the beam on the (x,y) plane. Each galvo can span a +- 20 degree FOV, by applying to them an analog input signal in the [-5, 5] V range. We generate two digital signals on our workstation, which are then converted to analog by the DAQ. The galvos are connected to their power unit and the DAQ through electrical wires. In turn, the DAQ is connected to the workstation through a USB port. To create a bi-dimensional laser grid on the wall, it is necessary to provide the analog signals corresponding to each location, p_m, m=1, ..., M, on the wall. The galvos however introduce spatial inaccuracies, therefore the laser may not be pointing to p_m, rather to p'_m, a point in the vicinity of the wanted one. In the reconstruction algorithm, we use the p'_m (m=1, ..., M) values and, to find these, we use a calibrated stereo camera system that is targeting the relay wall.

We use our laboratory setup as an example to explain how we extract the laser wall points. The stereo camera system is composed of two monochrome STC-MBA5MUSB3 Sentech cameras, mounted at a given inter-distance. Using variable focal length lenses, the CMOS cameras are focused on the relay wall and are looking at it at an angle. From the images, we can note the lens distortions at the edges of the pictures. To correct for lens distortion, we first have to calibrate the cameras and then we can use a set of predefined MATLAB functions.

We exploit MATLAB's camera calibration functions, based on [37, 38] and developed by the Computational Vision group at the California Institute of Technology (CalTech). To calibrate the cameras, we use a checkerboard poster composed of 19 x 18 black and white squares, whose side length is 5.20 cm on the relay wall. We proceed in capturing 15 sets of images of the checkerboard poster, where we show a subset of the acquired images; note that, for each image, the checkerboard occupies a different area of the cameras' FOV and we also change its orientation (w.r.t. to the cameras), in order to capture the lens distortions.

After the image acquisition phase, we remove the checkerboard pattern and run the calibration algorithm, which returns the intrinsic and extrinsic parameters of the stereo system. The intrinsic parameters comprise the internal parameters of the cameras, namely the principal points, focal length and skew. Conversely, the extrinsic parameters are external to the cameras and change depending on the scenario.

Using these parameters, we can remove the lens distortions, exploiting MATLAB's undistort image function.

Once the stereo system is calibrated, we proceed with the data acquisition phase. For each laser point on the wall p_m (m = 1, ..., M), we take a picture of the laser spot from the left and right cameras; we remove the lens distortions from these images and locate the spot position, expressed in pixel coordinates. Afterwards, we use MATLAB's triangulate function, to determine the point's 3D world coordinates, considering the left camera position as the origin, i.e.: O = (0,0,0).

In the end, we remark that the calibration procedure can be also done in a non-laboratory setup, depending on whether the relay wall is accessible or not. If the relay wall cannot be reached (i.e: unstable terrain between the imaging system location and the selected relay wall), a workaround consists in first estimating the distance between the portable imaging system and the relay wall (for example, through the on-board lidar), and then finding a suitable location where to place the system and the checkerboard at the estimated distance. The next step is to run the calibration algorithm as described above and finally reposition the portable NLOS system towards the desired relay wall. If the relay wall is accessible, one can repeat the procedure described above.

SPAD focus point on wall

The reconstruction algorithm also needs the SPAD focus position on the wall (cf. point q in Fig. 3). To locate it, we use a sub-system composed of a CW 50 mW, 532 nm laser beam (controlled by an Arduino Uno [39]), a 1" optical mirror, and two beamsplitters, each placed in front of the SPADs.

As it can be seen in Fig. 12, the beam travels from the laser pointer to the mirror, from which is redirected towards the first beamsplitter. Part of the light is reflected towards the relay wall and points at q_1, namely the Si SPAD focus location on the relay wall. A fraction of the beam travels to the other beamsplitter; again, a fraction of the light is reflected towards q_2 (where the InGaAs SPAD is focused), whereas the remaining part travels through and is blocked by a beam dumper. Using the stereo-camera system, we take a picture of these two points and extract their 3D coordinates.

During the data acquisition phase, the laser pointer is turned off, but the beamsplitters are left in their position. The majority of the photons scattered from the hidden scene travel through them and arrive at their respective SPAD, even though a small fraction is reflected, as shown in Fig. 12.
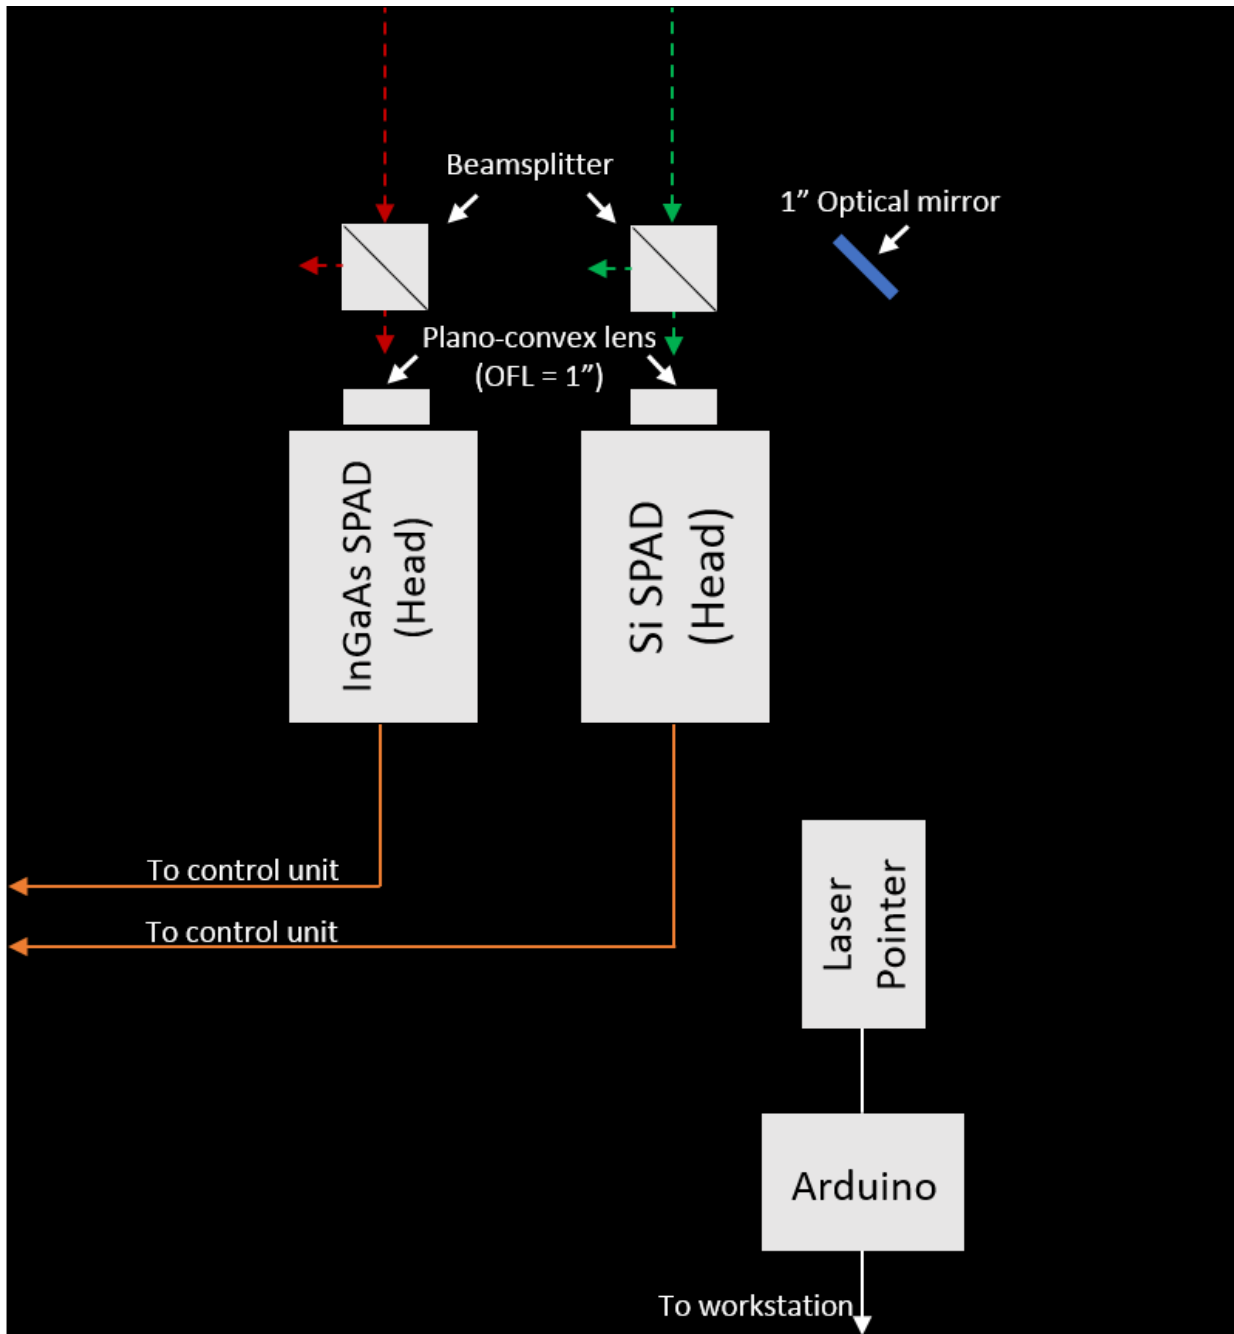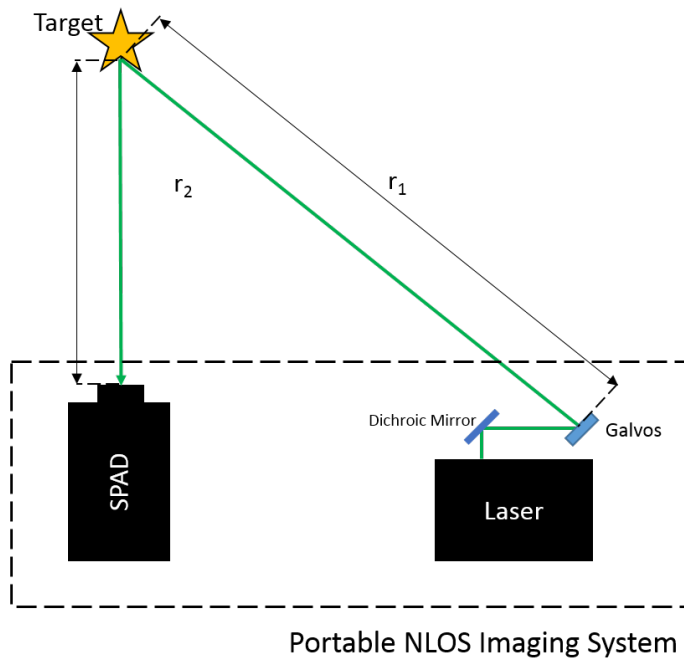

Figure 12: Calibration Laser Setup

In the end, this operation can be included in the automatic data acquisition process, since we have programmed the Arduino microcontroller to turn on and off the laser pointer, as necessary.

Portable NLOS Imaging System

To measure t_0, we use the simple scenario shown in Figure 13. In the resulting histogram we locate the time bin corresponding to the maximum peak, namely where the dummy target is located (cf. red circle) and measure the difference $t_0 = | t\_p - t\_m |$.

Recall from the previous discussion that the TCSPC generates a photon count vs. time bin histogram using the laser's Trigger Out as a synchronizing signal. This signal arrives at the TCSPC with a time delay that depends on the imaging system's electronics, as well as the length of the cables. The spec sheets of the equipment specify the minimum, maximum and typical delay values. However, to get an accurate measurement of the current system's delay, we proceed as follows.

Data Acquisition

A flowchart of the data acquisition process is shown in Fig. 14. Since we want to acquire a dataset considering multiple M laser wall positions and since M tends to be in the order of a few thousand points, we use an in-house LabVIEW application to automate the acquisition process (Fig. 15). Assuming that the hardware is turned on and correctly connected to the workstation, the following steps are key to a successful data acquisition (for each number, refer to the equivalent box in Fig. 15):

- Define where and under what name the dataset is going to be stored.
- Check if the TCSPC is working correctly
- `Initialization status' should read Hardw. ver.: 2.0; namely, the TCSPC is turned on and the program recognized the 2.0 hardware version,
- `Calibration' should read Calib. O.K.; namely, this operation internally calibrates the TCSPC's hardware,
- `Bin size' should read 4, namely the histogram's bin size 4 ps.
- Check that the application correctly loaded the file where the galvo positions are stored
- the box should display the total number of laser positions.
- Set the constant fraction discriminator (CFD) for the TCSPC's Channel 0 and Channel 1

- the CFD emits a logic pulse only when the signal coming from the considered channel reaches a specified percentage of its maximum
- if the selected thresholds are too high, the two vertical bars in the bottom 4 box do not show any number; if this is the case, the CFD value should be reduced, until a signal can be seen.
- Set how much time delay (in ms) there should between the the galvo moving to its new position and the data acquisition
- While moving to the new position, the galvos inevitably experience spatial jitter; waiting for a specific amount of time avoids acquiring data while the galvos are settling.
- Select the desired acquisition (exposure) time (in ms)
- higher exposure times generally increase the SNR, at the expenses of an increased total acquisition time
- The specific acquisition time is dependent on the scene and needs to be evaluated on a case-by-case basis
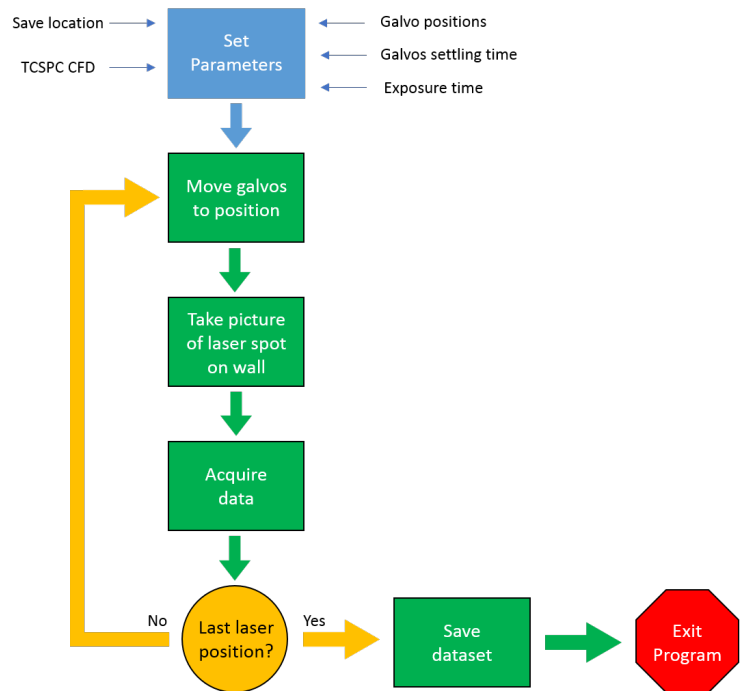
Figure 14: Flowchart for automated data acquisition.
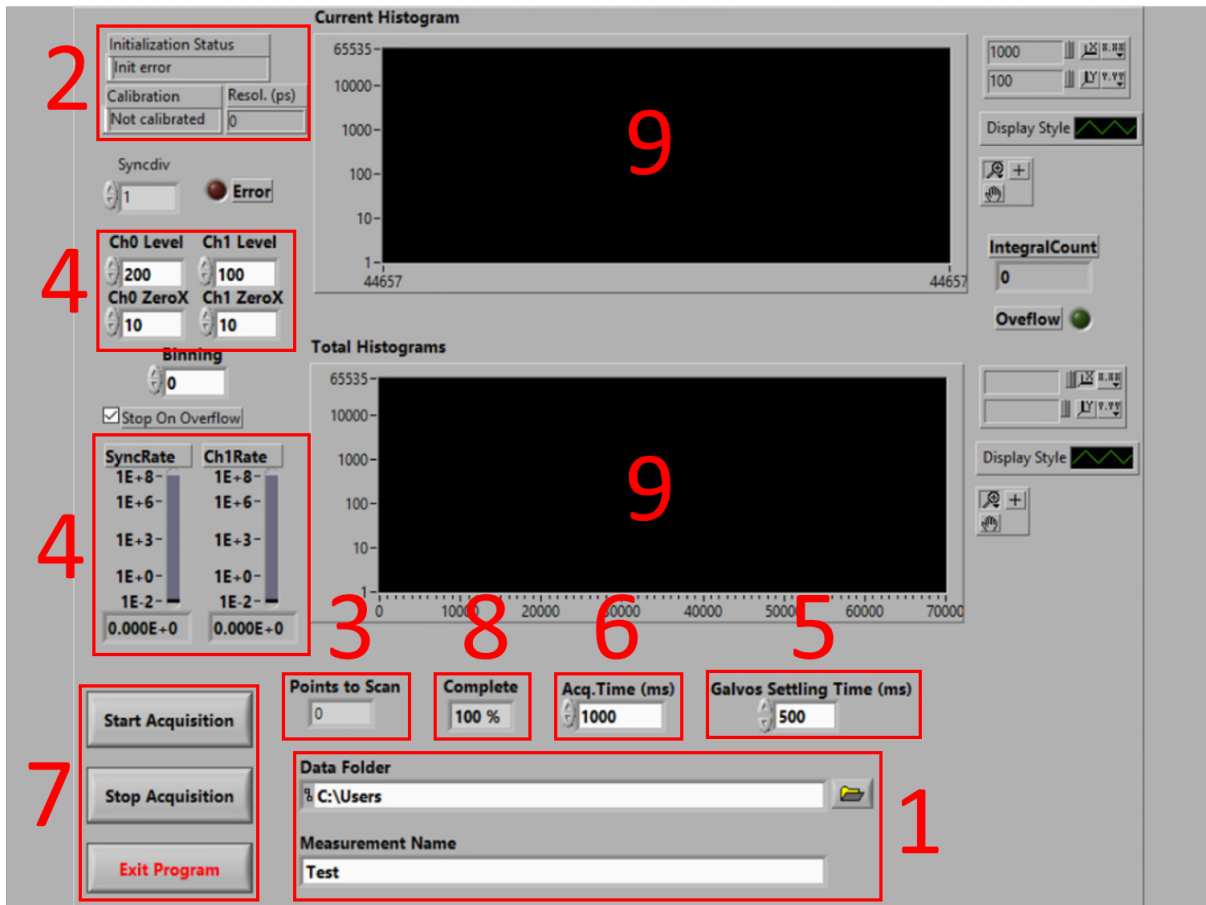
Figure 15: LabVIEW interface for automatic data acquisition.

Once the application has been set up, it is possible to start the data acquisition by pressing the Start Acquisition button (cf. box `7' in Fig. 15). The process can be monitored by the display box in `8', which shows the percentage of laser positions completed against total number of laser positions. The entire dataset is saved into a single binary file (".dat" extension).

Post-Acquisition Processing

Format

With an appropriate MATLAB script, we organize the data into a structure array (``.mat'' extension) that contains the following

- Data acquired from the LabVIEW application, organized as an M x N x T cube, where M is the number of laser positions on the wall, N is the number of SPAD positions on the wall (in our case, N=1) and T is the length of the histogram,

- 3D coordinates of the laser origin (L in Fig. 2); in our case, this corresponds to the galvo position
- Laser wall 3D coordinates, obtained from the calibrated stereo-camera system
- Norm vector of the laser wall positions
- 3D coordinates of the SPAD, namely where the ToF camera's active chip is located (R in Fig. 3)
- SPAD wall 3D coordinate, obtained using the laser pointer - beamsplitter system, and the stereo-camera system
- Norm vector of the SPAD wall positions
- System's internal delay, t_0
- TCSPC's histogram bin size

Laboratory Test Conclusion

Here, we have described a system composed of a picosecond laser that emits two different wavelengths, λ_{IR} = 1064 nm and λ_V = 532 nm, can reach up to P_{IR} = 2 W for λ_{IR} and P_V = 1 W for λ_V and its PRF can be tuned. Since the light source emits two wavelengths, we use a Si SPAD to collect the λ_V photons and an InGaAs SPAD to collect the λ_{IR} photons scattered from the hidden scene. This system is operational and currently being used in the lab.

Simulation engine

We developed an improved rendering engine to generate high resolution physically realistic simulations of time of flight light transport to create accurate computer models of the imaging experiment. This rendering engine is continuously improved and is currently running on the UW Center For High Throughput Computing supercomputer.

New Simulations - Synthetic Scene Model

As a first step in this project, while a new rendering and reconstruction pipeline was under development, we created updated renderings using the older rendering and reconstruction with an improved laplacian filter. Figure 16 shows a reconstructed point cloud (green) together with the ground truth (red). In the results presented in Phase 1 we were not able to reconstruct the walls of the cave and only saw an obstacle placed in the cave. In this reconstruction the upper walls are clearly visible while the floor is still not detected. There are two reasons for this:
1. The floor receives less light from the relay surface and the returned signal is much weaker. This makes the floor hard to reconstruct especially in the presence of the much stronger signal from the ceiling.
2. In this hand drawn cave model, the floor is perfectly smooth and has no features. A certain degree of roughness on the cave surface is required for the reconstruction algorithm to work well. This is most easily understood using the phasor field virtual wave model introduced earlier. According to this model the imaging process can be modeled as a virtual camera and light source positioned at the mount in the cave that projects and detects a phasor field wave - a wave that exists as a modulation on the intensity of the

light - to create an image of the cave. The wavelength of this phasor field wave is about 2 cm. A surface that is smooth on this size scale will act as a specular reflector to this virtual wave. In the model used here the cave floor is perfectly smooth. It acts as a perfect mirror to the phasor field wave and none of it is reflected back to the relay surface. This means that even if light from this surface returns to the relay surface, the modulation on the light intensity (i.e. the virtual wave) does not. The cave floor just returns a constant signal.
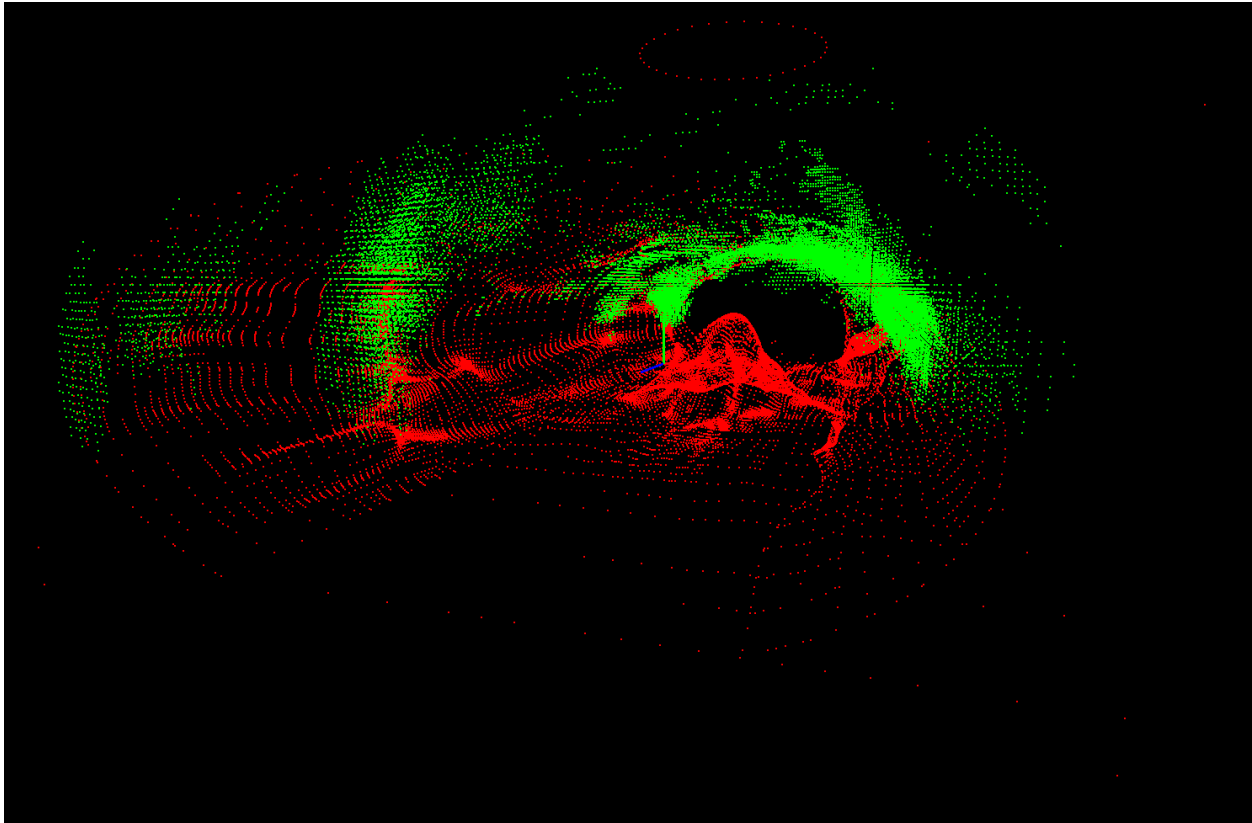
Figure 16: Reconstructed point cloud (green) together with the ground truth (red).
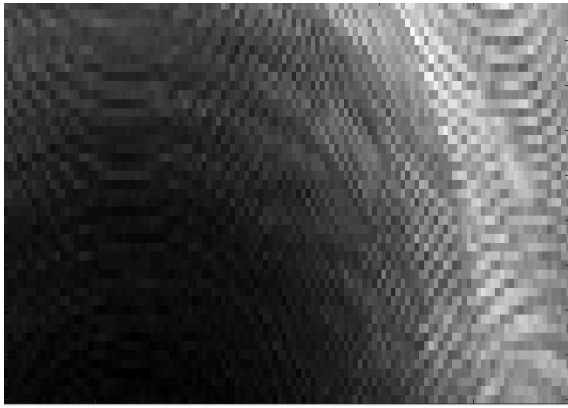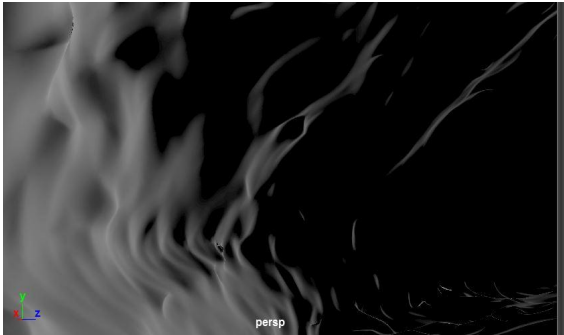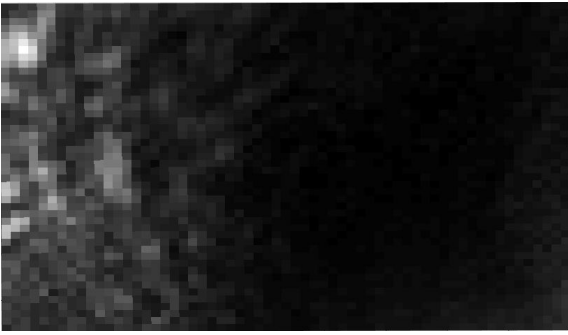
Field Test Campaign

We collected a LiDAR scan of Big Skylight cave in El Malpais National Monument, NM to provide a more realistic geometry for our computational model. We have converted the LiDAR point cloud to an OBJ file that can be loaded by our rendering engine, and have generated simulated results using this model. Reconstructions are shown below. We are able to recover albedo information from the scene with high resolution. Currently we use a 60 cm phasor field wavelength, resulting in a resolution over ten times worse than what our experimental system can do. Higher resolutions both in experiment and simulation are technically possible, but are hard to do on our current hardware due to the larger amount of data that has to be collected or simulated.

Reconstruction Results

Table 7 below shows reconstructions of different features in Big Skylight cave. The left column shows the ground truth acquired with a LiDAR system and the right our reconstruction. While there are clearly challenges with dynamic range when displaying the results, we can see that most features that are visible in the direct view are preserved in the reconstructions. Resolution could be further improved at the cost of additional computation, but is not currently limited by hardware capabilities. In these simulations we do not include realistic noise. We have however found in the past that monte carlo noise from the simulations is comparable to noise collected in the experiments. Realistic noise is added in the following Table.

Table 7: Example reconstructions

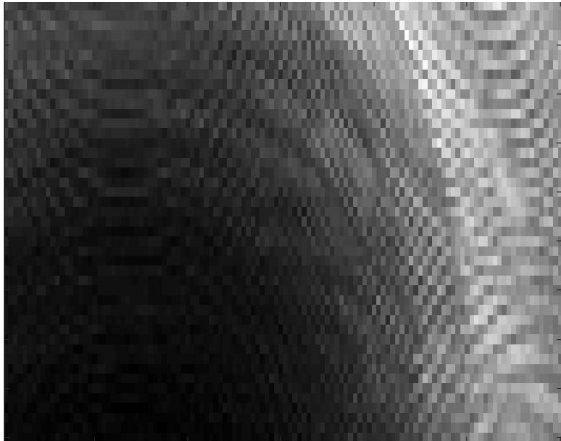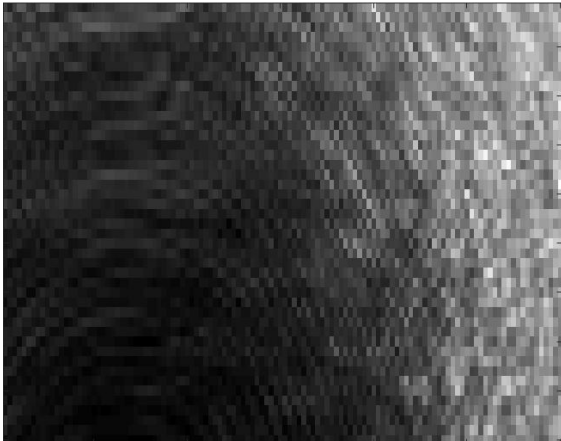|	Ground Truth	Reconstruction
a		
b		
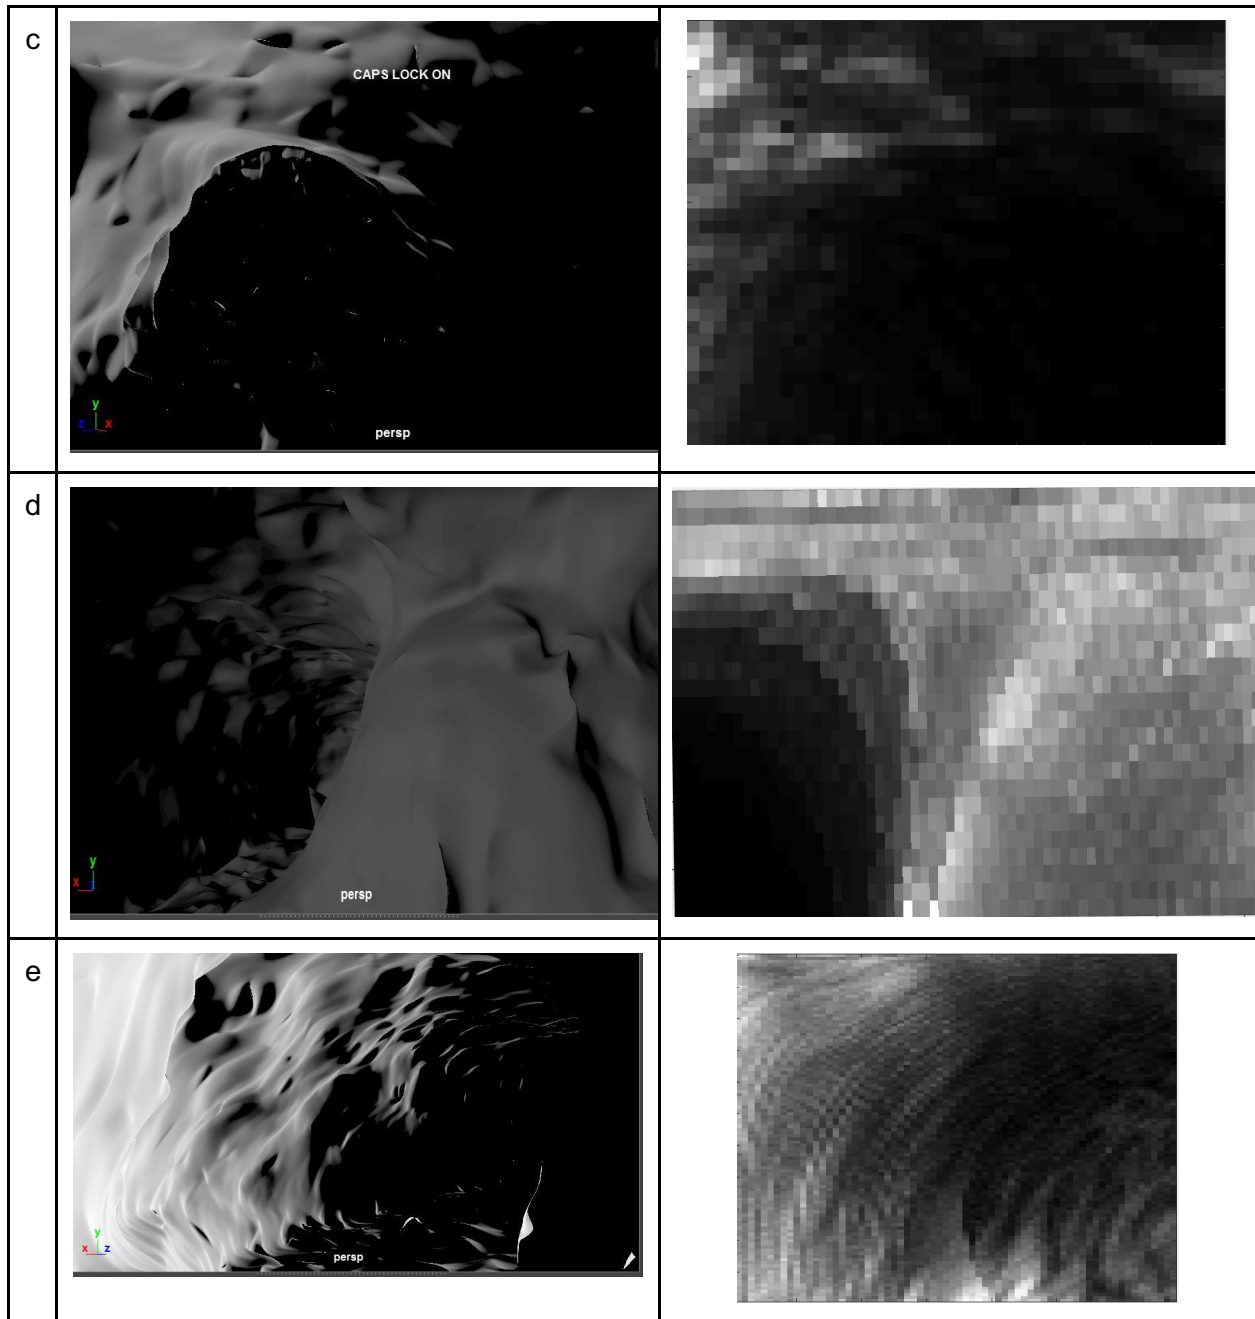
First results of simulations using LiDAR data of big skylight cave. (left column) Different views of the inside of the cave as seen from the relay surface under the cave entrance. The data is rendered using Blender from the LiDAR data of the cave. A single light source at the location of the relay wall illuminated the scene. (right column) Our reconstruction of the scene albedo for approximately the same view.

Table 8 below shows two of the reconstructions with realistic noise assuming a real imaging scenario. We assume a 2 watt average power laser illuminating a single spot in the cave for 1 second. The detector is a 128 by 128 pixel spad array with a 100% fill factor. Wall reflectance in the invisible cave is 1% and 8% on the visible cave floor. The cave is assumed to be illuminated by direct sunlight. The sunlight is partially removed by a filter with 0.1 nm passband which can be achieved with an unstabilized fabry-perot etalon. In this configuration, The data is dominated by poisson noise and contributions from sensor related noise and ambient light are negligible. The relevant signal is between 1 to 100 protons per time bin strong before the addition of noise which is in agreement with our preliminary estimates included in the proposals for phase 1 and phase 2 of the program.

While there is a notable degradation in image quality, main features of the cave remain discernable. This is in agreement with our notion that our reconstruction algorithm is extremely robust to noise, even without special de-noising strategies. We believe that significant further improvements are possible with better denoising. Beyond this the judicious use of prior assumptions can further improve reconstructions. There are a number of prior assumptions that can likely be made with great confidence. For example we can assume that our reconstruction is a surface, rather than a three dimensional volume. This would only be violated in the presence of strong subsurface scattering from semitransparent materials or by strong participating media like dense fog in the scene which are unlikely to be encountered on the moon. While using these priors appears to be unnecessary, they may allow us to significantly improve reconstruction quality or relax hardware requirements in future iterations of the design.

Table 8: Reconstruction noise examples

No Noise	Noise
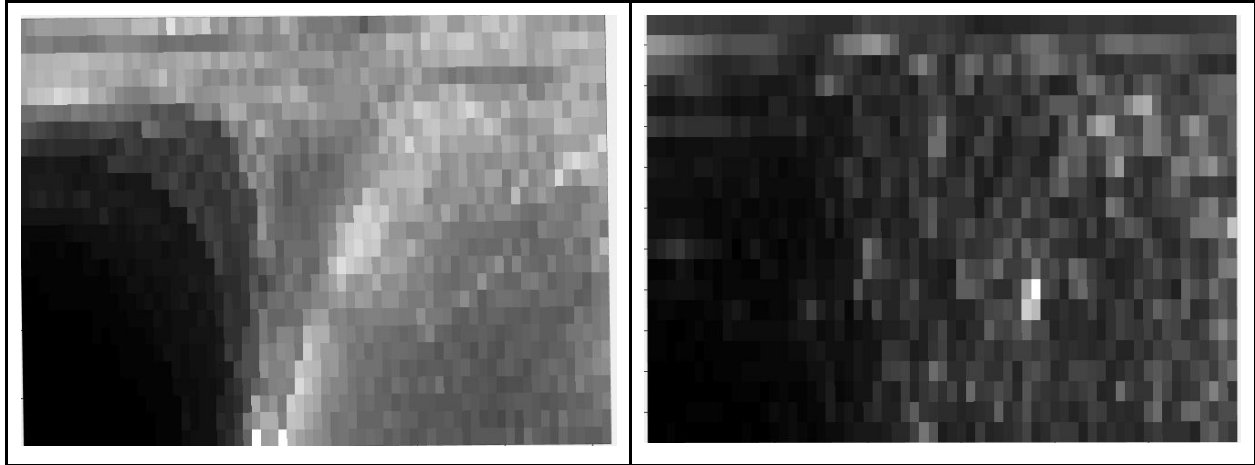
Reconstructions of sections (a) and (d) from above with a realistic amount of sensor noise assuming 5 km altitude, 1 s total exposure, 5% wall reflectance, and 1 Watt laser power, and direct sunlight.

Laser Effect on Vegetation

In preparation for taking our imaging system to field trials we checked the effect of our laser on vegetation in the cave. This is required to obtain a permit for imaging in El Malpais National Monument, where caves contain endangered vegetation. We obtained a species of moss similar to what is expected in the cave and exposed it to our laser system at close range. As the images below show, the moss suffers no visible effects from the laser light.
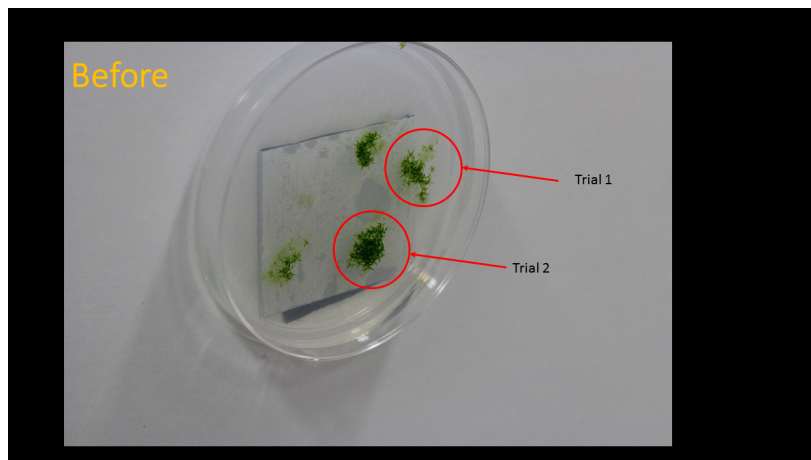

Figure 17: Samples before illumination

Figure 18: Samples after 10 s illumination at maximum power at 532 nm at 2 m range. No visible effect on the plant immediately after illumination or after 24 hours (shown).
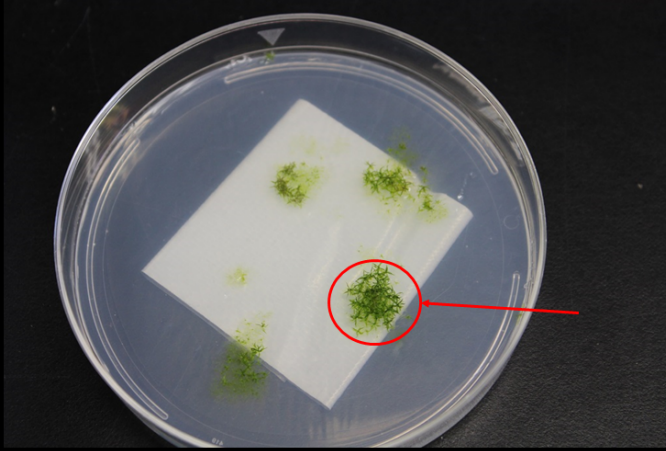

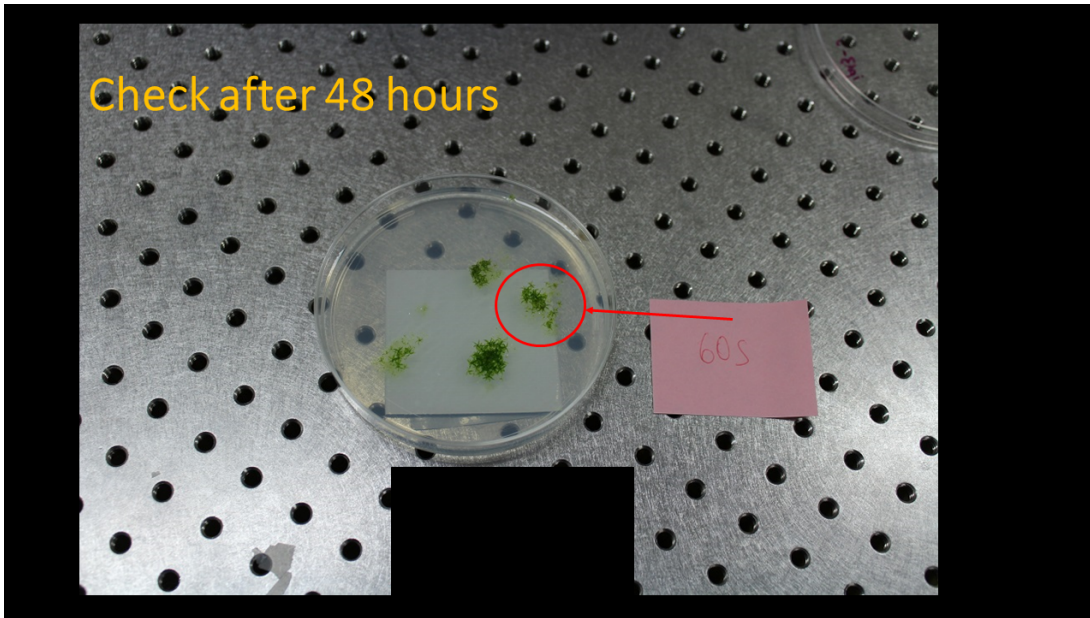*Figure 19: Samples after 10 s illumination at maximum power at 1064 nm at 2 m range. No visible effect on the plant immediately after illumination or after 48 hours (shown).*

Ultimate System Performance

There are several limitations to the performance of the final system: The time resolution of the laser and detector, the spatial resolution of the detector, the sensitivity and noise floor of the

detector, and the amount of ambient noise due to sunlight. As is explained in detail in our attached manuscript, the trade-offs and performance limitations of the method can be understood as the properties of a virtual camera and light source located at the bottom of the cave. Our physical imaging system probes the wavefront of a virtual wave at this aperture that interacts with the scene. The size of the virtual aperture and the virtual wavelength determine the performance of our method.

In the simulations shown above we used a relatively large virtual wavelength of about 60 cm. The reason for this is primarily that a higher resolution reconstruction would be computationally more difficult, consuming more time and memory (both reconstruction time and memory requirements grow roughly proportionally with the number of reconstructed 3D voxels), while not providing a lot of further fundamental insight about the method. The resolution is chosen to show the dominant features of the cave and achieve a 1 meter resolution about 10 - 20 meters into the cave.

Time Resolution: Our current SPAD systems and light sources have time resolutions of about 35 picoseconds leading to a combined time resolution of just below 70 picoseconds in our lab experiments. This allows us to generate a virtual wavelength of 4 cm. We thus can relax our requirements to time resolution to meet the resolution requirements.

Spatial Resolution: The sensor spatial resolution has to be such that the wall is resolved to at least 0.25 times the phasor field wavelength. In our case the diameter of our virtual aperture is about 20 meters requiring a focal plane array sensor with at least 128 by 128 pixels. Unclassified SPAD general purpose arrays with this resolution exist, but their fill factor and maximum count rate is far inferior to the single pixel sensors we use in our measurements. We will therefore require development of specialized SPAD array sensors. We are currently developing a suitable 16 by 16 pixel array for a different project and believe our design can be scaled up to at least 128 by 128 pixels.

Signal to Noise Ratio (SNR): Our lab experiments have shown that we are able to collect data with very high SNR (better than the simulation), even with strong ambient light, such as active room lighting or open windows. We have also observed that reconstructions degrade only slowly as noise is added and reasonable reconstructions can be obtained even in very high noise scenarios. As is often the case, definite statements regarding the ultimate performance of the method in situations with challenging SNRs is difficult. The noise itself consists of multiple different components and affects different reconstructions and features differently. It is also expected that we will be able to significantly improve our SNR performance by adding de-noising methods to our algorithm. Other than resolution - which can be described exhaustively in one concise and signal-independent parameter - statements about SNR are therefore more open ended. In this work, we have confirmed that cave surfaces remain visible even at SNR levels consistent with a 20W average optical power system observing the cave for 15 seconds from 10 kilometers away and direct sunlight. Exploring the limits of our method for more challenging scenarios will involve a multi-faceted approach involving large numbers of different

scenes, different measurement configurations, reconstructing different scene aspects, and exploring the effects of de-noising methods.

Overall, our studies have shown that non-line-of-sight imaging of lunar caves from low lunar orbit is possible and practical with minor further engineering advances such as a good mechanism to calibrate pointing of the imaging system at the cave and an improved SPAD sensor array. These further improvements can be implemented in a next level prototype that could be tested in real terrestrial caves.

Special thanks

Yongjae Kwon is a rising senior at the University of Michigan majoring in Information Science. Throughout his three consecutive summer internships at JPL, starting from 2013, he performed feasibility studies and cost analysis for Phase A of the Europa Clipper Mission. He contributed to the PERISCOPE proposal in the formulation stage where mission applications were reviewed. He is passionate about imagery analysis and remote sensing to investigate the effects of climate change.

Ben Ormond is a freshman at UC Berkeley double majoring in physics and computer science. In 2017 he 3d printed a model of the El Mapais cave and presented it to project team members at JPL.

James W. Ashley is a planetary scientist and systems engineer at JPL who was involved with the early imaging and characterization of the lunar mare and impact melt pits, and associated subsurface voids while a postdoc with the Lunar Rreconnaisance Orbiter Cameras at ASU. He assisted with the science background and exploration science content of this report.

Many graduate students contributed to this project under the leadership of Co-I Andreas Velten.

Marco La Manna (lab experiments)
Toan Lee (lab experiments)
Ji Hyun Nam (lab experiments)
Xiaochun Lui (reconstruction algorithm implementation)
Syed Azer Reza (reconstruction algorithm theory)
Jingyao Wu (big skylight cave simulations and reconstructions)
Jess Zeman (Phase 1 experiments and simulations)

References

[1] F. Zappa, S. Tisa, A. Tosi, and S. Cova, "Principles and features of single-photon avalanche diode arrays," Sensors Actuators A: Phys. 140, 103 – 112 (2007).

[2] M. Buttafava, G. Boso, A. Ruggeri, A. D. Mora, and A. Tosi, "Time-gated single-photon detection module with 110 ps transition time and up to 80 mhz repetition rate," Rev. Sci. Instruments 85, 083114 (2014).

[?] M.S. Robinson, J.W. Ashley, A.K. Boyd, R.V. Wagner, E.J. Speyerer, B. Ray Hawke, H. Hiesinger, C.H. van der Bogert, "Confirmation of sublunarean voids and thin layering in mare deposits," Planetary and Space Science. 69, 18-27 (2012).

[3] P. Sen, B. Chen, G. Garg, S. R. Marschner, M. Horowitz, M. Levoy, and H. P. A. Lensch, "Dual photography," ACM Trans. Graph. 24, 745–755 (2005).

[4] S. M. Seitz, Y. Matsushita, and K. N. Kutulakos, "A theory of inverse light transport," in "Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1," , vol. 2 (2005), vol. 2, pp. 1440–1447 Vol.

[5] A. Kirmani, T. Hutchison, J. Davis, and R. Raskar, "Looking around the corner using ultrafast transient imaging," Int. J. Comput. Vis. 95, 13–28 (2011).

[6] A. Velten, T. Willwacher, O. Gupta, A. Veeraraghavan, M. G. Bawendi, and R. Raskar, "Recovering three-dimensional shape around a corner using ultrafast time-of-flight imaging," Nat Commun 3, 745 (2012).

[7] O. Gupta, T. Willwacher, A. Velten, A. Veeraraghavan, and R. Raskar, "Reconstruction of hidden 3D shapes using diffuse reflections," Opt. Express 20, 19096–19108 (2012).

[8] D. Wu et al., Frequency Analysis of Transient Light Transport with Applications in Bare Sensor Imaging (Springer Berlin Heidelberg, Berlin, Heidelberg, 2012), pp. 542–555.

[9] F. Heide, M. B. Hullin, J. Gregson, and W. Heidrich, "Low-budget transient imaging using photonic mixer devices," ACM Trans. Graph. 32, 45:1–45:10 (2013).

[10] F. Heide, L. Xiao, W. Heidrich, and M. B. Hullin, "Diffuse mirrors: 3d reconstruction from diffuse indirect illumination using inexpensive time-of-flight sensors," in "2014 IEEE Conference on Computer Vision and Pattern Recognition," (2014), pp. 3222–3229.

[11] F. Heide, W. Heidrich, M. Hullin, and G. Wetzstein, "Doppler time-of-flight imaging," ACM Trans. Graph. 34, 36:1–36:11 (2015).

[12] M. Laurenzis and A. Velten, "Nonline-of-sight laser gated viewing of scattered photons," Opt. Eng. 53, 023102 (2014).

[13] M. Buttafava et al., "Non-line-of-sight imaging using a time-gated single photon avalanche diode," Opt. Express 23, 20997–21011 (2015).

[14] C.-Y. Tsai, K. N. Kutulakos, S. G. Narasimhan, and A. C. Sankaranarayanan, "The geometry of first-returning photons for non-line-of-sight imaging," in "IEEE Intl. Conf. Computer Vision and Pattern Recognition (CVPR)," (2017).

[15] M. La Manna, F. Kine, E. Breitbach, J. Jackson, and A. Velten, "Error back projection algorithms for non-line-of-sight imaging," MINDS@UW (2017).

[16] M. Gupta, S. K. Nayar, M. B. Hullin, and J. Martin, "Phasor imaging: A generalization of correlation-based time-of-flight imaging," ACM Trans. Graph. 34, 156:1–156:18 (2015).

[17] W. Goodman, Introduction to Fourier Optics (Mc-Graw Hill, New York, NY, 1996).

[18] W. Stutzman and G. Thiele, Antenna Theory and Design (Wiley, Hoboken, NJ, 1998), 2nd ed.

[19] S. J. Koppal, "Lambertian reflectance," in "Computer Vision," (Springer, 2014), pp. 441–443.

[20] D. Wu et al., Frequency Analysis of Transient Light Transport with Applications in Bare Sensor Imaging (Springer Berlin Heidelberg, Berlin, Heidelberg, 2012), pp. 542–555.

[b1] D. D. Ferris, Jr. and N. C. Currie, "Survey of current technologies for through-the-wall surveillance (TWS)," in Sensors, C3I, Information, and Training Technologies for Law Enforcement, ser. Proc. SPIE, E. M. Carapezza and D. B. Law, Eds., vol. 3577, Jan. 1999, pp. 62–72.

[b2] E. J. Baranoski, "Through wall imaging: Historical perspective and future directions," in 2008 IEEE International Conference on Acoustics, Speech and Signal Processing, March 2008, pp. 5173–5176.

[b3] M. Amin, Ed., Through-the-Wall Imaging Radar. Boca Raton, FL: CRC Press, 2010.

[b4] J. T. Kajiya, "The rendering equation," in Proceedings of the 13th Annual Conference on Computer Graphics and Interactive Techniques, ser. SIGGRAPH '86. New York, NY, USA: ACM, 1986, pp. 143–150. [Online]. Available: http://doi.acm.org/10.1145/15922.15902

[b5] S. K. Nayar, K. Ikeuchi, and T. Kanade, "Shape from interreflections," International Journal of Computer Vision, vol. 6, no. 3, pp. 173–195, 1991.

[b6] R. Ramamoorthi and P. Hanrahan, "A signal-processing framework for inverse rendering," in Proceedings of the 28th annual conference on Computer graphics and interactive techniques. ACM, 2001, pp. 117–128.

[b7] R. Ramamoorthi, "A signal-processing framework for forward and inverse rendering," Ph.D. dissertation, Stanford University, 2002.

[b8] S. M. Seitz, Y. Matsushita, and K. N. Kutulakos, "A theory of inverse light transport," in Computer Vision, 2005. ICCV 2005. Tenth IEEE International Conference on, vol. 2. IEEE, 2005, pp. 1440–1447.

[b9] D. Wu et al., Frequency Analysis of Transient Light Transport with Applications in Bare Sensor Imaging. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 542–555.

[b10] A. Jarabo, B. Masia, J. Marco, and D. Gutierrez, "Recent advances in transient imaging: A computer graphics and vision perspective," Visual Informatics, vol. 1, no. 1, pp. 65 – 79, 2017. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S2468502X17300104

[b11] C.-Y. Tsai, K. N. Kutulakos, S. G. Narasimhan, and A. C. Sankaranarayanan, "The geometry of first-returning photons for non-line-of-sight imaging," in IEEE Intl. Conf. Computer Vision and Pattern Recognition (CVPR), 2017.

[b12] A. Kirmani et al., "Looking around the corner using transient imaging," in 2009 IEEE 12th Int. Conf. Comput. Vision, September 2009, pp. 159–166. March 28, 2018 DRAFT24

[b13] ——, "Looking around the corner using ultrafast transient imaging," Int. J. Comput. Vision, vol. 95, no. 1, pp. 13–28, 2011. [Online]. Available: http://dx.doi.org/10.1007/s11263-011-0470-y

[b14] O. Steinvall, M. Elmqvist, and H. Larsson, "See around the corner using active imaging," in SPIE Security and Defence. International Society for Optics and Photonics, 2011, pp. 818 605–818 605.

[b15] O. Gupta et al., "Reconstruction of hidden 3D shapes using diffuse reflections," Opt. Express, vol. 20, no. 17, pp. 19 096– 19 108, Aug 2012.

[b16] A. Velten et al., "Recovering three-dimensional shape around a corner using ultrafast time-of-flight imaging," Nat Commun, vol. 3, p. 745, 03 2012. [Online]. Available: http://dx.doi.org/10.1038/ncomms1747

[b17] F. Heide, M. B. Hullin, J. Gregson, and W. Heidrich, "Low-budget transient imaging using photonic mixer devices," ACM Trans. Graph., vol. 32, no. 4, pp. 45:1–45:10, Jul. 2013.

[b18] F. Heide, L. Xiao, W. Heidrich, and M. B. Hullin, "Diffuse mirrors: 3D reconstruction from diffuse indirect illumination using inexpensive time-of-flight sensors," in 2014 IEEE Conf. Comput. Vision Pattern Recognition, June 2014, pp. 3222–3229.

[b19] F. Heide, W. Heidrich, M. Hullin, and G. Wetzstein, "Doppler time-of-flight imaging," in ACM SIGGRAPH 2015 Emerging Tech., ser. SIGGRAPH '15. ACM, 2015, pp. 9:1–9:1. [Online]. Available: http://doi.acm.org/10.1145/2782782.2792497

[b20] A. Kadambi, H. Zhao, B. Shi, and R. Raskar, "Occluded imaging with time-of-flight sensors," ACM Trans. Graph., vol. 35, no. 2, pp. 15:1–15:12, Mar. 2016. [Online]. Available: http://doi.acm.org/10.1145/2836164

[b21] J. Klein, C. Peters, J. Martín, M. Laurenzis, and M. B. Hullin, "Tracking objects outside the line of sight using 2d intensity images," Scientific Reports, vol. 6, 2016.

[b22] A. K. Pediredla, M. Buttafava, A. Tosi, O. Cossairt, and A. Veeraraghavan, "Reconstructing rooms using photon echoes: A plane based model and reconstruction algorithm for looking around the corner," in 2017 IEEE International Conference on Computational Photography (ICCP), May 2017, pp. 1–12.

[b23] V. Arellano, D. Gutierrez, and A. Jarabo, "Fast back-projection for non-line of sight reconstruction," in ACM SIGGRAPH 2017 Posters, ser. SIGGRAPH '17. New York, NY, USA: ACM, 2017, pp. 79:1–79:2. [Online]. Available: http://doi.acm.org/10.1145/3102163.3102241

[b24] M. La Manna, F. Kine, E. Breitbach, J. Jackson, and A. Velten, "Error backprojection algorithhms for non-line-of-sight imaging," MINDS@UW, Aug 2017. [Online]. Available: https://minds.wisconsin.edu/handle/1793/77851

[b25] F. Zappa, S. Tisa, A. Tosi, and S. Cova, "Principles and features of single-photon avalanche diode arrays," Sensors and Actuators A: Physical, vol. 140, no. 1, pp. 103 – 112, 2007.

[b26] A. Tosi, A. D. Frera, A. B. Shehata, and C. Scarcella, "Fully programmable single-photon detection module for ingaas/inp single-photon avalanche diodes with clean and sub-nanosecond gating transitions," Review of Scientific Instruments, vol. 83, no. 1, p. 013104, 2012. [Online]. Available: https://doi.org/10.1063/1.3675579

[b27] M. Buttafava, G. Boso, A. Ruggeri, A. D. Mora, and A. Tosi, "Time-gated single-photon detection module with 110 ps transition time and up to 80 MHz repetition rate," Review of Scientific Instruments, vol. 85, no. 8, p. 083114, 2014. [Online]. Available: https://doi.org/10.1063/1.4893385

[b28] M. Buttafava et al., "Non-line-of-sight imaging using a time-gated single photon avalanche diode," Opt. Express, vol. 23, no. 16, pp. 20 997–21 011, Aug 2015. [Online]. Available: http://www.opticsexpress.org/abstract.cfm?URI=oe-23-16-20997

[b29] A. Kak and M. Slaney, Principles of Computerized Tomographic Imaging.

[b30] Sentech America Inc., "STC-MBA5MUSB3 Datasheet," IEEE Press, 1998.
http://www.sentechamerica.com/En/Cameras/USB/ STC-MBA5MUSB3, [Online].

[b31] Onefive GmbH, "Katana HP datasheet," http://www.onefive.com/img/products/katana-hp datasheet.pdf, [Online]. March 28, 2018 DRAFT25

[b32] Micro Photon Devices, "InGaAs SPAD - Gated Datasheet," http://www.micro-photon-devices.com/Docs/Datasheet/ InGaAs Datasheet.pdf, [Online].

[b33] M. Wahl, "Time-Correlated Single Photon Counting," PicoQuant GmbH, Tech. Rep., 2014. [Online]. Available: https://www.picoquant.com/images/uploads/page/files/7253/technote tcspc.pdf

[b34] Pico Quant, "PicoHarp300 - Datasheet," http://www.micro-photon-devices.com/Docs/Datasheet/InGaAs Datasheet.pdf, [Online].

[b35] Micro Photon Devices, "Picosecond Delayer," http://www.micro-photon-devices.com/Docs/Datasheet/PSD MOD.pdf, [Online].

[b36] National Instruments, "NI USB-6001 Low-Cost DAQ USB Device Manual," http://www.ni.com/pdf/manuals/374369a.pdf, [Online].

[b37] Z. Zhang, "A flexible new technique for camera calibration," IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 22, no. 11, pp. 1330–1334, Nov 2000.

[b38] J. Heikkila and O. Silven, "A four-step camera calibration procedure with implicit image correction," in Proceedings of the 1997 Conference on Computer Vision and Pattern Recognition (CVPR '97), ser. CVPR '97. Washington, DC, USA: IEEE Computer Society, 1997, pp. 1106–. [Online]. Available: http://dl.acm.org/citation.cfm?id=794189.794489

[b39] Arduino, "Arduino Uno Rev3," https://store.arduino.cc/usa/arduino-uno-rev3, [Online].

J. H. Nasmyth and J. Carpenter, "The Moon: Considered as a planet, a world and a satellite," John Murray, Albemarle Street, London, 1874.

D. L. Evans, T. G. Farr, and J. B. Adams, "Spectral reflectance of weathered terrestrial and Martian surfaces," Proc. Lunar and Planet. Sci. Conf. XII, Section 2, Houston, TX, March 16-20, 1981, pp. 1473-1479, 1982.

L. Kerber, I. Nesnas, L. Keszthelyi, J.W. Head, B. Denevi, P.O. Hayne, K. Mitchell, J.W. Ashley, J.L. Whitten, A.M. Stickle, M. Paton, K. Donaldson-Hanna, R.C. Anderson, D. Needham, P. Isaacson and L. Jozwiak, "Moon Diver: A Discovery mission concept for understanding the history of the Mare basalts through the exploration of a lunar Mare pit," in Lunar Planet. Sci. LIX, Abstract #1956, 2018.

A. Parness, N. Abcouwer, C. Fuller, N. Wiltsie, J. Nash and B. Kennedy, "LEMUR 3: A limbed climbing robot for extreme terrain mobility in space," in. Proc. of IEEE Int. Conf. Robotics and Automation (ICRA), Singapore, 2017.

I. A. D. Nesnas, P. Abad-Manterola, J. A. Edlund and J. W. Burdick, "Axel Mobility Platform for Steep Terrain Excursions on Planetary Surfaces," in Proceedings of the IEEE Aerospace Conference, Big Sky, Montana, 2008.

R. Seu, R. J. Phillips, D. Biccari, R. Orosei, A. Masdea, G. Picardi, A. Safaeinili, B. A. Campbell, J. J. Plaut, L. Marinangeli, S. E. Smrekar and D. C. Nunes, "SHARAD sounding radar on the Mars Reconnaissance Orbiter", J. Geophys. Res., 112, E05S05, 2007.

Cushing, G.E., Titus, T.N., Wynne, J.J., Christensen, P.R., 2007. THEMIS observes possible cave skylights on Mars. Geophysical Research Letters 34, L17201