Quantitative data analysis methods for 3D microstructure characterization of Solid Oxide Cells

Jørgensen, Peter Stanley; Larsen, Rasmus; Bowen, Jacob R.; Hansen, Karin Vels; Lassen, Niels Christian Krieger; Wallenberg, Reine

Publication date:
2010

Document Version
Publisher's PDF, also known as Version of record

Quantitative Data Analysis Methods for 3D Microstructure Characterization of Solid Oxide Cells

Peter Stanley Jørgensen

Kongens Lyngby 2010
IMM-PHD-2010-231

Technical University of Denmark
Informatics and Mathematical Modelling
Building 321, DK-2800 Kongens Lyngby, Denmark
Phone +45 45253351, Fax +45 45882673
reception@imm.dtu.dk
www.imm.dtu.dk

IMM-PHD: ISSN 0909-3192

Summary

The performance of electrochemical ceramic devices such as solid oxide fuel and electrolyser cells depends on the distribution of constituent phases on the micro or nano scale, also known as the microstructure. The microstructure governs key properties such as ion, electron and gas transport through percolating networks and reaction rates at the triple phase boundaries. Quantitative analysis of microstructure is thus important both in research and development of optimal microstructure design and fabrication. Three dimensional microstructure characterization in particular holds great promise for gaining further fundamental understanding of how microstructure affects performance.

In this work, methods for automatic 3D characterization of microstructure are studied: from the acquisition of 3D image data by focused ion beam tomography to the extraction of quantitative measures that characterize the microstructure. The methods are exemplified by the analysis of Ni-YSZ and LSC-CGO electrode samples.

Automatic methods for preprocessing the raw 3D image data are developed. The preprocessing steps correct for errors introduced by the image acquisition by the focused ion beam serial sectioning. Alignment of the individual image slices is performed by automatic detection of fiducial marks. Uneven illumination is corrected by fitting hypersurfaces to the spatial intensity variation in the 3D image data.

Routine use of quantitative three dimensional analysis of microstructure is generally restricted by the time consuming task of manually delineating structures within each image slice or the quality of manual and automatic segmentation

schemes. To solve this, a framework for the automatic segmentation of 3D image data is developed. The technique is based on a level set method and uses numerical approximations to partial differential equations to evolve a 3D surface to capture the phase boundaries. Vector fields derived from the experimentally acquired data are used as the driving forces. The framework performs the segmentation in 3D rather than on a slice by slice basis. It naturally supplies sub-voxel accuracy of segmented surfaces and allows constraints on the surface curvature to enforce a smooth surface in the segmentation.

A high accuracy method is developed for calculating two phase boundary surface areas and triple phase boundary length of triple phase systems. The calculations are based on sub-voxel accuracy segmentations of the constituent phases. The method performs a three phase polygonization of the interface boundaries which results in a non-manifold mesh of connected faces. The triple phase boundaries can be extracted from the mesh as connected curve loops without branches. The accuracy of the method is analyzed by calculations on geometrical primitives.

A suite of methods is developed for characterizing the shape and connectivity of phase networks. The methods utilize the fast marching method to compute distance maps and optimal paths in the microstructure network. The extracted measurements are suited for the quantitative comparison and evaluation of microstructures. The quantitative measures characterize properties of network path tortuosity, network thickness, transport path width and dead ends.

Resumé

Brændselsceller og elektrolyseceller kan produceres af keramiske materialer. Deres elektrokemiske ydeevne er afhængig af hvorledes de enkelte bestanddele af materialet er fordelt på mikro eller nano niveau, også kendt som mikrostrukturen. Mikrostrukturen er bestemmende for vigtige egenskaber så som ion, elektron og gas transport gennem sammenhængende netværk, samt reaktionshastigheden ved trefasegrænserne. Kvantitativ analyse af mikrostrukturen er derfor vigtig inden for forskning og til udvikling af optimale metoder til design og fremstilling af optimal mikrostruktur. Tredimensionel karakterisering i særdeleshed er lovende for en dybere forståelse af hvordan mikrostrukturen påvirker ydeevnen.

I denne afhandling studeres metoder til automatisk karakterisering af mikrostruktur, fra dannelse af 3D billede data ved fokuseret ion stråle tomografi til måling af kvantitative værdier der karakteriserer mikrostrukturen. Metoderne er eksemplificeret ved analyse af Ni-YSZ og LSC-CGO elektrode prøver.

Der er blevet udviklet automatiske metoder til at pre-processerer de rå 3D billede data. Pre-processeringen korrigerer for fejl der opstår ved dannelsen af billeddata ved fokuseret ion stråle tomografi. Korrektion af billedforskydninger sker ved automatisk detektering af referencemærker i billederne. Varierende belysning af billeddata bliver korrigeret ved at fitte en hyperflade til den rumlige intensitets variation i 3D billeddata.

Rutinemæssig brug af kvantitativ tredimensionel analyse af mikrostruktur er generelt begrænset af det store tidsforbrug som er forbundet med manuel optegning af strukturer i billederne eller kvaliteten af manuelle og automatiske segmenterings metoder. En overordnet metode til automatisk at foretage segmenteringen af 3D billede data er blevet udviklet for at løse dette problem.

Metoden er baseret på en *level set* metode og anvender numeriske approksimationer til partielle differentialligninger til at forme en 3D overflade så den passer overens med fase grænserne i mikrostrukturen. Vektorfelter udledt fra de eksperimentelt optagede billeddata bruges som den drivende kraft for at udvikle overfladen. Metoden udfører segmenteringen i 3D i stedet for på de enkelte billeder hver for sig. Metoden repræsenterer grænsefladerne med sub-voxel præcision og gør det muligt at indsætte begrænsninger for krumningen af overfladen for at opretholde en glat overflade i segmenteringen.

En metode med høj nøjagtighed er blevet udviklet til at udregne tofasegrænse overflade arealer og trefasegrænse længde for trefase systemer. Udregningerne er baseret på en segmentering af de faser mikrostrukturen består af med sub-voxel nøjagtighed. Metoden foretager en tre fase polygonisering af fase grænsefladerne hvilket resulterer i et net af sammenhængende polygoner. Tre fase grænserne kan udtrækkes fra polygon nettet som sammenhængende kurveringe uden forgreninger. Nøjagtigheden af metoden analyseres ved udregninger på kugleflader.

En samling af metoder er blevet udviklet til at karakterisere form og struktur for netværk af sammenhængende faser. Metoderne anvender *the fast marching method* til at udregne afstandskort og optimale veje i fase netværkene i mikrostrukturen. De udregnede mål er egnet til at foretage en kvantitativ sammenligning og bedømmelse af mikrostrukturer. De kvantitative mål karakteriserer egenskaber så som hvor snoede netværksstrukturerne er, deres tykkelse og mængden af blinde ender.

Preface

This thesis was prepared at the Technical University of Denmark in partial fulfillment of the requirements for acquiring the Ph.D. degree. The thesis work was carried out partially at Risø - DTU and partially at DTU - Informatics during the period 2007–2010.

The thesis deals with methods for extracting quantitative measurements from 3D tomography image data of microstructure. The main focus is on developing the computational methods that allow accurate and relevant information to be automatically extracted.

Lyngby, March 2010

Peter Stanley jørgensen

Papers included in this thesis

Chapter 5 P. S. Jørgensen, K. V. Hansen, R. Larsen, J. R. Bowen. A framework for automatic segmentation in three dimensions of microstructural tomography data. Ultramicroscopy, 2010, Volume 110, p. 216-228.

Chapter 6 P. S. Jørgensen, K. V. Hansen, R. Larsen, J. R. Bowen. High accuracy interface characterization of three phase material systems in three dimensions. (Submitted).

Chapter 7 P. S. Jørgensen, K. V. Hansen, R. Larsen, J. R. Bowen. Geometrical characterization of interconnected phase networks in three dimensions. (Submitted).

Other related papers not included in the thesis

- P. R. Shearing, L. E. Howard, P. S. Jørgensen, N. P. Brandon, S. J. Harris. Characterization of the 3-dimensional microstructure of a graphite negative electrode from a Li-ion battery. Electrochemistry Communications, 2010, Volume 12, p. 374-377.

Acknowledgments

This work was supported financially by The Programme Commission on Sustainable Energy and Environment, The Danish Council for Strategic Research, via the Strategic Electrochemistry Research Center (SERC) (www.serc.dk), contract no. 2104-06-0011.

I would like to thank a number of people who have contributed to making this thesis better and much more enjoyable to write:

All the great people at my two workplaces: Risø DTU and DTU Informatics. You are too many to name. Thanks for helping me out, for fruitful and silly discussions and for generally making me smile. This really would have been much less enjoyable without you.

The people at Imperial College in London for taking good care of me during my stay in the summer of 2009. In particular *Paul Shearing* for being a great host.

My five supervisors (in alphabetical order): *Jacob Ross Bowen*, *Karin Vels Hansen*, *Rasmus Larsen*, *Niels Christian Krieger Lassen* and *Reine Wallenberg*. I appreciate your guidance and help through the exciting and, for me at the time, new fields of electron microscopy and electrochemistry. Thank you for good advice and for always being positive and just generally nice people.

My fantastic family for supporting me through good and bad times (Mmmmmmmmmmmmmmmmmmmmmmmmm). Special thanks to my very young nephew *Malte* who has been a great source of motivation and happy thoughts during those long days of writing.

Contents

Chapter 1

Introduction

This work deals with computational methods for quantitative characterization of microstructure. The main motivation for developing these methods has been to quantify important microstructure characteristics that affect the performance of solid oxide cells (SOC). However, the majority of the work presented here is not tied to the analysis of SOC and the developed methods can be applied to the analysis of other structures. The focus is on developing the experimental and computational methodologies that allow the extraction of quantitative parameters. The combined results of this work should be seen as steps towards the creation of an advanced automated tool for quantitative 3D microstructural analysis.

1.1 Thesis overview

The following chapters will describe each of the steps required to obtain quantitative 3D measurements from a sample. The central chapters covering segmentation, interface characterization and network characterization consist of one published paper and two submitted manuscripts.

- Chapter 1 gives a motivational introduction to SOCs. The concept of

microstructure is introduced and its influence on SOC performance is described.

- Chapter 2 gives an introduction to focused ion beam (FIB) tomography. FIB tomography has been used to obtain 3D image data throughout this work. The chapter describes instrumentation, image formation and experimental techniques.
- Chapter 3 deals with methods for automatic correction of drift artifacts induced by the image acquisition process.
- Chapter 4 describes a basic method for automatic correction of non-uniform illumination.
- Chapter 5 covers 3D image segmentation and consists of a paper [1] published in Ultramicroscopy. The chapter describes how a level set method can be used to segment the 3D image data with a number of desired properties.
- Chapter 6 covers methods for quantitative characterization of phase interfaces and consists of a paper manuscript submitted for publication. The chapter describes how an interface polygonization method can be used to extract measurements of surface area and triple phase boundaries.
- Chapter 7 covers methods for quantitative characterization of phase networks and consists of a paper manuscript submitted for publication. The chapter describes the calculation of parameters that characterize the shape and connectivity of network structures.
- Chapter 8 provides perspectives and outlook on the combined results of the previous chapters. Some additional SOFC modeling results are given to provide perspective on the use of the calculated parameters for modeling.
- Chapter 9 presents the conclusions drawn in this thesis.

1.2 Fuel cells

A fuel cell is an electrochemical device that converts chemical energy into electricity and heat. The electricity and heat is generated by reactions between a fuel and an oxidant at an electrolyte. The reactants flow into the cell and the reaction products flow out of the cell. A fuel cell is similar to a battery in how it works, but where the battery contains a finite amount of reactants the fuel cell can in theory be operated indefinitely with a constant supply of fuel. In practice, degradation of the component with operating time limits the lifespan.

A combustion engine likewise utilizes the chemical energy stored in a fuel. However, in the combustion engine fuel is burned to release heat while in the fuel cell, electricity is produced directly along with heat. Consider the simplest combustion example: the conversion of hydrogen and oxygen into water (1.1).

$$H_2 + \frac{1}{2}O_2 \rightleftharpoons H_2O \tag{1.1}$$

When hydrogen molecules and oxygen molecules collide, a reaction occurs. The hydrogen-hydrogen and oxygen-oxygen bonds are broken and new hydrogen-oxygen bonds are formed by transfer of electrons. The energy of the produced water molecules is lower than that of the hydrogen and oxygen molecules. The energy difference is released as heat. To produce electricity the heat must first be converted into mechanical energy and then from mechanical energy into electricity, a process that is potentially inefficient.

A fuel cell works by spatially separating the hydrogen and oxygen reactants. The combustion reaction is thus split into two electrochemical half reactions. For a solid oxide fuel cell (SOFC) the reactions are:

$$H_2\ (g) + O^{2-} \rightarrow H_2O\ (g) + 2e^- \tag{1.2}$$

$$\frac{1}{2}O_2\ (g) + 2e^- \rightarrow O^{2-} \tag{1.3}$$

Spatial separation is achieved through the use of an electrolyte that allows ions to pass through but not electrons. The electrical current created by the electrons moving from the fuel species to the oxidant species can then be used to power electrical devices. It is the ability of fuel cells to continuously convert a fuel directly into electricity, with high efficiency, that makes fuel cells attractive [2].

A number of different fuel cell types currently exist. They are named after the material that is used as electrolyte. These include: alkaline, polymer exchange membrane, phosphoric acid, molten carbonate and solid oxide fuel cells. Each type differ in their characteristics such as operating temperature, electric efficiency and fuel composition demands. Of these, only the SOFC will be described here.

1.3 Solid oxide fuel cells

An illustration of an SOFC supplied with hydrogen and air can be seen in figure 1.1. An SOFC consists of three main components.

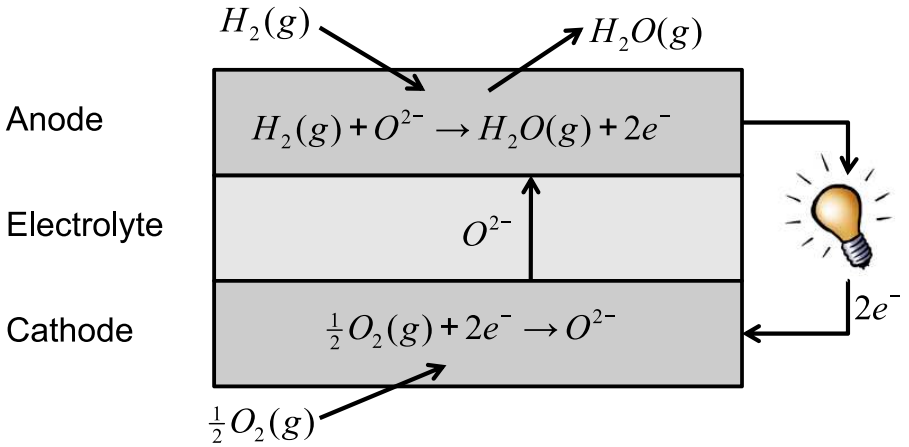

Figure 1.1: Illustration of the basic operation of an SOFC with hydrogen and air.

- **The anode** is the electrode where the oxidation reaction 1.2 takes place. It is a porous structure that consists of an electron conducting phase and an ion conducting phase.
- **The electrolyte** transports the dissolved ions between the cathode and the anode electrodes. The electrolyte in an SOFC is made from a thin solid ceramic membrane that allows ions to pass through it and acts as a physical barrier for electrons and gases.
- **The cathode** is the electrode where the reduction reaction 1.3 takes place. Like the anode it consists of an electron conducting phase and an ion conducting phase.

The driving force in an SOFC is the difference in oxygen activity between the two electrodes. Oxygen is delivered to the cathode by an external feed. The oxygen is then dissociated at the cathode by accepting two electrons (reaction 1.3). The oxide ions migrate through the electrolyte to the anode where they react with the hydrogen by the release of electrons. The product of this reaction (1.2) is water and surplus electrons. The electrons are then conducted back to the cathode through an external circuit where a load can be introduced to draw power.

1.3.1 Advantages and problems

An SOFC has a number of advantages [3, 4, 5] which makes them attractive compared to combustion engines and other fuel cell types.

- SOFCs have high efficiency and can run on many different fuels, such as natural gas or methanol.
- All the components of the cell are solid state, including the electrolyte. An SOFC thus contains no moving parts, which means less maintenance and less noise.
- The high operation temperature means that the high quality excess heat can be used for combined heat and power systems or for additional electricity generation in a hybrid system.
- Fuel cells are highly modular. SOFC efficiency is relatively independent of size and large systems can be constructed by combining smaller modules. Each module can be replaced independently of other modules for easy maintenance or replacement.

SOFC systems have however not yet had their commercial breakthrough. The main obstacle for the widespread use of SOFCs on the energy market is overall production cost in relation to performance and durability. However, the past 20 years have seen a dramatic increase in SOFC research worldwide and all aspects of the problems are being actively pursued: cheaper and/or higher performing materials, less material use (thinner cells), improved durability and optimized production methods.

1.3.2 Materials

The choice of materials in the electrolyte and the electrodes must be made with careful considerations. The materials must perform well in their primary role of electron conductor, ion conductor and/or membrane. However, they must also meet other requirements such as: high catalytic activity, compatible thermal expansion coefficients, low cost and acceptable behavior in reducing and/or oxidizing environments.

An often used material for the electrolyte is yttria stabilized zirconia (YSZ). YSZ is often also used as the ion conducting phase in both the anode and the cathode. Ni is used as the electron conductor in the anode due to its excellent

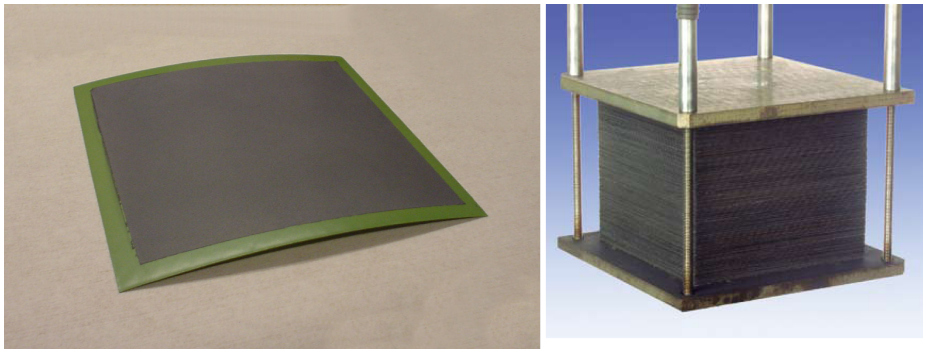

Figure 1.2: (Left) A planar fuel cell. The dark active area is 12 x 12 cm. (Right) A stack of 75 cells in series.

ability to catalytically dissociate hydrogen and withstand the high operating temperature in reducing atmosphere. Strontium-doped lanthanum manganite (LSM) has high electronic conductivity in oxidizing atmospheres and is often used as the electron conducting phase in the cathode. Other cathode materials are being investigated for better performance and reduced operating temperature. One such material combination is strontium-substituted lanthanum cobaltite (LSC) and gadolinia-substituted ceria (CGO).

An SOFC operates at high temperatures compared to other types of fuel cells. High temperature improves the kinetic performance of the cell but also results in many practical challenges relating to durability. Current operating temperatures are between 500 and 1000 °C with development moving towards lower temperatures.

The thickness of an SOFC is less than a millimeter. To increase the voltage, planar SOFCs are typically used in series in a stack. The number of cells in the stack can be scaled to meet specific voltage requirements or module sizes. Figure 1.2 shows a planar SOFC next to a stack of a 75 cells.

1.4 Energy perspectives of solid oxide cells

In recent years there has been a growing acceptance of the need to move from fossil fuel dependent energy technologies to renewable energy technologies. As a standalone technology SOFCs could offer improved fuel efficiency and less CO_2 and NO_x emission over traditional electricity production methods. However, a fuel is still required for SOFCs to produce electricity and heat. For an SOFC

to be a part of a renewable energy system the fuel must be produced in a sustainable way.

One method of acquiring renewable fuels is from biomass. Fuels such as hydrogen may soon be produced from organic waste materials [6]. Other sources of renewable energy, such as wind turbines or solar cells, produce electricity directly. However, since they are dependent on wind and clear daytime skies, their power generation is unpredictable and excess electricity generation can currently not be stored. One solution to the problem of excess energy storage could be to convert the electricity into hydrogen by means of electrolysis.

Electrolysis cells utilize the reverse process of fuel cells. Water is split electrochemically into hydrogen and oxygen by applying electrical energy. Electrolysis can be performed by a solid oxide electrolyser cell [7] (SOEC). An SOEC is essentially just an SOFC run in reverse and as such the materials used and production methods are similar. The term solid oxide cell (SOC) is used as a common reference to SOFC and SOEC. Thus, the term SOC is used henceforth for both.

A future vision for how the world's energy demands could be met by renewable energy is the hydrogen society. This vision is build around hydrogen as an energy carrier [8]. Hydrogen would be produced through biomass or electrolysis, transported to where the energy was needed and then be converted to electricity and heat by a fuel cell. Small combined heat and power units [9] could be used in single households to produce both electricity and heat and fuel cell powered electric cars would have water as its only emission.

1.5 SOFC performance

All reactions in an SOFC occur near the phase interfaces. At these interfaces the reactions take place in both directions. This means that water can be created from hydrogen and oxygen ions, but also that water can dissociate into hydrogen and oxygen ions. A certain activation energy is required for each of these two reactions to occur. This activation barrier is, however, typically not equal for the two reaction directions. The probability that each of the reactions occurs is a function of the required activation energy. The reaction rates of the forward and backward reactions are thus dependent on the local number of reactant species. The net reaction rate is the difference between the forward and backward reaction rates.

The current generated by an SOFC is directly proportional to the speed of

the electrochemical reactions and thus the amount of fuel consumed. The area specific current, j, is most often used since it provides a size invariant measure of current. There are several ways of increasing the area specific current [10]:

- Decrease the activation barrier by choosing an electrode material with high catalytic activity.
- Increase reactant concentration by increasing pressure.
- Increase the temperature. High temperature increases the probability that the reactions occur.
- Increase reaction site density. The reactions can only occur where the electron conducting, ion conducting and gas phases meet, also known as the triple phase boundaries (TPBs). For an electrode with a mixed ionic electronic conductor (MIEC) the reactions occur at the MIEC/gas two phase boundaries. The reactions at the reaction sites can only occur at a finite rate. Increasing the overall TPB and/or two phase boundary density thus improves reaction rates.

Increasing temperature and pressure results in many challenges concerning leakage and material choices. New materials are being investigated but as previously stated they must be compatible with other components and operating conditions. The increase in reaction site density relates mainly to how the cells are produced. The TPB density is steadily being optimized through better recipes and production methods.

1.5.1 Efficiency

The area specific current is only one aspect of SOFC performance, the voltage of the cell is equally important. The voltage will generally drop as current is produced. This means that the voltage of an SOFC can be seen as a measure of its efficiency (power generated per unit of fuel).

The voltage without any current load is generally referred to as the open circuit voltage (OCV). The difference between the OCV and the voltage at a certain current is called the overpotential or overvoltage, η. The overpotential can be written as:

$$\eta = R_i j \tag{1.4}$$

where j is the area specific current density and R_i is the area specific internal resistance. The internal resistance can be stated as the sum of the resistances of many smaller terms relating to specific losses from different processes [10]. Here, the internal resistance is divided into two main contributions:

- Ohmic resistance is the resistance that originates from the transport of charge through the cell. The ohmic resistance is a result of both electron and ion conduction through electrodes and electrolyte. The resistance increases with the length traveled by electrons and ions. SOFC electrodes in general and the electrolyte in particularly are thus designed to be thin to decrease ohmic resistance.

- Polarization resistance is the resistance that originates from reaction rate limitations. The reaction rates increase with increasing reactant concentration. To keep the concentration of reactants high the reaction product (water in the anode) must be transported away and new reactants must be supplied to the active sites (H_2 and O^{2-} in the anode). This mass transport occurs through bulk motion by convection and concentration gradient driven diffusion. Convection transport occurs mainly in the external flow channels at the outer edge of the electrodes. Mass transport near the electrolyte occurs almost exclusively by diffusion.

These resistances can be minimized by making the pathways that molecules, ions and electrons need to travel shorter and wider. The shape of the pathways is controlled by the electrode microstructure.

1.6 Microstructure

SOFCs are produced from mixtures of very fine powders of the constituent materials. The powders are mixed with a binder into a slurry or paste and used to form the shape of a specific component. Different methods of processing the slurry are used, such as tape casting [11] or screen printing [12, 13]. After the forming process the component is sintered under a controlled atmosphere which allows the powder particles to partially fuse and form connected structures. The production of an SOFC involves a number of steps where the individual layers; anode, electrolyte and cathode are added one after another. Many different methods and steps are used but the processes of making a slurry, forming it and sintering are similar. Even though, the steps are similar they are not necessarily simple and they have a large impact on the finished cell.

The final distribution of the individual phases (electron conductor, ion conductor and gas phase) in the electrodes is called the microstructure of the electrode. The very fine powders used creates phase structures in the microstructure that are in the order of 100s of nanometers. In comparison to this small scale the production methods are bulky. This means that the microstructure can not be controlled directly but is a function of powder particle size [14, 15], slurry composition [16, 17], the forming method and the sintering process [18]. The actual microstructure produced thus has a stochastic nature.

Important microstructure parameters can be divided into two major groups: reaction site quantity/quality parameters and transport efficiency parameters. Figure 1.3 shows a 2D illustration of how these parameters affect SOFC performance, exemplified by a small section of the electrolyte and the anode.

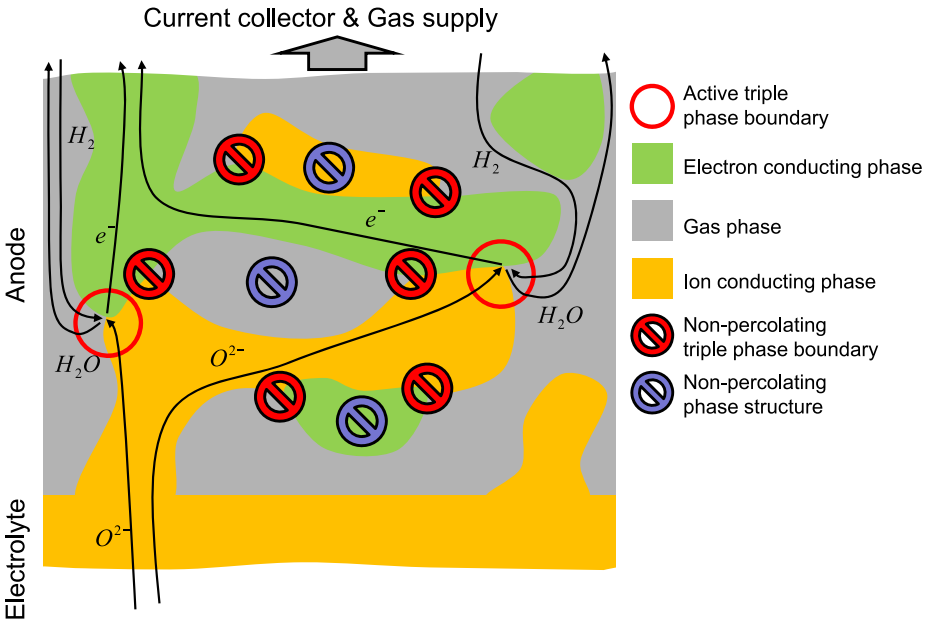

Figure 1.3: Illustration of the reactions at the active TPBs close to the electrolyte/anode interface.

- **Reaction sites**. In figure 1.3 there are only two sites where the reactions can occur, marked as the active triple phase boundaries. These sites can be reached by oxide ions from the electrolyte through the oxide ion conducting material and by hydrogen molecules from the gas supply through the pores. At the same time electrons can be conducted through the electron conducting phase to the current collector. All the reactants can thus

reach the TPB site and the products can escape the cell.

There are six additional TPBs in the structure in figure 1.3. However, these TPB sites lack a connection to either the gas supply, the electrolyte, or the current collector. A phase structure that does not have a connection to its respective source or destination is denoted a non-percolating structure. Non-percolating structures are generally undesirable since they do not contribute to reactions or charge/mass transport. The material could thus have been utilized better elsewhere in the structure.

- **Transport efficiency**. To minimize ohmic resistance and polarization resistance it is desirable to have transport networks (pathway structures) in the microstructure that facilitate rapid transport of electrons, ions and gas. For mass transport and diffusion in particular the effectiveness of the transport network is dependent on many aspects. For instance, twisting or tortuous pathways increase pathway length and narrow pathways can act as bottlenecks.

A good microstructure is characterized by having high TPB density (TPB length per unit volume), short pathways with few bottlenecks and good percolation of phase structures. These characteristics are challenging to achieve simultaneously.

An optimal microstructure utilizes constituent phase material to maximum effect. Not only should all network structures be percolating but each part of the structure should have optimal dimensions. For instance, pore space occupied by large cavities might be more efficiently used to form several narrower pathways in another location. With current industrially relevant production methods this level of optimization is not possible and microstructure optimization must be performed by altering recipes and bulk production methods. However, even if an optimal microstructure can not be achieved, it is still important to be able to determine how good a given microstructure is in respect to the desired characteristics. By systematically comparing new recipes and production methods with the current state of the art microstructure with improved characteristics can be steadily achieved.

It is possible to evaluate a given cell by directly measuring its performance. However, because of the complex performance dependencies it is challenging to isolate the microstructure's influence and determine which microstructure characteristics the performance is attributed to.

1.7 Motivation for the thesis

The analysis of microstructure has traditionally mainly been performed by inspection of scanning electron microscopy images of polished surfaces [19]. This technique provides 2D information about the microstructure.

Methods of stereology can be used to obtain unbiased and quantitative measurements of 3D properties from measurements on 2D planar sections. Important quantities such as surface areas and phase fractions can be obtained by these methods. However, to obtain more advanced quantities such as particle size distributions assumptions on the shape of the phase structures must be made. Phase structures of arbitrary shape and networks of complex connectivity limit the type and accuracy of the measurements that can be extracted by stereological methods.

The connectivity of the phases in SOC microstructure is of great importance as illustrated in figure 1.3. By employing 3D image information in the analysis the interconnectivity of phases can be analyzed as well. However, to be able to make systematic comparisons between different microstructures it is not sufficient to visually inspect the structures. There are two reasons for this. 1) The visual inspection of 3D information is challenging since humans can only perceive surfaces and not dense volumes. 2) The comparisons should be based on quantitative and repeatable parameters rather than qualitative anecdotal descriptions. The calculation of microstructure parameters is important in understanding the influence of microstructure on performance, but also for systematic optimization of recipes and production methods.

3D image data contains a large amount of information and is not well suited to be analyzed manually. The motivation for this study has been to apply computational algorithms to the analysis of 3D data. The focus has been on two main aspects: decreasing the human labor requirements by automation of the analysis and calculating accurate parameters that describe important microstructure characteristics.

Chapter 2

Focused ion beam tomography

Abstract

Focused ion beam (FIB) tomography is an established experimental method that allows 3D microstructure information to be collected from a sample. The method uses a scanning electron microscope (SEM) to acquire high resolution surface images of the sample surface and focused ion beam milling to remove thin sections of the sample. An overview of FIB literature is presented along with the theory and practice of SEM and FIB systems. Lastly, the experimental methodology that allows acquisition of 3D data by a FIB/SEM system is described.

2.1 Introduction

This chapter will describe the instrumentation, theory and practical techniques used in acquiring 3D image data by focused ion beam (FIB) tomography. Electron microscopy and FIB are large and important research fields on their own with a wide range of applications. This chapter will thus focus on describing the

aspects of FIB tomography that have implications for the methods described in the following chapters.

FIB tomography generally refers to the combination of a FIB system used as a micro machining tool together with an imaging technology such as scanning electron microscopy (SEM), FIB imaging or microanalysis. SEM imaging has been used for all FIB tomography applications in this text. At the basic level 3D reconstructions are obtained by sequentially acquiring a surface image of the sample followed by the removal of a thin section of the sample by the FIB.

2.1.1 Other tomography techniques

A number of alternative tomography techniques exist for acquiring 3D volumes of different materials. They differ in their resolution, image quality, size of reconstructed volume and speed.

- Serial sectioning techniques [20, 21] based on automatic polishing of the sample can be used to acquire rapid 3D reconstructions by light microscope imaging with a layer separation down to 0.2 μm [22]. Similar serial sectioning techniques can be used with SEM imaging to obtain high 2D image resolution, although the need to remove the sample from the SEM chamber after each section makes the process laborious.
- X-ray tomography [23] can be performed by many different types of x-ray sources. However, to obtain high spatial resolution a synchrotron source is generally used for its high brightness, intensity and beam coherency. Sub-micrometer resolution has been shown on biological samples [24]. Smaller x-ray systems have been applied to the analysis of the pore phase of SOFC anodes [25] with sub-50 nm resolution. The main advantage of x-ray tomography is that it is non-destructive, allowing further subsequent analysis of samples to be performed.
- Transmission Electron Microscopy (TEM) nanotomography [26] or 3D-TEM can be used to obtain nanometer scale spatial resolution. The technique is however limited to thin samples (typically significantly less than 1 μm).

Compared to these tomography techniques FIB tomography operates in the important 10 nm to 10 μm resolution range [27]. FIB tomography with SEM imaging generally provides good image quality and sample sizes can be scaled with resolution by increasing the FIB milling current.

2.1.2 FIB state of the art

The earliest use of FIBs can be traced back to 1974 where Seliger & Fleming [28] first demonstrated their potential. Krohn & Ringo [29] tested the feasibility of cesium, gallium and mercury liquid sources and found the most suited to be gallium, which is still the most common liquid metal ion source in modern FIB systems. Early FIB systems in the 1980s found use in the semiconductor industry for device fabrication and repair techniques [30]. Kirk et al. [31] reported the use of FIB for preparation of cross-sectional specimens and Overwijk et al. [32] introduced the popular "lift out" technique for TEM lamella preparation [33]. TEM lamella preparation is still one of the main applications of modern FIB systems [34] along with nano-tool fabrication [35, 36] and preparation [37] for chemical (energy dispersive spectroscopy) or crystallographic (electron backscatter diffraction) analysis.

FIB tomography in its modern form was introduced by Inkson et al. [38, 39] in 2001. Holzer et al. [40] did an early three-dimensional analysis of porous $BaTiO_3$ ceramics and reported a voxel dimension down to 5.9 x 6.8 x 16.6 nm.

FIB tomography has during the last 10 years found widespread use in a wide range of applications in materials science. One application is the study of cement structures where Holzer, Munch and others have performed several thorough analyses [41, 42, 43, 44]. Another example is the study of particle and crystal morphologies in metals where several people have contributed, for instance in [27, 45, 39].

The analysis of the microstructure of SOFC electrodes using FIB tomography was first introduced by Wilson et al. [46]. The analysis of SOFC electrode microstructure has since then seen significant research interest [47, 48, 49, 50, 51, 52, 53]. The SOFC literature and other related contributions will be discussed further in the following chapters as they relate to the relevant topic.

2.2 Scanning electron microscopy

SEM is a well established and widely used technology for acquiring high magnification surface images of a specimen. Modern SEM systems can achieve magnifications in a range of many orders of magnitude with resolution down to the sub-nanometer level.

SEM images can have many similarities to light microscopy images but the

image formation is fundamentally different. These differences can lead to misinterpretation of images if the electron-specimen interactions are not carefully considered. This section will give a brief overview of instrumentation and image formation to form a basis for interpreting the intensity structure in SEM images. For a more comprehensive coverage the reader is referred to existing literature [54, 55].

2.2.1 SEM instrumentation

An SEM consists of two primary components (illustrated in figure 2.1):

- **An electron column**. The column consists of an electron gun that acts as a source of electrons and accelerates these to energies typically in the 0.5-30 keV range. A V-shaped tungsten filament has historically been used as the electron source in the past. However, the more recent field emission gun (FEG) electron source offers a more coherent beam with orders of magnitude higher brightness especially at low accelerating voltage. The FEG works by field electron emission over a sharply pointed tip by applying a strong potential gradient. The electrons emitted from the source are focused through a series of lenses that condense and focus the electrons to form a thin electron beam at the specimen surface. In order to form an image a deflector system is used to direct the electron beam in a scanning rectangular raster pattern across the specimen surface. It is the size of the raster pattern, and not the electron lenses that determines the magnification in the acquired images. However, resolution is controlled by the focused spot size of the electron beam and the accelerating voltage.
- **Electron detectors**. The interaction of the electron beam with the specimen surface results in a number of electromagnetic radiation phenomena. Different detectors are used to detect the radiation signal synchronously with the raster pattern, thus creating a signal intensity map/image.

2.2.2 Electron-specimen interactions

When the electron beam impinges on the specimen a multitude of interactions occur. It is these interactions that give SEM its great versatility. The electrons in the incident beam are referred to as primary electrons. The interactions between the primary electrons and the electron shells of the atoms of the specimen can be divided into two classes:

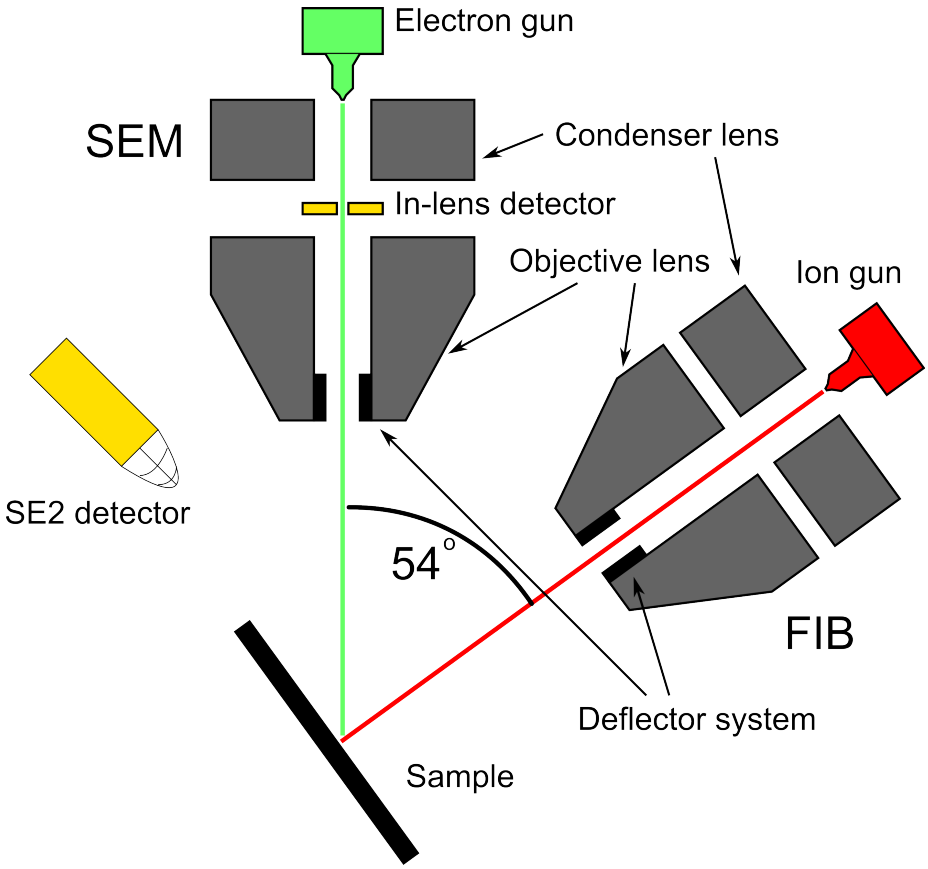

Figure 2.1: Illustration of the main components of a dual beam FIB/SEM system.

- Elastic scattering events change the trajectory of incident electrons without changing the kinetic energy of the electrons. These events result in the phenomenon of backscattered electrons (BSE) which are primary electrons that through one or more elastic scattering events escape the specimen surface with a high kinetic energy.
- Inelastic scattering events result in a transfer of energy from the incident electrons to the atoms of the specimen. Inelastic scattering leads to the generation of many forms of electromagnetic radiation of which only secondary electrons will be described here. Secondary electrons are created by ejection of an electron from an atoms valence or conduction band.

Each primary electron that enters the surface of the specimen will be subjected to a series of elastic and inelastic scattering events. Each inelastic scattering event reduces the energy of the electron which means that the primary electron will either be captured by the specimen or escape the specimen as a BSE.

The scattering events limit the penetration of the electron into the sample. The scattering of the primary electrons is a stochastic process where each electron will interact with the specimen through different trajectories. The region of the specimen over which the primary electrons interact with the specimen is known as the interaction volume. The size and shape of the interaction volume depends on the energy of the incident electron beam and the atomic number of the specimen.

The secondary electrons that are created from the inelastic scattering events typically have energies lower than 10 eV which mean that they travel only a short distance in the specimen (5-50 nm) dependent on specimen composition. This means that, for secondary electrons to escape the specimen and be detectable, they must be created by scattering events near the surface of the sample. The secondary electrons that escape the sample can be divided into three types illustrated in figure 2.2:

- SE_1 are produced as the primary electrons enter the specimen. They contain high resolution information because they are produced where the primary electrons enter the sample.
- SE_2 are produced as backscattered electrons leave the sample. Because of the large interaction volume of backscattered electrons (accelerating voltage dependent) the SE_2 carries less location specific information.
- SE_3 are produced as backscattered electrons that have escaped the specimen hit the internal surfaces of the microscope chamber which induce secondary electrons proportional to the BSE yield.

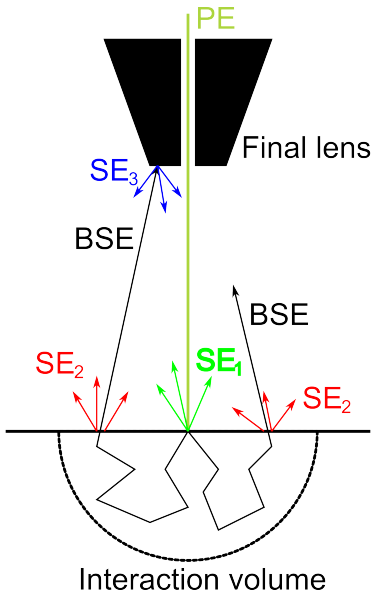

Figure 2.2: Electron signals.

Although their origins are different the SE_1, SE_2 and SE_3 have identical properties and the three signals are not easily separated.

The electrons (both BSE and SE) emitted from the specimen have a wide distribution of energies. A somewhat arbitrary upper threshold of 50 eV is typically used to distinguish between SEs and BSEs. The yield of electrons emitted from the surface is given by the backscatter coefficient η and the secondary electron coefficient δ as

$$\eta = \frac{n_{BSE}}{n_B} \quad , \quad \delta = \frac{n_{SE}}{n_B} \tag{2.1}$$

where n_B is the number of beam electrons incident on the specimen, n_{BSE} is the number of BSE emitted from the sample and n_{SE} is the number of SE emitted from the sample.

The yields of both BSE and SE have complex dependencies on microscope settings, specimen composition, surface properties and surface charge. However, some typical dependencies can be described. The backscatter coefficient is primarily dependent on the atomic number of the specimen in the interaction volume and the tilt angle of the beam on the specimen surface (deviation from

perpendicular). The backscatter coefficient ranges from around 0.1 for small atomic number and small tilt angle to around 0.8 for large atomic number and higher tilt angles.

The secondary electron coefficient is dependent on the beam energy, the tilt angle and to a lesser degree the atomic number of the specimen. The yield of SEs is typically several times smaller than the yield of BSEs. However, the total yield of electrons $\eta + \delta$ can exceed 1.

2.2.3 Image contrast

The dependency, of η and δ on both the atomic number of the specimen and the tilt angle, gives rise to two important forms of contrast: compositional contrast (atomic number) and topographic contrast. The contrast obtained is dependent on how the yield of electrons at a scan spot is detected and converted into an image intensity. The detection of the electron signal can be done with a range of detectors each with different strengths. Detectors can vary between microscopes, the description below is thus based on the Zeiss XB 1540 microscope that has been used for all experimental work.

- The Everhart Thornley detector also referred to as the SE2 detector is mounted on the wall of the specimen chamber and detects a combination of BSE and SE (SE_1, SE_2 and SE_3). The detector works by allowing electrons to strike a scintillator material after being accelerated through a large positive potential (10 to 12 kV). This positive potential allows low energy SE as well as BSE that reach the detector to generate light in the scintillator. The produced photons are then conducted outside the microscope chamber, through a vacuum seal, where it again is transformed into an electrical signal that can be recorded. To prevent the large potential in the detector from affecting the electron beam a Faraday cage surrounds the detector. A separate bias potential can be applied to the Faraday cage to increase the SE detection efficiency. Dependent on the bias being negative or positive the bias either results in a rejection of SE or an increase in detection efficiency. For large positive bias a large fraction of emitted SE from the specimen are detected including those being emitted in directions that would not otherwise hit the detector (see figure 2.3). Only BSE that are emitted in the small solid angle that covers the detector contribute directly to the signal. However, the SE2 detector signal can contain a significant indirect BSE component in the form of SE_3 generated by BSE hitting the chamber walls.

  The detector is very successful for imaging of non-planar specimens since

the placement of the detector offers great topological contrast. Intuitive interpretation of images is possible similar to how we normally perceive illuminated objects. The SE2 detector image can be perceived as observed down from the electron gun and illuminated by a light source at the SE2 detector.

- The In-lens detector is positioned in the beam path before the final lens. The detector utilizes an electrostatic field in the column used for decelerating the electron beam to the final accelerating voltage. The electrostatic field absorbs and re-accelerates incoming SEs towards the In-lens detector where they hit a scintillator similar to the SE2 detector.

  The In-lens detector detects almost exclusively SE_1 and SE_2. It detects more SE_1 than the SE2 detector and almost no SE_3 due to its position. The signal thus contains more surface information than the SE2 detector since the SE_1 and SE_2 are generated near the surface of the sample. Because the In-lens detector is placed directly above the specimen the detector images contain much less topographic contrast than the SE2 detector. As a result the In-lens image appears flatter, which is often an advantage for a quantitative analysis.

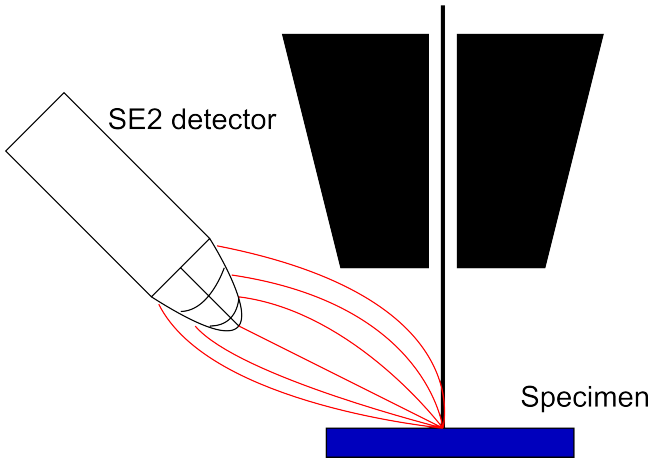

Figure 2.3: Illustration of the increased SE detection efficiency of the SE2 detector by applying a positive bias.

Other detectors such as the BSE detector and the energy selective and angle selective backscattered electron (EsB, AsB) detector can be used to acquire good compositional contrast since they detect only backscattered electrons. The position of the BSE detector obstructs the FIB in the Zeiss XB 1540 microscope, meaning that they can not be used simultaneously. This work has not had access

to experimental facilities with an EsB or AsB detector. As a consequence these two detectors are not covered further in this text.

To be able to draw any information from the gray scale images obtained from the detector signals, the image intensity must vary dependent on the local composition and shape of the sample where the electron beam impinges. Image contrast consists of three primary components.

- Compositional contrast is contrast between regions of the specimen with different atomic number. As mentioned above high atomic number results in a high backscatter coefficient. The detected signal, from detectors that detect BSE, is thus stronger when the beam impinges on a part of the specimen with a high atomic number than when it impinges on a part with a lower atomic number. This in turn results in an intensity difference at the corresponding image locations. The yield of SE is less dependent on atomic number but detection of SE_3 indirectly increases the signal from BSE.

- Topographic contrast is contrast between regions of the specimen with different angle of incidence between the specimen and the beam. The yield of both BSE and SE2 are dependent on the angle of incidence. The recorded image intensity is thus dependent on both the angle of incidence between the specimen and the beam but also the angle of incidence between the surface and the detector. It is this dependency that causes the shadowing effects seen in many SE2 detector images. If the beam hits a surface oriented away from the SE2 detector fewer electrons will hit the detector than if the surface had been oriented towards the detector. The use of a strong positive bias on the SE2 detector will reduce the shadowing effect since SE with trajectories away from the detector can be detected.

- Channeling contrast, or crystallographic contrast, is contrast generated by the orientation of the local crystal lattice relative to the incident beam. The penetration depth of the primary electrons is dependent on the crystal lattice orientation which in turn effects the yield and thus the contrast.

The implications of image contrast for the interpretation of the compositional structure of the analyzed SOFC samples is discussed further in chapter 5.

2.3 The focused ion beam

A FIB system [56] is similar to the system that creates the beam of electrons in a SEM. However, instead of electrons the beam consists of positively charged ions. A FIB system consists of two primary components:

- **The liquid metal ion source (LMIS)**. The LMIS typically contains a tungsten (W) needle attached to a reservoir of the metal source material. The most common metal is gallium (Ga), which has a low melting point (29.8 °C). The liquid Ga flows to wet the W needle and an electric field applied to the tip of the needle causes the liquid Ga to form a point source of 2-5 nm in diameter. A large extraction voltage (on the order of 10^8 V/cm) is applied to pull Ga from the tip and the Ga is ionized by field evaporation.
- **The ion column.** The extracted Ga^+ ions are accelerated through a potential (5-50 keV) down the ion column and focused into a thin beam on the sample surface through condenser and object lenses. Beam currents ranging from a few pA to several nA can be obtained by varying aperture sizes.

2.3.1 Ion-specimen interactions

The ion-specimen interactions can result in a number of events similar to the electron-specimen interactions. Three additional types of events that are important for FIB tomography are described here:

- Sputtering of atoms from the specimen surface can occur as a result of momentum being transferred from the incident ions to the atoms on the surface of the specimen. If the transferred kinetic energy is large enough to overcome the surface binding energy, the atom is ejected from the surface. The sputter yield, Y, is defined as the number of ejected particles per incident ion and typically varies in the range 10^{-1} to 10^2 depending on material and incident angle.

  The sputtering of atoms from the specimen is the primary functionality of the FIB system in FIB tomography as it enables the FIB to be used for the milling of sequential sections.
- Redeposition of sputtered atoms can occur when the FIB is used to mill deep and/or narrow trenches. A sputtered atom leaves the specimen with

a finite kinetic energy. If the atom along its trajectory again collides with the specimen it may be redeposited on the specimen. The redeposition rate is dependent on the kinetic energy of sputtered atoms, the sputter yield, the geometry of the milled features and the properties of the specimen.

- Gas deposition can be utilized to deposit controlled layers of metals or insulators. A controlled amount of precursor gas is injected into the specimen chamber close to the incident point of the FIB. By scanning an area with the ion beam, the gas will be decomposed into volatile and non-volatile components. The non-volatile component remains on the surface as a deposition layer. Typical deposition materials are metals such as W or Pt. The main usage of gas deposition for FIB tomography is to form a layer on top of the investigated sample that protects the sample. The surface layer also serves as a good surface for placing reference marks on (see section 2.5.1).

The products of the ion-specimen interactions can also be used for imaging in a similar way as in the SEM. However, this functionality has not been utilized in this work and will thus not be covered here.

2.4 FIB/SEM systems

SEM and FIB systems can be very useful tools on their own but the combination of the two in a single system offers many advanced capabilities. A FIB/SEM system such as the Zeiss XB 1540 consists of a FEGSEM column and a FIB column mounted with an angle of 54° between them. The sample is placed at the coincidence point between the two beams. The setup is illustrated in figure 2.1. High spatial resolution surface images can be obtained by the SEM and thin layers of material can be sputtered from the sample by the FIB. A FIB/SEM system allows rapid acquisition of FIB tomography data by automating the iterative task of image acquisition and sample milling while keeping the sample stationary.

2.5 Experimental FIB tomography techniques

A FIB/SEM system is a general tool used not only for FIB tomography. FIB tomography is thus often not the main purpose of a FIB/SEM system and certain experimental methodology is needed to acquire 3D data. Any SEM or

FIB parameter values reported in this section are based on experimental work on SOFC samples. Optimal parameter values could be dramatically different for other types of samples.

The data acquisition itself consists of two main parts: sample preparation and sectioning automation.

2.5.1 Sample preparation

The sample preparation (see figure 2.4) starts by polishing the sample such that the rectangular cuboid of interest, that the data will be acquired from, is close to the surface of the sample. A rectangular section of interest on the surface of the sample is selected. This rectangle corresponds to the top face of the rectangular cuboid of interest.

A protective layer of material (Pt or SiO_2 has been used) is deposited in a thin layer on the rectangle of interest (0.5-2 μm). The layer serves partially as protection of the surface from the ion beam during sample preparation, and partially as a dense surface to mill reference marks into.

Thin lines are milled in the surface of the protective layer. These fiducial marks serve as invariant structures that can be used for image alignment and to measure slice thickness (see section 2.6 and chapter 3). The fiducial marks are shown in the SEM view of figure 2.5

Trenches are milled on three sides of the rectangle of interest (see figure 2.4). The trenches serve two purposes. 1) They allow a free path for the electron beam to strike the front side of the excavated cuboid. 2) They allow sputtered material to be redeposited away from the sample such that the electron beam path is kept clear and redeposition on the cuboid itself is minimized. The depth of the trenches is determined based on the desired dimensions of the final cuboid. An experimental example of the setup is shown in figure 2.5

2.5.2 Sectioning automation

For the image acquisition process the sample is setup as in figure 2.4. The SEM views the sample surface at an angle of 36° from perpendicular. The acquisition of images proceeds in a loop that runs until data has been acquired from a sufficiently large section of the samples. The loop generally contains the following steps.

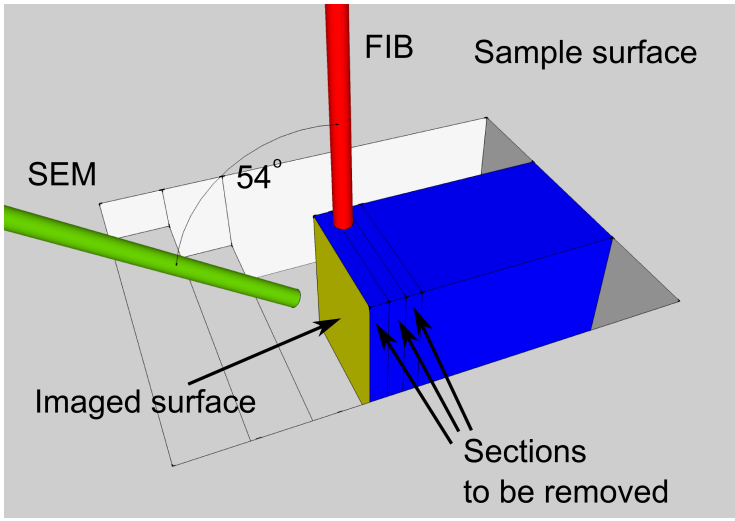

Figure 2.4: Illustration of the sample setup for automated sectioning.

1. **Remove a section from the sample**. The milling of a section is performed by continuously scanning the FIB in a line for a period of time. The time used to mill one section is dependent on the desired milling depth, the FIB current and the material of the sample. Increasing the FIB current increases the sectioning speed but also results in less smooth surfaces. The FIB current should thus be scaled with the size of the sample and the required voxel resolution. For voxel resolutions of approx 50 x 50 x 50 nm and cuboid dimensions of approx 20 x 20 x 20 μm FIB currents of 200-500 pA have been used for data acquisition in this work. Milling and dimension parameters are typically chosen to keep the removal of a layer on the order of 10s of seconds.

   After the milling of each line the FIB beam is translated to a new scan line slightly further into the sample. It can be an advantage to set the distance between scan lines lower than the desired layer thickness in the final data set. With small scan line increments (10 nm) and comparatively large FIB current (500 pA) the progress in each scan line is not visible due to the spread of the FIB. This means that the milling of the sample effectively occurs continuously rather than in discrete sections. Thus, the imaging of the sample surface does not need to be performed in synchronization with the completion of a FIB line scan. The layer thickness can then be controlled directly by the amount of time milled between images.

2. **Acquire an image of the new sample surface**. After the specified amount of the material has been removed the FIB is blanked (deflected

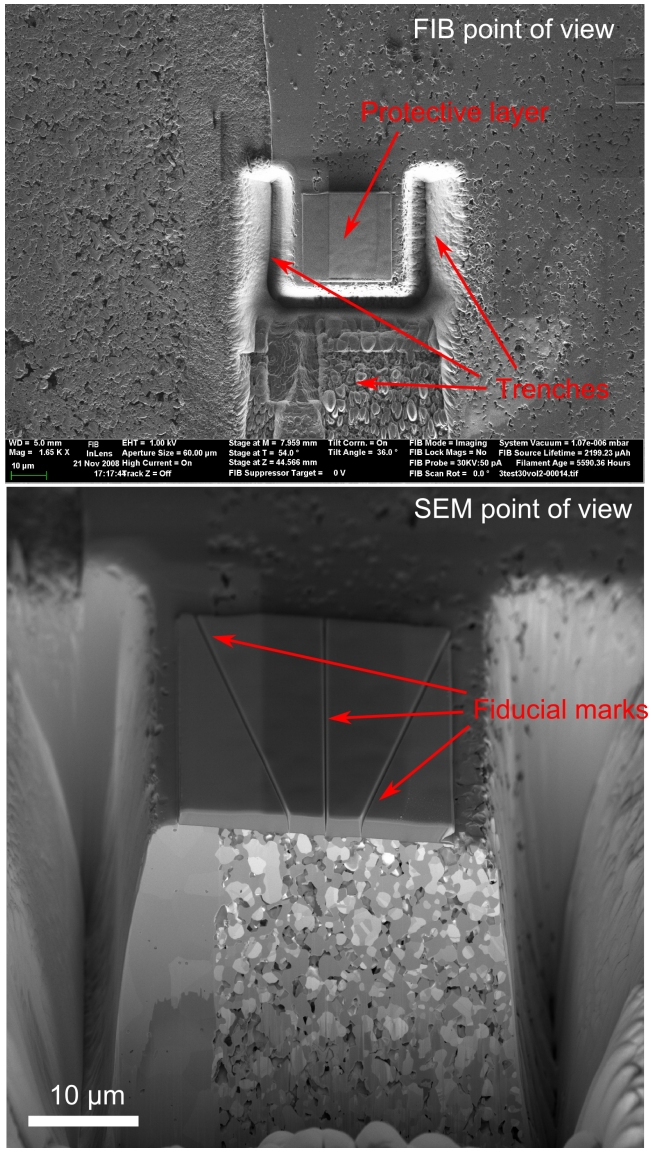

Figure 2.5: The sample preparation setup for FIB tomography as seen from the viewpoint of the FIB and the SEM. The FIB viewpoint image was acquired before automated serial sectioning began. The SEM viewpoint image was acquired during serial sectioning.

away from the sample) and the milling paused. The surface is scanned in the desired resolution with the SEM and simultaneous images are collected by one or more detectors. Acquiring images by more than one detector has certain advantages as is discussed in chapter 5. The time needed for image acquisition is on the order of 10s of seconds and dependent on the selected scan speed of the SEM. Slow scan speeds result in higher signal-to-noise ratio since the intensity at each pixel is based on more electrons. After the image acquisition step the FIB milling is resumed and the removal of a new section starts.

The time required for one iteration of the loop is typically 1-2 min. The loop is typically run for several hundred iterations as time permits. The result is a stack of collected images. For future reference x- and y-axes are introduced to denote the pixel column and row directions in the image stack. The z-axis denotes the image number direction. A pixel thus refers to an image coordinate (x,y) and a voxel to a stack coordinate (x,y,z).

2.6 Experimentally induced data analysis challenges

The geometry of the experimental setup for FIB tomography and the substantial time requirements result in a number of undesired artifacts in the raw image data. These artifacts must be corrected before a further analysis of the 3D structure can be performed.

- **Anisotropic pixel size**. The 36° angle of incidence of the SEM beam on the sample surface means that if no corrections are made each pixel will represent a rectangular part of the sample surface rather than a square one. This geometrical artifact can however be corrected for directly in the microscope software. After the correction the pixel dimension is the same as if the image had been acquired from a position perpendicular to the imaged sample surface. This correction option has been used for all imaging in this work.

- **Systematic drift**. The systematic drift of the imaging plane in consecutive images is a geometrical artifact similar to the anisotropic pixel size. If the raw image stack is viewed in sequence the boundary between the sample and the protective layer will appear as moving upwards in the image. The drift of the sample in the image is a consequence of the 36° angle

of incidence. The systematic drift is a direct function of how much of the sample is removed between consecutive images.

- **Random drift**. The random drift is caused by the slow movement of the sample relative to the SEM and FIB beams. Random drift can be the result of many factors: charge accumulation on the sample, mechanical movement of the stage, thermal changes in the sample and the microscope chamber and external vibrations or electromagnetic fluctuations. The random drift is smooth and unpredictable in nature as will be shown in chapter 3. It is not possible to separate the systematic drift from the random drift and as such the total drift is unpredictable. Chapter 3 describes how drift along the x and y axes may be automatically corrected.

  Drift along the z-axis is a major concern since strong drift in this direction can cause the FIB milling to leave thin flakes of material behind that obscure the SEM view. It is often this phenomena that ultimately limits how large continuous data blocks that can be collected. Minor drift along the z-axis will result in an uneven spacing between images. This drift can however be measured and corrected if needed (see chapter 3).

- **Non-uniform illumination**. The trench setup can cause a shadowing effect in images near the bottom or sides. The effect is caused by the trenches limiting the amount of electrons that reach the detectors. This results in non-uniform illumination of images where the upper part of the image is brighter than the lower part. For visual inspection the effects can be subtle but the non-uniform illumination must be taken into account for automatic analysis of the data. Methods for non-uniform illumination correction are covered in chapter 4

Chapter 3

Stack alignment

Abstract

Image data acquired by FIB tomography contain drift artifacts induced by the geometry of the data acquisition setup and movement of the sample relative to the SEM and FIB beams. Methods for systematically correcting these artifacts are presented. The methods work by detecting reference lines in the image data by the Hough transform. The images are then aligned to a common reference coordinate based on a spline fit of the detected line locations. The methodology is demonstrated on a LSC-CGO cathode sample. The methodology is shown to be robust to detection failures on individual images and has limited but adequate detection accuracy.

3.1 Introduction

This chapter presents methods for automatic correction of alignment caused by systematic and random drift as introduced in section 2.6. The alignment of the image stack is an important requirement to obtain accurate quantitative results in the later steps of the analysis. However, the alignment itself is not necessarily

challenging and builds on well established principles. The description here is included for completeness and describes the method that was used for alignment of image data in this work. The methodology is not fully matured and is still being improved.

Several contributors have used reference markers to align microscopy images. Mangan et al. [57, 58] described the use of fiducial marks as invariant markers that can be used to align images. Screier et al. [59] and Sutton et al. [60] analyzed the use of image correlation for sub-pixel accuracy surface deformation measurements. Many contributers have described image alignment for use in FIB tomography but few give specific details of the procedure. Kubis et al. [61] detected FIB z-direction drift by aligning a line milled in a surface deposited platinum layer. Holzer et al. [40] performed an automatic alignment of the FIB after each slice by pattern recognition algorithms, thus correcting for z-drift. Drift in the x-y plane drift was done during data processing by transforming the images to the cross correlation maximum.

The two central steps of the different alignment methodologies are the same. Detection of reference marks and transformation of the images such that the reference marks coincide in all images. The alignment can be done either during milling, as a post processing step or in combination. The methodology described here performs alignment as a post processing step.

3.2 The alignment problem

Figure 3.1 illustrates the alignment problem. The images are acquired from a stationary position at an angle of 36°. This means that the region of interest of the acquired image will have a higher position in consecutive images. Additionally, random drift (see section 2.6) can cause the region of interest to move unpredictably between consecutive images. To obtain data that corresponds to a rectangular cuboid of the sample space the drift corrected region of interest (Red rectangle in figure 3.1) must be extracted from each image. The image alignment is performed by translating each image such that the positions of a certain invariant reference mark coincide in all images.

Figure 3.1 also shows how large the random drift can become. In the latter image (right hand side) the sample has moved significantly to the right compared to in the first image. If only systematic drift was affecting the images only vertical drift would be observable. The entire region of interest must be visible on all images in order to be able to reconstruct a cuboid without missing data. To avoid this problem the initial SEM image dimensions are setup with a padding

region around the region of interest.

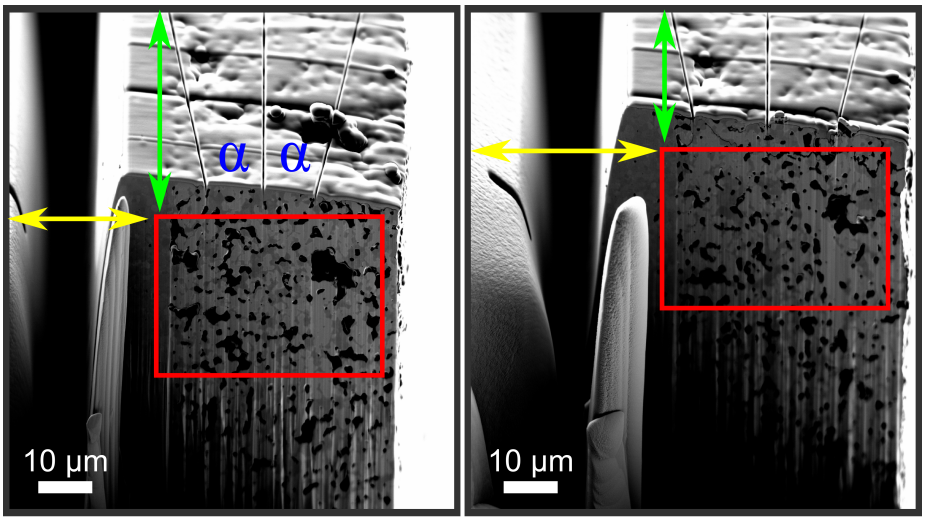

Figure 3.1: The alignment problem. The left image is acquired at the beginning of the image acquisition process and the right image towards the end. The two red rectangles represent the region of interest to be extracted from each image. The three milled lines used as reference marks are visible on the surface of the sample.

3.3 Reference marks

A pattern of three lines milled in the sample surface is used as reference marks (see figure 2.5 and 3.1). The pattern contains a center line in the same plane as the SEM and FIB beams. This line will appear as a vertical line in all images. Two additional lines are milled in the sample surface to the left and right of the center line at an angle α (20° has been used). A fourth line is indirectly formed by the boundary between the flat sample surface and the gas deposited protective surface layer. These four lines form the basis for measuring and correcting for the drift as illustrated in figure 3.2. The image alignment is independent of the physical scale of the sample. It is thus convenient to base the image alignment on the row and column coordinates of the image pixels.

The intersection point (x_{row}, x_{col}) between the center line and the surface boundary layer (green and cyan colored lines in figure 3.2) is used as the refer-

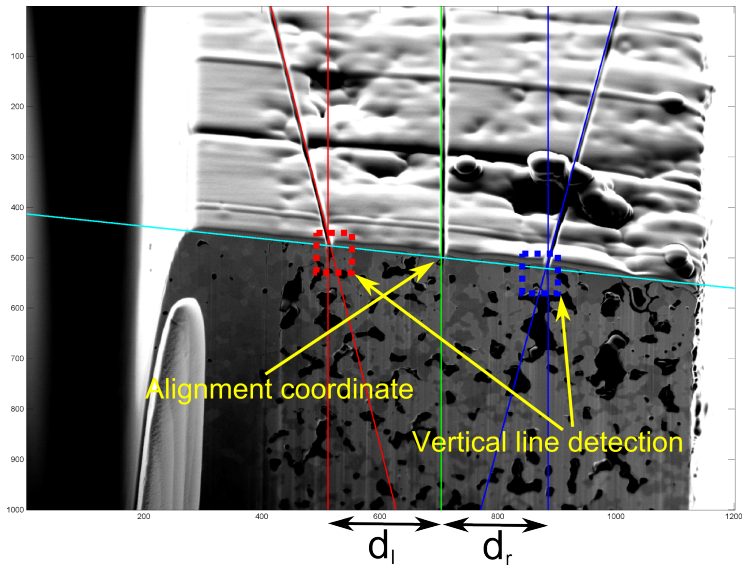

Figure 3.2: Detection of fiducial marks. See the text for explanation.

ence point to which all images are aligned. The intersection point is extracted in image coordinates and stored for further analysis.

The left (red) and right (blue) lines in the protective surface layer are used to calculate the milling depth. The angle, α, between the center lines and the left and right lines is known. The distance between the locations where the center and left/right lines intersect the milling plane can thus be used to calculate the milling depth. The milling depth, d_{depth}, measured in pixels can be calculated as

$$d_{depth} = \frac{d_{side}}{\sin\alpha\cos\beta} \tag{3.1}$$

where d_{side} is the distance (in pixels) from the center line to the left (d_l) or right (d_r) side lines. The term $\cos\beta$ corrects for the off horizontal tilt angle, β, of the sample surface/protective layer boundary line. The milling depth in physical units is obtained by multiplying d_{depth} with the horizontal pixel spacing. The milling depth is calculated from the arbitrary intersection point of the center line with one of the side lines. The relative depth of the current image in the stack is measured by subtracting the milling depth of the first image from the milling depth of the current image.

The thickness of the protective layer causes the side lines and the sample/protective layer boundary to not intersect in the milling plane. The intersection between these lines can thus not be used directly in the milling depth calculations. The lines milled in the protective layer cause vertical line artifacts to appear in the protective layer and the upper part of the sample due to uneven milling speed. The column coordinate of the vertical lines are used as the side line/milling plane intersection locations as shown by the vertical red and blue lines in figure 3.2. One side line is enough to calculate the milling depth, two increase redundancy and thus robustness.

3.4 Line detection

The detection of lines in the images is performed by the Hough transform (HT) [62] which makes the detection of lines simple. The HT is also employed in the automatic analysis of electron backscattering patterns [63]. The HT transforms an image of intensities into a parameterization of the lines in the image. The line parameterization describes a line by the length ρ and angle θ of its normal vector through the origin (illustrated in figure 3.3). The transformed image is thus a 2D parameter space where a line in the original image has the strongest response at a point in (θ,ρ) coordinates.

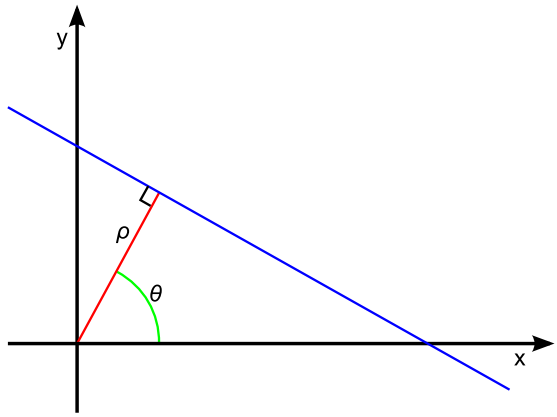

Figure 3.3: The line parameterization used in the Hough transform.

The detection of lines in the images is done by calculating the HT and finding the maxima. The HT will contain multiple local maxima. However, since the approximate angle and location of each reference line is known the search range can be reduced to a small rectangular region in θ-ρ space. This makes the

identification of each line robust, and the line detection of the three protection layer reference lines rarely fail. To increase robustness, the detection of the vertical side lines (red and blue vertical lines in figure 3.2) is performed on a smaller part of the whole image corresponding to the red and blue boxes in 3.2. The location of the smaller image parts are calculated relative to the intersection of the side reference lines and the sample/protection layer line.

3.5 Drift fitting

The detection of the alignment coordinates (x_{row}, x_{col}) and the side line distances d_r and d_l by the HT is robust but not necessarily accurate depending on the parameter resolution of the HT. To improve the accuracy and robustness the calculated coordinates and distances are not used directly. Instead, a spline is fitted to the detected locations. The spline f minimizes the expression

$$p\sum_{j=1}^{n}[y(j) - f(x(j))]^2 + (1-p)\int [D^2 f(t)]^2 dt \tag{3.2}$$

where $y(j)$ is the j^{th} data value at data site $x(j)$, n is the number of data sites, p is an input parameter between 0 and 1 and $D^2 f$ denotes the second derivative of f. The parameter p controls the curvature penalty such that low p values result in a rigid spline fit. Higher p values allow more curvature and thus a more interpolative fit.

To decrease the influence of possible detection failures the residuals with the 5% highest absolute values are removed after the first fit and a new spline fit is performed on the pruned data. The pruning makes the fitted alignment and distance data robust to detection failures on individual images.

3.6 Image alignment

The alignment of the image data is performed in a two pass setup. The first pass extracts the alignment coordinates and side line distances. The second pass performs the alignment. The final aligned stack of images is created by selecting a rectangular region in the first image and then translating consecutive images by their relative x_{row} and x_{col} coordinate differences from the first image.

3.7 Results

Figures 3.4, 3.5 and 3.6 show the detected values and fitted curves for a series of 365 images from a LSC-CGO cathode data set. A spline parameter of $p = 10^{-4}$ was used in all figures. The residuals of the fit without outliers are shown in the lower part of each figure. Note that the fit with and without outliers are very similar. The fit without outliers can thus be difficult to distinguish from the fit with outliers.

Figure 3.4 shows the vertical component of the alignment coordinate. The values on the vertical axis are the image row number where the alignment reference point was detected. The figure shows a number of outliers but the overall trend of the detected row locations is a curve with almost constant slope as we would expect from the systematic drift. Some minor random drift is observed as the slight curvature variation of the fitted curves. There is clear structure to be seen in the residuals. By observing a series of line detection results similar to the one in figure 3.2 the structure was found to be caused by the line detection routine. The vertical alignment component is mainly affected by the detection of the sample/protective layer boundary. However, this line is often interrupted by pore structures in the surface of the sample. This means that the true boundary occasionally will not have the highest HT value and an alternative line will be detected.

Figure 3.5 shows the horizontal component of the alignment coordinate. The values on the vertical axis are the image column number where the alignment reference point was detected. A constant detected column number would be expected if no sample drift was present. However, the figure shows considerable horizontal drift. The residuals in figure 3.5 are smaller, in the order of 1 pixel, and show less structure than the residuals for the vertical drift fit in figure 3.4. The better fit is caused by the vertical reference line being much clearer in the images. The HT detection is thus much more accurate and consistent.

Figure 3.6 shows the calculated milling depth for each image. The depth is calculated based on the average side distance $(d_l + d_r)/2$ and a pixel size of 58.6 x 58.6 nm. A single large outlier is found at image number 54. The outlier is seen to have a large effect on the fit with outliers but not on the fit without outliers. The residuals are large and show clear structure. The size of the residuals is explained by the low angle between the center and side reference lines. The low angle means that even small changes in the side distance d_{side} correspond to a large change in the calculated milling depth. The structure in the residuals is explained by the detection of the vertical side lines. Because the protection layer is relatively thin the line to be detected by the HT is relatively short and extends into the microstructure itself. This means that if a well defined vertical

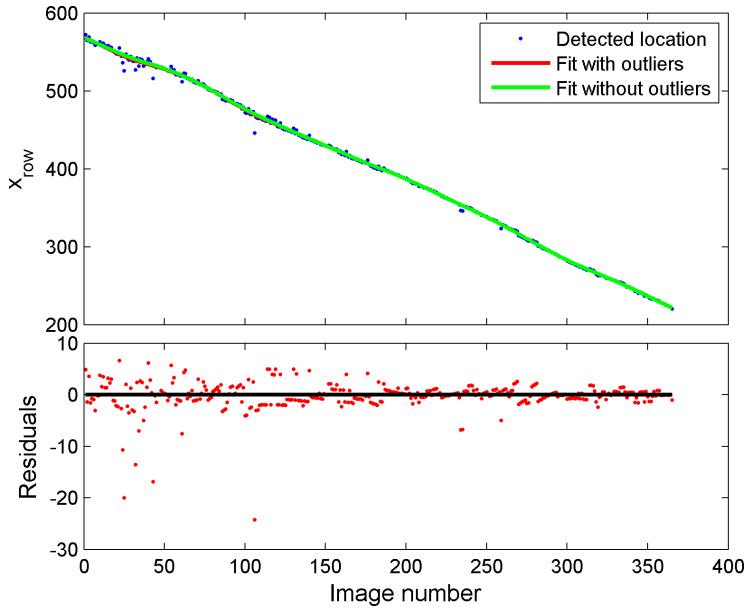

Figure 3.4: Detected image row coordinate values and fit. The residuals are calculated from the fit without outliers but contain all data points.

line structure is present in the microstructure near the protective layer the HT will detect this line instead of the vertical line in the protective layer. The overall trend of the fitted curve does not give any indication of significant z-axis drift as seen by the constant slope of the fitted curve.

The fitted coordinates of figure 3.4 and 3.5 were used to align the 365 images in the data set. A 3D box rendering of an aligned and cropped cuboid can be seen in figure 3.7. The image structure shows smooth phase boundaries without discontinuities in all three directions.

3.8 Discussion

The crossbeam XB 1540 microscope software used in this work has functionality to automatically detect a reference mark and perform an automatic drift correction. However, this functionality was found to be prone to failure and as such it could not be used for automatic correction. It was decided to base the entire alignment on the reference lines instead.

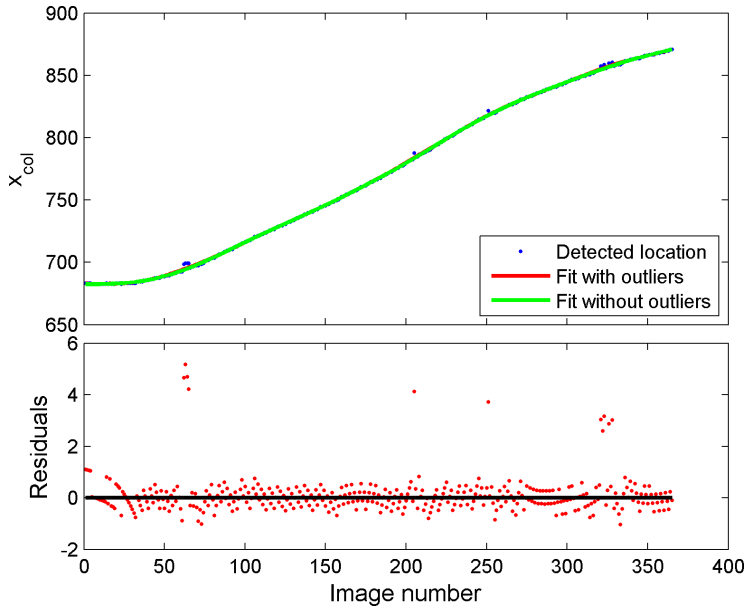

Figure 3.5: Detected image column coordinate values and fit. The residuals are calculated from the fit without outliers but contain all data points.

Since no invariant marker with high vertical stability is milled directly the vertical alignment has to be based on the sample/protective layer boundary. This border has turned out to not be an optimal marker in particular when a thin protective layer is applied. A horizontally milled line at the far end of the sample surface would likely perform better. Even though the horizontal detection accuracy is low the use of fitted curves instead of directly detected coordinates means that the alignment is performed on smooth coordinate changes. Smooth coordinate changes are important to obtain smooth phase boundaries in the aligned image data. The raw detected coordinates contain noise that if transferred to the alignment will result in a similar noise in the alignment of the phase boundaries. The alignment of the images by use of the sample/protective layer boundary is considered to be adequate but not optimal.

The detection of the milling depth is too inaccurate to be used directly to correct for potential z-axis drift. However, the milling depth plot provides a way of verifying uniform image spacing. An accurate and robust correction of the FIB during image acquisition would be optimal since this would allow the FIB to be adjusted for uniform image spacing.

For optimal results only the z-drift should be corrected during image acquisition.

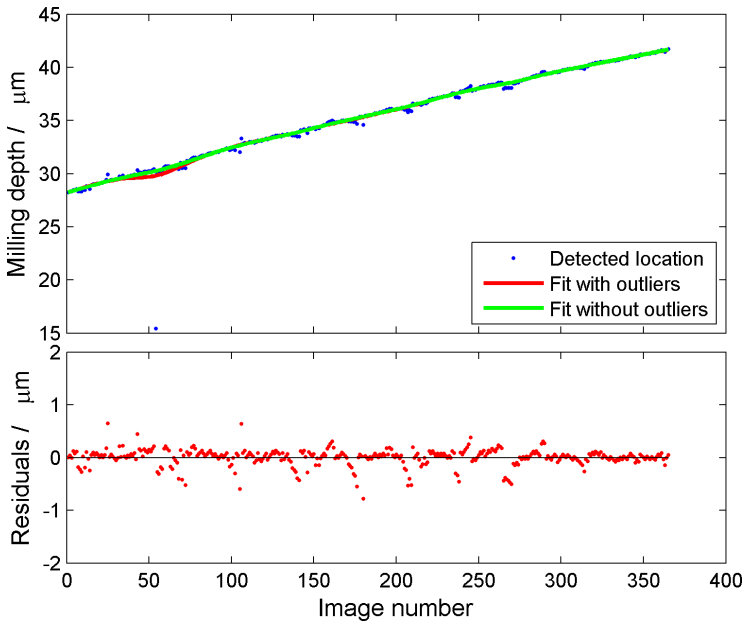

Figure 3.6: Detected milling depth values and fit. Note that the outlier at image number 54 has been excluded from the residual plot for visualization purposes.

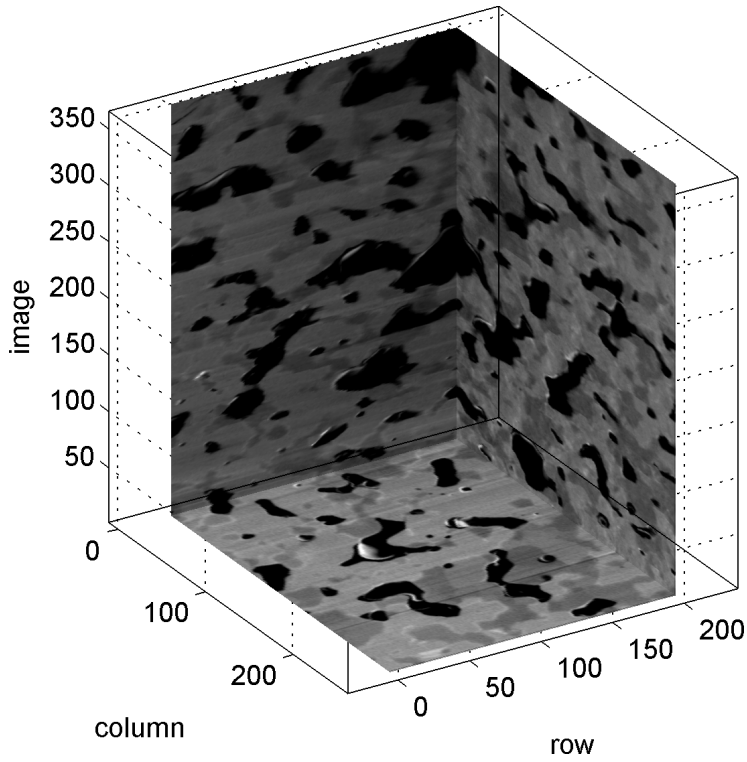

Figure 3.7: A rectangular cuboid of the aligned image data in row, column, image number coordinates

The ability to fit the detected coordinates as a post processing step provides increased detection robustness and can reduce detection noise.

3.9 Conclusion

A method for aligning serial sectioning image data was presented. The method performs the alignment by the detection of reference lines milled into a protective layer on the sample surface. The detection of the lines is made simple by the Hough transform. The alignment is based on a spline fitted to the detected coordinates rather than the raw coordinates. The spline fit with outlier removal provides robustness for detection failure in individual images.

The method was shown to provide adequate drift correction on an LSC-CGO cathode sample. The ability to calculate milling depth based on the distance between the reference lines provides a way of verifying uniform image spacing.

Chapter 4

Non-uniform illumination correction

Abstract

Non-uniform illumination of image data acquired by FIB tomography is caused by shadowing of electron detectors due to the geometry of the image acquisition setup. The effect is seen as a spatial gradual change of the overall intensity level of all phases inside the acquired cuboid of voxels. A method for automatic non-uniform illumination is presented. The method works by estimating the illumination field by fitting a hypersurface to sampled voxel intensities. The estimation of the illumination field allows the non-uniform illumination to be corrected. The methodology is fully automatic but relies on very basic subroutines. The method is shown to perform well on an LSC-CGO cathode SE2 detector data set.

4.1 Introduction

Non-uniform illumination as described in section 2.6 causes shadowing effects in the image data caused by the geometry of the image acquisition setup. Illu-

mination in this chapter refers to the ability of emitted electrons to reach the detector from a specific sample surface location. The effect can clearly be seen in figure 3.1 towards the bottom left of the images. The effect of non-uniform illumination is often subtle to the human eye but can have a large detrimental effect on an intensity based segmentation scheme.

The non-uniform illumination problem can be reduced by milling larger trenches or applying lift out techniques [64, 50] where the cuboid of the sample to be analyzed is completely removed from the bulk sample. These methods have been reported to almost completely remove the effect. They, however, further complicate the image acquisition process by an increased need for manual control and FIB milling.

This chapter describes a basic method for automatic correction of non-uniform illumination effects. Similar methodologies for illumination correction have found use in medical image data [65] but the application of these methods to FIB tomography data has to the author's knowledge not previously been reported.

The principle behind the presented non-uniform illumination correction method consists of three main steps. (1) Sampling of the spatial illumination changes. (2) Estimation of the illumination field by fitting a hypersurface to the samples. (3) Correction of the original intensity levels based on the fitted illumination field. The method takes as input an aligned stack of image data as described in chapter 3 and performs the correction on the whole 3D cuboid at once, rather than on an image by image basis.

4.2 Intensity sampling

To be able to estimate the illumination field, sample locations that are expected to have similar intensities under uniform illumination must be identified. The methodology described here relies on random spatial sampling and is motivated by two assumptions.

- The illumination has a slow spatial variation. This has been observed to be a valid assumption on experimental data.
- The individual phase structures in the cuboid are small compared to the cuboid volume and each phase has an approximately uniform spatial distribution.

The random sampling proceeds by selecting a large number, n, of randomly selected voxels from the aligned data cuboid. The idea behind the random sampling is that even though the random sampling is made from all phases the overall trend of the sampled intensities will still describe the slowly varying illumination field.

4.3 Illumination field fitting

A hypersurface is fitted to the n randomly sampled voxel intensities. Because of the random sampling it is important that the fitted hypersurface is rigid. If the fitted hypersurface is too flexible it will fit to the local phase structures rather than the slowly varying illumination field.

A simple regression model is used to fit the data:

$$\begin{aligned}\mathbf{I} =& \alpha + \mathbf{x}_r\beta_1 + \mathbf{x}_c\beta_2 + \mathbf{x}_i\beta_3 + \mathbf{x}_r\mathbf{x}_c\beta_4 + \mathbf{x}_r\mathbf{x}_i\beta_5 + \mathbf{x}_c\mathbf{x}_i\beta_6 + \mathbf{x}_r\mathbf{x}_c\mathbf{x}_i\beta_7 \\ &+ \mathbf{x}_r^2\beta_8 + \mathbf{x}_c^2\beta_9 + \mathbf{x}_i\beta_{10} + \epsilon \end{aligned} \tag{4.1}$$

where $\mathbf{I}$, $\mathbf{x}_r$, $\mathbf{x}_c$ and $\mathbf{x}_i$ are n-element vectors containing the intensity values and the (row, column, image number) coordinates of the voxel-samples respectively. The n-element vector ϵ is an error term. The unknown parameters α and $\beta_1 \ldots \beta_{10}$ are estimated by ordinary least squares. The regression model only contains first order, second order and mixture terms and thus represents a quite rigid hypersurface.

4.4 An iterative method

The random sampling scheme described in section 4.2 samples all phases in the cuboid under the assumption that the illumination field affects all phases equally. This is however not a good assumption if the cuboid contains pores since the pore phase is likely to have close to zero intensity in the entire cuboid independent of illumination. This means that the hypersurface least squares fit will be more influenced by the pore samples in regions with high illumination than in regions with low illumination because the intensity difference between the pore phase and other phases will be larger. A related problem is the presence of high intensity milling artifacts. These artifacts have the appearance of high

intensity vertical bands and mostly occur in the lower parts of the images thus causing a spatial hypersurface bias similar to the effect of the pore phase.

A sample pruning method is introduced to remove voxel-samples corresponding to the pore phase or high intensity image artifacts. The pruning method is an iterative scheme that alternates between the two steps of fitting and voxel-sample pruning.

1. A hypersurface is fitted to the vector of voxel-samples and the residuals of the fit are calculated.
2. The fraction, p_o, of residuals with the highest absolute values is removed from the vector of voxel-samples.

The scheme is continued until only a fraction, σ_p, of the original n samples remains. An example of the pruning scheme is given in section 4.6.

4.5 Illumination correction

A multiplicative illumination model is assumed

$$I_m(\mathbf{x}) = I_u(\mathbf{x}) I_f(\mathbf{x}) \tag{4.2}$$

$$\Leftrightarrow I_u(\mathbf{x}) = \frac{I_m(\mathbf{x})}{I_f(\mathbf{x})} \tag{4.3}$$

where $I_m(\mathbf{x})$ is the measured intensity, $I_u(\mathbf{x})$ is the uniform illumination intensity and $I_f(\mathbf{x})$ is the value of the fitted illumination field at the voxel $\mathbf{x} = (x_r, x_c, x_i)$. The uniformly illuminated cuboid of voxels is obtained from equation 4.3 by dividing the original voxel intensity with the intensity of the fitted hypersurface illumination field evaluated at the same voxel.

A multiplicative model is believed to be reasonable since a certain fraction of electrons emitted at a sample location is expected to be blocked primarily due to the angular distribution of electron trajectories. This, however, means that the illumination field can not be expected to be the same for different detectors at different geometrical positions. A separate non-uniform illumination correction is thus necessary for each set of detector image data.

4.6 Results

The correction method was applied to an aligned SE2 detector data cuboid from an LSC-CGO cathode sample. The cuboid had dimensions 401 x 301 x 250 voxels (row, column, image number). An initial 1 million voxel-samples were taken randomly from the cuboid ($n = 10^6$), the pruning was performed by removing 5 % of the samples in each iteration ($p_o = 0.05$) and the pruning iterations were continued until 40 % of the original samples remained ($\sigma_p = 0.4$).

Figure 4.1 illustrates the pruning process. Note that the histograms show the distribution of residuals of the hypersurface fit to the current samples from the entire cuboid while the sampling is only shown in one image. The three rows represent iteration 1, 8 and 18 (the last). After 8 iterations the pruning has removed all the pore samples from the vector of voxel-samples. After 18 iterations almost all samples are from the CGO phase (the brightest phase) and local bright artifacts are excluded. The absence of LSC phase voxels in the final sampling is a direct consequence of the continued pruning. After the pores have been pruned the LSC phase voxel-samples correspond to the fit residuals with the largest absolute value.

Figure 4.2 shows the result of the non-uniform illumination correction on image 10, 125 and 241 of the data cuboid. The new color map is introduced to better show the effect of the illumination changes. The original image data (left column) shows a drop in overall intensity from the top right towards the lower left and additionally an overall drop in intensity by increasing image number. The corrected image data shows a more uniform illumination of both the LSC phase (at approx. 0.7 intensity) and the CGO phase (at approx. 1.0 intensity) both across each image and between images. The vertical high intensity artifacts are clearly visible as the red structures in the lower part of the images.

Figure 4.3 shows a histogram of the intensity of all voxels in the cuboid before and after the illumination correction. The voxel intensities before the correction shows one broad peak centered at approximately 0.28. This peak corresponds to the combined voxels in the LSC and CGO phase. The histogram of the corrected voxel intensities shows better separation of the LSC and CGO phase voxels and an emerging two peak structure.

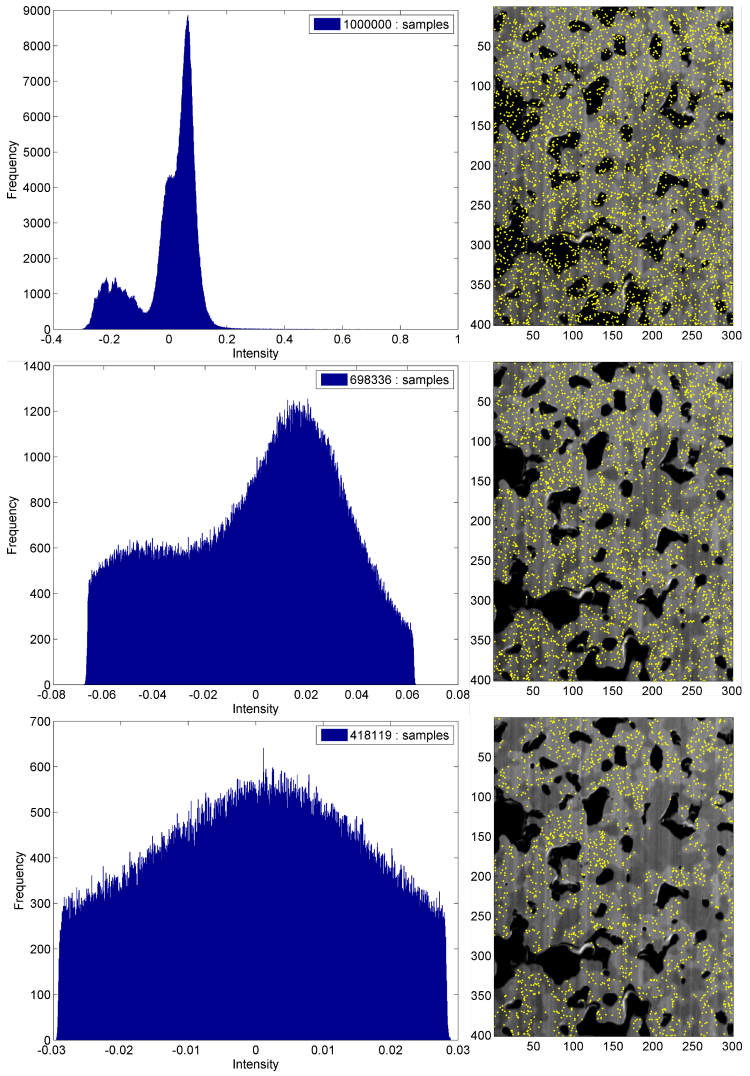

Figure 4.1: Iterative sample pruning. (Left) The histogram of the fit residuals for the entire cuboid. (Right) The corresponding spatial sampling shown on image number 10. The three rows represent 1.00, 0.70 and 0.42 million samples.

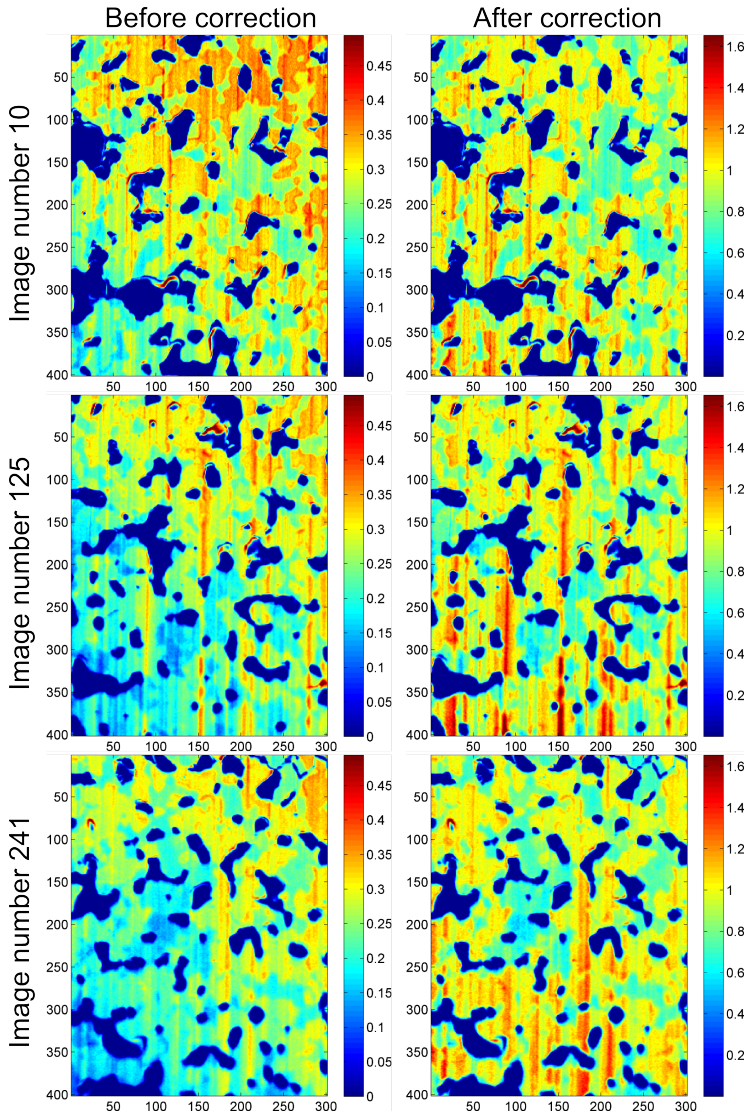

Figure 4.2: The image intensity structure before and after non-uniform illumination correction shown on image number 10, 125 and 241.

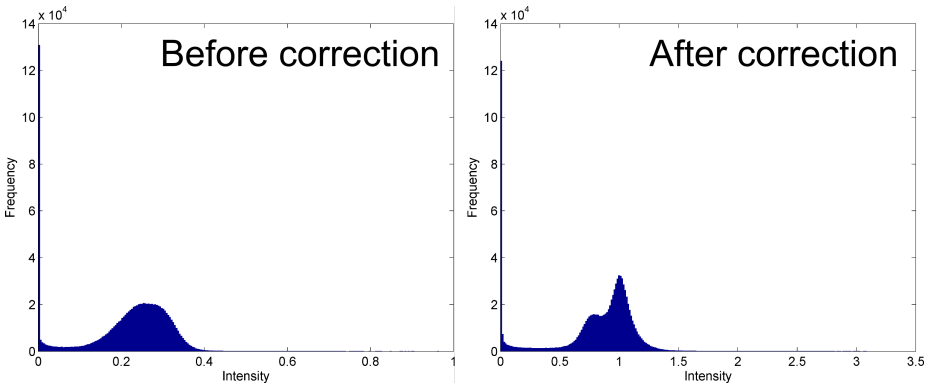

Figure 4.3: Histograms of voxel intensities in the cuboid before and after non-uniform illumination correction.

4.7 Discussion

Performing the illumination correction simultaneously in 3D has a number of advantages. The 3D sampling makes the fit less sensitive to local intensity structure changes in individual images. If the illumination field was fitted on each image one at a time the illumination field would likely contain field discontinuities between images. The field discontinuities would then result in intensity discontinuities in the corrected cuboid, thus introducing intensity artifacts from the illumination correction itself. By fitting a single hypersurface to the illumination field of the entire cuboid at once a continuous field is ensured in all three dimensions.

The example used in this chapter contains severe artifacts towards the lower parts of each image. These artifacts will typically be too severe to perform an accurate segmentation of each phase in the image and the cuboid should be cropped to not include these parts of the images. It is generally a good idea to perform the non-uniform illumination correction on a cuboid that is larger than the one to be used in the final 3D analysis. The hypersurface fit near the boundaries of the cuboid is supported by fewer data points compared to the center of the cuboid. The fit thus represents higher uncertainty since the illumination field variation structure outside the cuboid is unknown.

The presented methodology should be seen as an outline of a general non-uniform correction method. While the overall scheme of sampling, fitting and iteratively pruning voxel-samples is believed to be a sound strategy, the individual steps as presented here are very basic.

The type of hypersurface used for the fit should be investigated. The predetermined terms of the presented regression model does not provide much flexibility to fit to different types of illumination field structures. A spline based hypersurface with high curvature penalization is expected to perform better.

The iterative pruning should be carried out taking the spatial distribution of the samples into consideration to ensure an approximately uniform spatial pruning. The pruning as presented here could potentially remove all samples from a large region of the cuboid leaving the fit in the region unsupported by any data points. This was, however, not observed to be a problem on any of the experimental data that the method has been applied to.

The iterative pruning of samples on the shown LSC-CGO data succeeds because the number of voxels from the pore phase and high intensity artifacts is small compared to the number of LSC and CGO phase voxels. However, the methodology is not guaranteed to work on any data set. For instance, if the sample has high porosity the iterative pruning could result in the removal of high intensity phase voxels instead of the pore phase voxels since the high intensity phase voxels would constitute the residuals with the largest absolute values. The hypersurface would thus fit to the pore phase rather than the high intensity phases and the methodology would fail. One solution could be to simply remove all voxel-samples below a certain intensity threshold before applying the pruning scheme. Further studies and development of the sampling and pruning schemes is needed to guarantee successful sampling of future unknown data sets with varying porosity.

4.8 Conclusion

A method for automatic non-uniform illumination was presented. The method works by estimating the illumination field by fitting a hypersurface to sampled voxel intensities. An iterative scheme was applied that prunes the voxel-samples by alternating the fitting of a hypersurface and the removal of the residuals with the largest absolute values.

The method was shown to perform well on an LSC-CGO cathode SE2 detector data set. The corrected image data showed a uniform illumination after correction and better separation of the LSC and CGO phase intensities in the cuboid. Several suggestions for improving robustness of the method were given.

Chapter 5

A framework for automatic segmentation in three dimensions of microstructural tomography data

P. S. Jørgensen[a], **K. V. Hansen**[a], **R. Larsen**[b], **J. R. Bowen**[a]

(a) Fuel Cells and Solid State Chemistry Division, Risø National Laboratory for Sustainable Energy, Technical University of Denmark, Frederiksborgvej 399, 4000 Roskilde, Denmark.

(b) Department of Informatics and Mathematical modelling, Technical University of Denmark, Richard Petersens Plads, Building 321, Denmark.

Published in *Ultramicroscopy, 2010, Volume 110, p. 216-228*

abstract

Routine use of quantitative three dimensional analysis of material microstructure by in particular, focused ion beam (FIB) serial sectioning is generally restricted by the time consuming task of manually delineating structures within each image slice or the quality of manual and automatic segmentation schemes. We present here a framework for performing automatic segmentation of complex microstructures using a level set method. The technique is based on numerical approximations to partial differential equations to evolve a 3D surface to capture the phase boundaries. Vector fields derived from the experimentally acquired data are used as the driving forces. The framework performs the segmentation in 3D rather than on a slice by slice basis. It naturally supplies sub-voxel precision of segmented surfaces and allows constraints on the surface curvature to enforce a smooth surface in the segmentation. Two applications of the framework are illustrated using solid oxide cell materials as examples.

5.1 Introduction

Microstructural tomography data can be acquired through a number of different methods such as micro CT, serial sectioning by automated mechanical polishing and focused ion beam (FIB) tomography. The focus of this work is on FIB tomography but the methods described here are applicable to segmentation of tomography data in general.

FIB tomography is an established method for the analysis of nanoscale structures. The serial sectioning technique has been used often to acquire high quality data [46, 66, 67, 68]. The technique sequentially acquires an image of the sample surface followed by the removal of a thin layer of the sample by FIB milling.

Acquisition of high quality high resolution images of a sample is only the first step in obtaining accurate quantitative 3D measurements of constituent phase properties. The pipeline that processes the sample and extracts the wanted measurements in the end can be divided into the following equally important steps:

1. *Image acquisition*: The acquisition of images from a microscope. The

exact method is dependent on the equipment and the examined specimen.

2. *Data alignment*: The raw series of microscope images are typically aligned, cropped, and combined to create a voxel array (forthwith referred to as the voxel grid).

3. *Non-uniform illumination correction*: The geometry of the sample and the detector in the microscope will often cause a disparity between the number of electrons reaching the detector from one location on the sample surface compared to another. This disparity results in a shadowing effect on the resulting images where the global intensity of all phases changes gradually across the image. This bias field is often of no consequence in visual inspection, because the human visual system automatically compensates for the effect. For many automatic methods [65] it is however critical to correct for the bias in order to obtain accurate results. Non-uniform illumination correction equalizes the intensity level of sample areas with the same physical properties, over all voxels.

4. *Phase segmentation*: To be able to calculate accurate measures of phase properties it is necessary to perform a partition of the space, spanned by the voxel grid, into disjoint regions corresponding to each phase. At the most basic level, segmentation simply corresponds to assigning a phase label to each voxel in the grid.

5. *Measurement extraction*: The calculation of phase property measurements from the segmentation. For instance, volume fractions of phases, interface areas between sets of two phase or triple phase boundary curve lengths.

This work concentrates on phase segmentation. We present a framework that performs the segmentation simultaneously in 3D while keeping the boundary of the segmented region smooth. The framework incorporates sub-voxel precision of segmentation boundaries and avoids many of the segmentation errors often encountered with traditional thresholding schemes.

The research in FIB tomography data analysis so far can roughly be divided into two levels:

1. The first has worked on reconstructing high quality 3D voxel volumes of the acquired data incorporating methods of image alignment, interpolation and noise reduction [69, 45, 70]. This level of 3D reconstruction allows visualizing structures of interest but further processing is needed to perform quantitative measures.

2. The next level has sought to perform measurements of physical properties of the sample under investigation by segmenting the data into phases of interest and subsequently performing calculations on the segmented data to obtain the wanted measurements [47, 49].

The degree of difficulty when performing segmentation is highly dependent on the nature of the sample, and the structures to be characterized within the sample. In the case where each phase corresponds to a distinct intensity level of the image data, sufficiently good results can be obtained by simple thresholding of intensities. The contrast variations associated with most microstructural images are usually more complex. These structures can have large overlaps between the intensity of different phases, or the intensity of a phase can depend on the physical structure of its surroundings. In cases where a consistent texture difference between phases is present statistical methods have been shown to perform well [71].

Phase segmentation of FIB tomography data with complex structure has so far been performed primarily by manual delineation of phases [47, 72]. Performing the segmentation by hand is laborious, subjective and retards the use of FIB tomography for routine analysis of complex structures. Specialized techniques for automatic segmentation in specific applications have been developed, for instance by Munch et al. [43]. Common to the automatic methods is that they work serially image by image, often using a combination of filtering, thresholding and morphological operations [73, 74] using software packages like Matlab^{TM}or Amira^{TM}.

Here, we take the segmentation process to the next dimension. The segmentation framework presented below seeks to remove the reliance on 2D images for 3D segmentation purposes. Here, 2D images are the precursor to the 3D voxel grid; once constructed there are many advantages to treating the data as a 3D grid of data without any preferred viewing orientation. The framework introduces sub-voxel precision in the segmentation such that the final boundaries of a segmented phase need not lie exactly halfway between two voxels. The sub-voxel precision is strengthened by the possibility to penalize high curvature in the segmented phase boundary, thus naturally being able to represent smooth phase boundaries present in the original sample.

5.2 Phase segmentation

The framework is based on an active contours approach [75] and uses level set methods [76] to both evolve and represent the segmentation boundary. The idea

of using active contours and level set methods to automatically locate wanted structures in image data has been used both often and successfully in medical image analysis, primarily to automatically segment organs or brain tissue in images acquired by Magnetic Resonance Imaging, for instance in [77]. We are inspired by this methodology and adopt its use for FIB tomography.

Level set methodology was initiated by Osher and Sethian in 1988 [76] but has been extended over the years by many contributors. A comprehensive review of level set methods is given in [78]. The key idea for level set segmentation is to use a subset of voxels in the data set to define a robust initial classification of a phase. The subset contains those voxels that are known from prior knowledge to very likely be a certain phase. If we take porosity phase segmentation as an example, as we present in the following, the small subset could be voxels with a very low intensity. Very low intensity voxels are highly likely to correspond to the porosity phase and as such they are a robust guess for an initial general location of that phase. We call this subset of voxels the seed for that phase. Any classification scheme can in principle be used to create the seed. The seed is however only the base for further segmentation. The parameters of the classification scheme can thus be set at a moderate level that only includes correctly classified voxels in the seed, rather than having to weigh the number of false negatives against false positives.

Instead of performing the segmentation in one pass like threshold segmentation the level set methodology uses an iterative approach. The surface of the seed is iteratively evolved based on the intensity structure of the voxel intensities around it until a steady state is achieved. It is important to note that after the seed has been created the surface of the evolving phase class is no longer tied to the surface of the voxels but is a real and, implicit 3D surface.

In the following an outline of the framework is presented. A rigorous mathematical representation and implementation in detail is provided in Appendix A.

5.2.1 Surface representation

The most intuitive and easily understood way of representing moving phase surfaces is to think of a number of markers or buoys equally spaced along the surface. In this way we define the surface by connecting each buoy to its neighbour. The boundary can then be manipulated by moving each buoy one at a time. This technique has two main problems. The first is handling the splitting and merging of surfaces. The second is keeping an equidistant spacing between the buoys when they move. There exist techniques for handling these problems

and the method is generally successful in 2D. In 3D however, equidistant spacing of buoys becomes challenging and the buoy technique in general becomes complicated.

As an alternative to moving buoys we can define the surface as an implicit surface. In a discretized or explicit representation the individual points that belong to the interface are recorded. In an implicit representation the interface is defined as a certain iso-contour of a function. This gives meaning to the level set name. The interface is the *set* of locations at a certain *level* of some function.

The common choice for the level set function (ϕ) is the signed distance function. That is, the level set function at any location is defined as the distance to the closest point on the interface, negative on the inside of the surface and positive on the outside of the surface. A 2D example of an implicit contour representation is shown in figure 5.1.

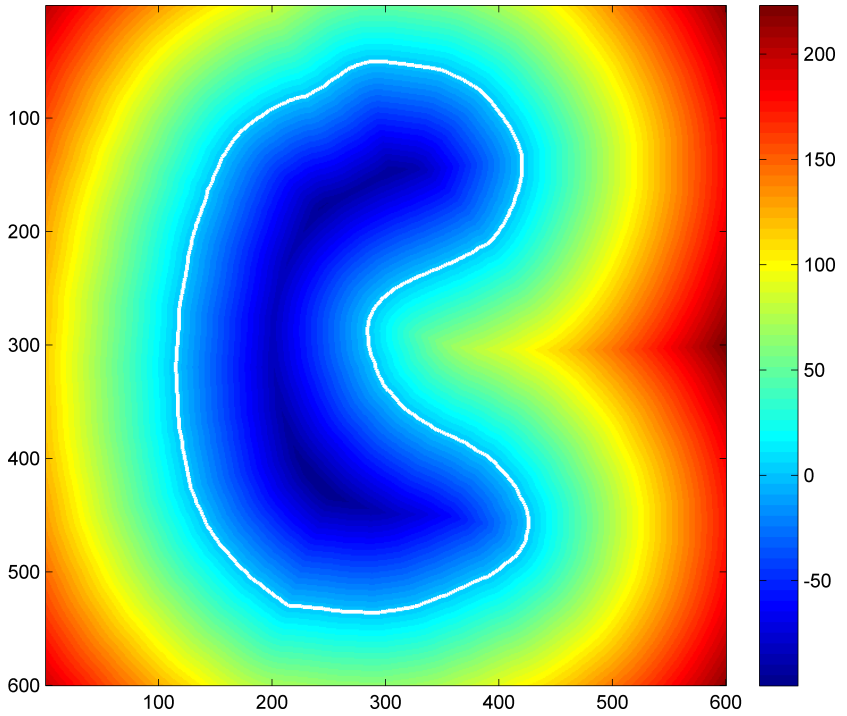

Figure 5.1: An implicit representation of a contour (in white) by the zero level set of a signed distance function. The image scale is in pixels.

Constraining ϕ to be a signed distance function has a number of advantages:

- The signed distance function is smooth near the interface. This property provides numerical stability when estimating partial derivatives using finite difference approximations.
- Implicit surfaces automatically handle splitting and merging of regions; there are no buoys to keep track of and the methodology extends trivially to 3D and beyond.
- There exist rapid methods of evaluating signed distance functions on a Cartesian grid.
- A signed distance function has the property of unit length gradients, which simplifies the evaluation of the level set equation.

5.2.2 Surface evolution

We add dynamics to the interface by letting the level set function evolve over time. In the original *active contours* approach [75] evolution was performed by summing up the contribution of velocities influencing each buoy, see figure 5.2. The level set method essentially works in the same way but instead of manipulating the surface directly, the entire level set function that implicitly contains the surface, is evolved.

The evolution of ϕ is performed by iteratively solving a simple partial differential equation (PDE):

$$\phi_t + \mathbf{V} \cdot \nabla\phi + \mathbf{a} = \mathbf{b}\kappa \tag{5.1}$$

where $\mathbf{a}$ and $\mathbf{b}$ are scalar fields, $\mathbf{V}$ is a vector field and κ is the mean curvature. The scalar fields $\mathbf{a}$ and $\mathbf{b}$ are often used as constants to weight the different terms against each other, but the more general formulation as scalar fields allows more flexibility in some cases.

Each of the terms $\mathbf{V} \cdot \nabla\phi$, $\mathbf{a}$ and $\mathbf{b}\kappa$ can be thought of as velocity vectors acting on the interface, trying to push it in a certain direction. The concept is depicted in figure 5.2.

The $\mathbf{V} \cdot \nabla\phi$ term describes evolution of the surface by the vector field $\mathbf{V}$. This vector field can be designed to evolve the surface in any direction. Note that in the example in figure 5.2 the vector field is constant across the domain, but is rarely so in practice.

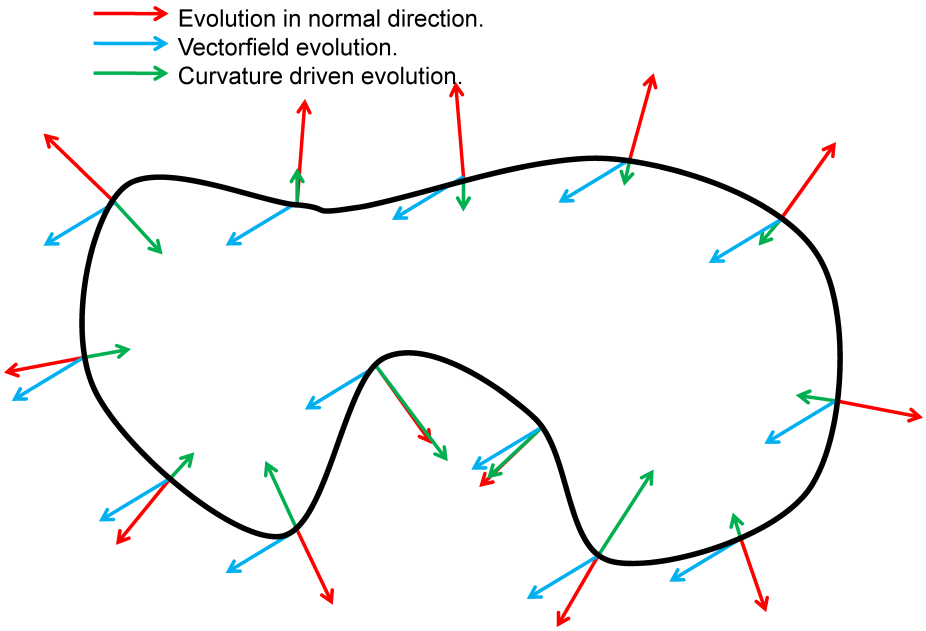

Figure 5.2: The evolution of the surface by velocity vectors

The **a** term describes evolution of the surface in the direction normal to the surface. This term will thus make the surface expand in areas where **a** is positive and contract in areas where **a** is negative. In the example in figure 5.2, **a** is set to a positive constant.

The $\mathbf{b}\kappa$ term describes evolution of the surface by mean curvature. If **b** is positive this term has a smoothing effect on the surface since it will contract convex surface areas and expand concave surface areas.

5.2.3 Solving the level set equation

So far the level set method has been described in continuous space. To be able to solve the problem on a computer a Cartesian grid is defined that spans the spatial domain. The level set equation is solved using numerical methods, replacing derivatives by finite difference approximations. The notion of a time step is used to describe the evolution of the surface from one iteration to the next. The size of the time step Δt thus reflects the amount of change that take place from one iteration to the next.

To evolve the surface the time step Δt forward in time, the update rule:

$$\phi^{n+1} = \phi^n - \Delta t \left(\mathbf{V} \cdot \nabla \phi^n + \mathbf{a} - \mathbf{b}\kappa^n \right) \tag{5.2}$$

is applied at each grid position. The term ϕ^{n+1} is the value of the level set function at the next time step, ϕ^n is the value of the level set function at the current time step and superscript n denotes evaluation at the n^{th} time step.

At the basic level, a level set segmentation algorithm begins by creating an initial signed distance function from the seed. The update rule is then applied iteratively for all voxels until the evolution of the surface converges to a steady state. See Appendix A for details.

5.2.4 Designing scalar fields and vector fields

Equation (5.2) has three spatially dependent fields used to control the evolution of the level set: the vector field $\mathbf{V}$ and the two scalar fields $\mathbf{a}$ and $\mathbf{b}$. These fields will typically be derived from the intensity level of the underlying voxel intensities.

The scalar field $\mathbf{b}$ is the simplest to define. It controls the penalty on the mean curvature of the interface. It is often set as a constant for the entire domain, based on prior knowledge of the smoothness of the structures to be segmented. If there is prior knowledge of the smoothness of the structure as a function of position in the sample, it can be modelled with a spatially changing **b.**

The scalar field $\mathbf{a}$ controls the evolution in the normal direction. This scalar field is typically set as a direct function of the corresponding intensity in the voxel grid. At the basic level the term can be used to perform a thresholding like segmentation. Let τ be the intensity threshold that we would use in a simple thresholding segmentation and $I_{\mathbf{x}}$ the intensity value at $\mathbf{x}$. The scalar field $\mathbf{a}$ is then set at a positive value where $I_{\mathbf{x}}$ is above the threshold and at a negative value where $I_{\mathbf{x}}$ is below the threshold. The initial surface from the seed will then keep expanding while it crosses voxels with an intensity value above the threshold and will stop when it reaches a value below the threshold. The advantage over normal thresholding segmentation is that we can set the seed independently of the threshold. As mentioned previously the seed can be made from voxels that are very likely to be a specific phase. By seeding robustly and then expanding the interface as described above we effectively perform a thresholding where all parts of the final segmentation are connected to the initial

seed. This seed driven procedure enhances the robustness of most traditional thresholding schemes and curvature constraints can easily be added.

A good class of functions for mapping intensities to values of $\mathbf{a}$ has the form

$$\mathbf{a}(I_x) = \sigma \tanh(\lambda(I_x - \tau_a)) \tag{5.3}$$

where σ is a scaling parameter, τ_a controls the threshold where the function changes sign and λ controls the sharpness of the transition from negative values to positive values. The parameter σ will typically only need to be changed if the intensity value range changes. The λ parameter can be set to reflect the confidence we have in the threshold to provide a good segmentation on its own. For small values of λ there is an almost linear mapping where $\mathbf{a}(I_x)$ changes slowly around the threshold τ_a. For very high values of λ the mapping is similar to a Heaviside step function and is best used in situations where there is complete confidence that the intensity threshold will provide a good segmentation on its own. Figure 5.3 shows plots of the function for different values of λ.

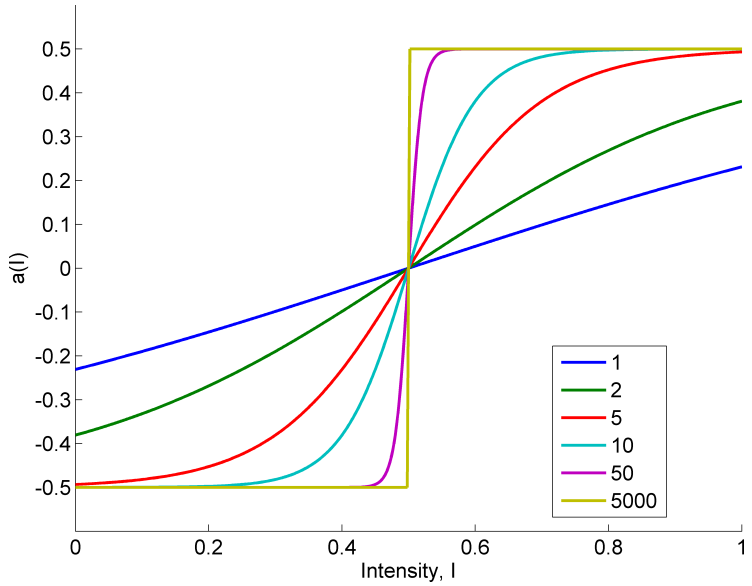

Figure 5.3: Plots of equation (5.3) for 6 different values of the sharpness parameter λ for intensities in the range $[0, 1]$. $\sigma = 0.5$ and $\tau_a = 0.5$.

The $\mathbf{V}$ vector field controls the evolution in an externally defined direction. It can be used to direct the interface towards or away from specific intensity

structures in the voxel grid. An often used application is edge attraction, where the vector field will direct the surface towards high intensity gradients, also known as edges. Let $g_{\mathbf{x}}$ denote the gradient of intensities at location $\mathbf{x}$. For a voxel grid, $g_{\mathbf{x}}$ will be a 3 component vector in the direction of steepest intensity increase. A new scalar field is created by storing the magnitude of the gradient $|g_{\mathbf{x}}|$ at each location. This scalar field has high values at locations corresponding to edges in the voxel grid. The $\mathbf{V}$ vector field can then be defined as $\mathbf{V} = g(|g_{\mathbf{x}}|)$, the gradient of the intensity gradient magnitude. This expression for $\mathbf{V}$ will attract the interface to edges in the voxel grid. The concept is illustrated in figure 5.4. Multiple expressions can be added together as we shall see in the example below.

5.3 Results

The two application examples presented below are both from solid oxide cell (SOC) materials. Phase segmentation of constituent phases in SOC electrodes and the subsequent quantitative analysis of phase microstructure is a valuable tool in the development and understanding of SOCs. They are presented here both as the motivation for the development of the methodology and as practical segmentation examples of how algorithm parameters can be designed. An interpretation of the segmentation results in an electrochemical context is thus beyond the scope of this work. The example in section 5.3.1 is used to illustrate how a novel design of the externally generated vector field may be used to perform a segmentation of data with a challenging intensity structure rather than to show a best practice of porosity segmentation. The example in section 5.3.2 represents a less challenging segmentation problem showing how intensity information can be used together with seeding and edge information to perform a robust coherent segmentation. The reader is referred to Appendix B for experimental details regarding image acquisition.

5.3.1 Porosity segmentation

The first sample is a porous cermet of Nickel/Yttria Stabilized Zirconia (Ni/YSZ) that functions as the active hydrogen electrode in SOCs manufactured at Risø-DTU [79, 80]. The empty pores in the sample occasionally result in a high electron count deep within them such that some surfaces, below the raw image plane are visible. This effect (see figure 5.5A) causes some regions that correspond to pores in the raw image plane to have intensities similar to the Ni or YSZ phase.

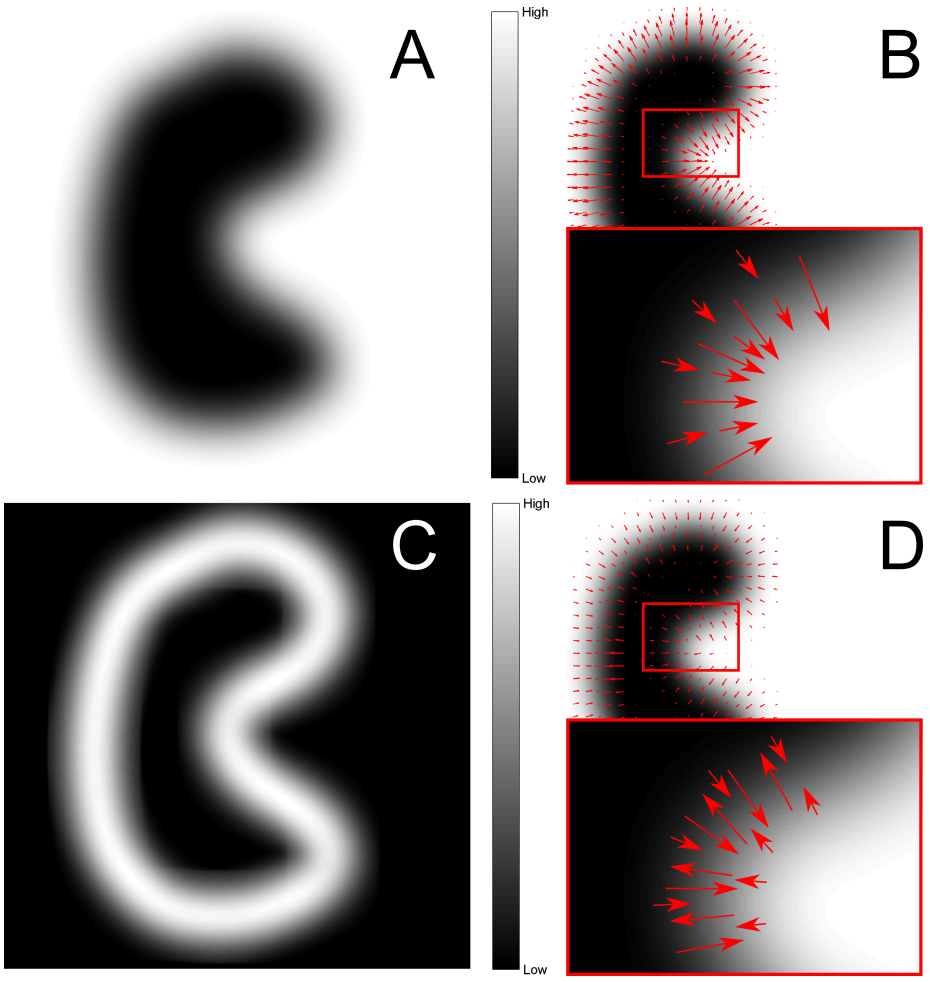

Figure 5.4: Creation of an edge attraction vector field. (A) The original image showing a low intensity structure on a high intensity background. (B) The gradient vector field superimposed on the original image. The inset shows an enlarged view of the region inside the red rectangle. (C) The magnitude of the gradient vector field. (D) The edge attraction vector field obtained as the gradient of the intensity gradient magnitude. The inset shows an enlarged view of the region inside the red rectangle.

Figure 5.5 shows the challenge in the segmentation. Figure 5.5A shows an example of the raw image data from the Ni/YSZ electrode data set with the result of the level set porosity segmentation superimposed. The contours are created from a cut through the surface of the full 3D segmentation. The dark phase is the porosity phase, the brightest phase is Ni and the slightly darker phase is the YSZ phase. The Ni and YSZ phases have similar intensities but the Ni phase can generally be distinguished by giving the impression of being raised compared to the YSZ phase. Note how deeper lying Ni/YSZ phases can be seen through the porosity phase, causing the porosity intensities to be heterogeneous. Figure 5.5(B-C) show porosity segmentation by thresholding at 0.38 and manual segmentation. Figure 5.5D shows a comparison of the level set and thresholding result with the manual segmentation. Figure 5.5E shows the histogram of intensities of voxels segmented as porosity by the level set method. Figure 5.5F shows a histogram of intensities in the raw image data.

Performing a threshold-like segmentation of the porosity will result in a large number of false positives as indicated in Figure 5.5(B-D) because the porosity phase contains a large portion of high intensity voxels. A threshold at 0.38 was used for this comparison based on simple inspection of the histogram in Figure 5.5F. The intensity histogram and the location of errors in figure 5.5F suggest that increasing the threshold value will result in more false negatives and decreasing the threshold will result in more false positives. By inspection of thresholds made at different values this was verified to be the case. The threshold might thus not be optimal but it shows the trade-off between false positives and false negatives that must be made.

The level set segmentation was performed by first creating a seed from voxels with low intensity. This ensures that only true porosity phase voxels are included in the seed. The curvature penalizing scaling parameter b is set as a constant 0.1 for the entire domain to reflect the overall smoothness of the wanted physical phase boundary. The scalar field $\mathbf{a}$ controlling the evolution in the direction normal to the surface is constructed as a function of the voxel intensity at each voxel position, $\mathbf{a}(I_{\mathbf{x}}) = 0.38 - I_{\mathbf{x}}$. The $\mathbf{a}$ term is thus positive for intensities lower than 0.38 and negative for intensities higher than 0.38, where voxel intensities are in the range 0 to 1. The design goal for $\mathbf{a}$ is to allow the initial seed to expand into nearby low intensity areas and provide a global penalty for expanding into high intensity areas. The actual intensity value where $\mathbf{a}$ changes sign is not critical since the $\mathbf{a}$ contribution will be negligible near this value and the evolution by externally generated vector field will be the main contributor.

The externally generated vector field $\mathbf{V}$ is the weighted sum of two vector field contributions: an edge attraction field $\mathbf{V}_{edge}$ constructed as outlined in section 5.2.4 above and an upslope expansion field $\mathbf{V}_{slope}$, thus $\mathbf{V} = c_{edge}\mathbf{V}_{edge} + c_{slope}\mathbf{V}_{slope}$. The upslope expansion field is essentially the intensity gradient,

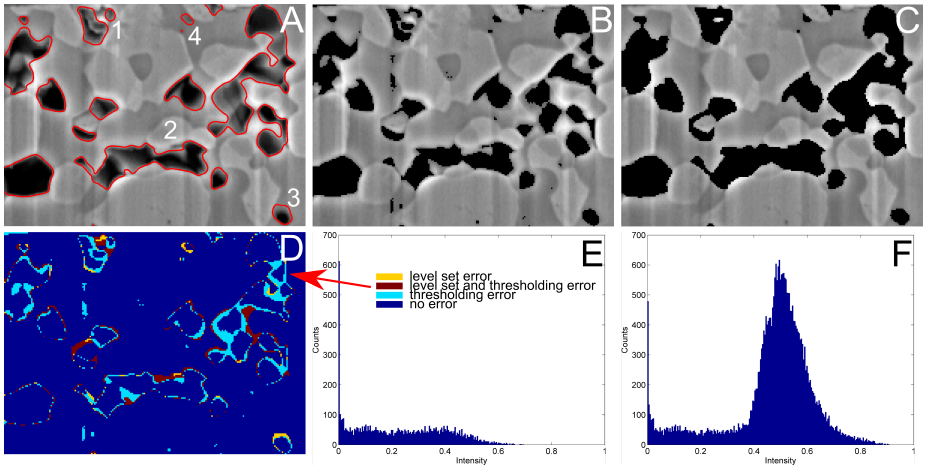

Figure 5.5: The challenge of the porosity segmentation. (A) An example of the raw image data from the Ni/YSZ electrode data set with the result of the level set porosity segmentation superimposed. (B) Porosity segmentation by thresholding at 0.38. Black regions have values below the threshold. (C) Porosity segmentation done by hand. (D) Comparison of level set segmentation and thresholding segmentation with the manual segmentation. The colour coding shows which method has errors in a particular location. (E) The histogram of intensities of voxels segmented as porosity by the level set method. (F) The histogram of intensities in the raw image data.

and causes the surface to expand across rising intensities and be repelled by falling intensities. The vector field exploits the fact that the seed is created around local intensity minima and that the phase structures that can be seen at the bottom of the pore phase generally have smooth surfaces that rise in intensity towards the true porosity surface. The weighting parameter c_{edge} was set to 0.04, c_{slope} was set to 0.45 and the seeding intensity threshold τ_{seed} was set to 0.2917 corresponding to 70 in the original 8 bit encoding.

Since the value 0.38 was used for both the value where $\mathbf{a}$ changes the sign and the threshold in the comparison the segmentation in figure 5.5B also represents the locations in the image where $\mathbf{a}$ has a positive value. This shows the significance of the upslope contribution to the externally generated vector field since the surface will only expand based on the $\mathbf{a}$ term to the extent shown in figure 5.5B. The rest of the expansion is mainly driven by $\mathbf{V}_{slope}$, pushing the surface towards rising intensity. Observe how the level set segmentation contains a significant amount of voxels with intensities above the 0.38 mark in figure 5.5E.

The comparison of the level set method with manual segmentation has previously been discussed in [80]. The errors shown in Figure 5.5D correspond to a voxel misclassification rate (false positives and false negatives) of 14.4% for the level set segmentation and 32.2% for the threshold segmentation and a relative error of segmented porosity area of -9.59% for the level set method and -28.91% for thresholding. Note the distinction between the amount of wrongly classified voxels and the segmented porosity area. It is possible to tune parameters for most segmentation methods (including level set segmentation and thresholding) such that the error in segmented volume is zero. This does however not necessarily mean that the segmentation is correct but rather that the amount of false positives equals the amount of false negatives. The segmented phase interfaces will ultimately be used for calculations of microstructural measures such as triple phase boundaries and percolation. For these calculations, it is arguably more important that the segmentation is topologically correct and that the segmented interface boundaries match the physical interface boundaries.

The first and last iteration of the segmentation can be seen in figure 5.6. Observe the coherent structure of the segmented surface. The red contour on the upwards facing side of the cubes illustrates how bright intensities inside a pore are eventually included in the segmentation as the evolution converges in cube B. Figure 5.7 shows the volume fraction of the segmented region as a function of the iteration number. The segmentation has an asymptotic like convergence. This slow convergence is caused by the $\mathbf{V}_{slope}$ vector field. As illustrated above, $\mathbf{a}$ will be negative in high intensity areas thus making the surface contract. The $\mathbf{V}_{slope}$ field will make the surface expand into voxels with rising intensity. At the converged state the two terms cancel each other to make the interface stationary. The iterations leading up to this equilibrium will however move the

interface by an increasingly small amount thus resulting in slow convergence.

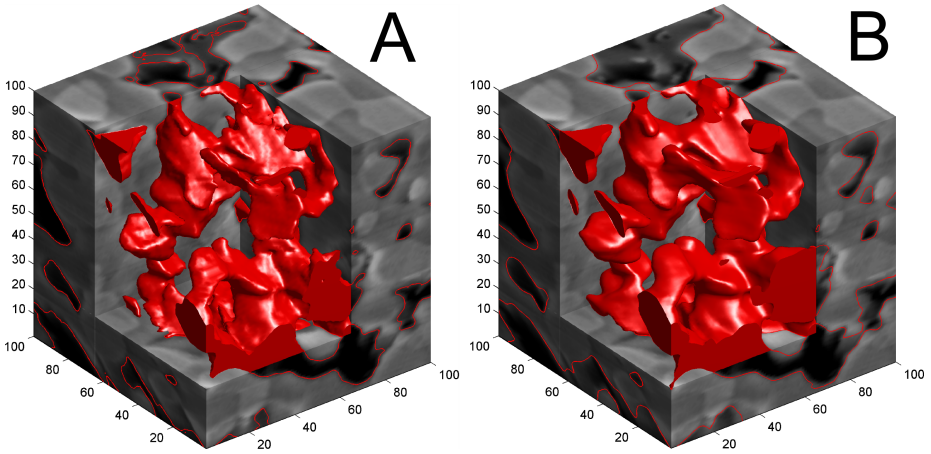

Figure 5.6: A 3D cube view of the porosity segmentation on a subset of the Ni/YSZ anode data at the first and last iteration. The sides of the cube show the voxel intensities of the raw image data. The red contour lines denote the intersection between the segmented surface and the boundary of the voxel cube. The red surface is a 3D rendering of the segmented porosity surface. The 2 cubes A and B correspond to iteration 1 and 57 (the final iteration). The cube has dimensions 100 $\times$ 100 $\times$ 100 voxels, corresponding to 4.19 $\times$ 4.19 $\times$ 4.71 μm.

5.3.2 Robust thresholding

The second sample is a composite SOC cathode consisting of 50/50 weight-percent strontium-substituted lanthanum cobaltite (LSC) and gadolinia-substituted ceria (CGO), similar to [81].

The data set was acquired using simultaneous acquisition of data from the In-Lens and SE2 detectors (see figure 5.8 and Appendix B). This provides a bimodal vector data set, where two values are available at each voxel. In this example, the two modalities will be used sequentially but the natural extension is to utilize the distribution of phase intensities in the two dimensional space spanned by the two modalities.

The porosity phase is segmented as described above utilizing the SE2 modality data. The CGO phase is segmented utilizing the InLens modality data and modifying the normal expansion scalar field to repel the evolving surface from

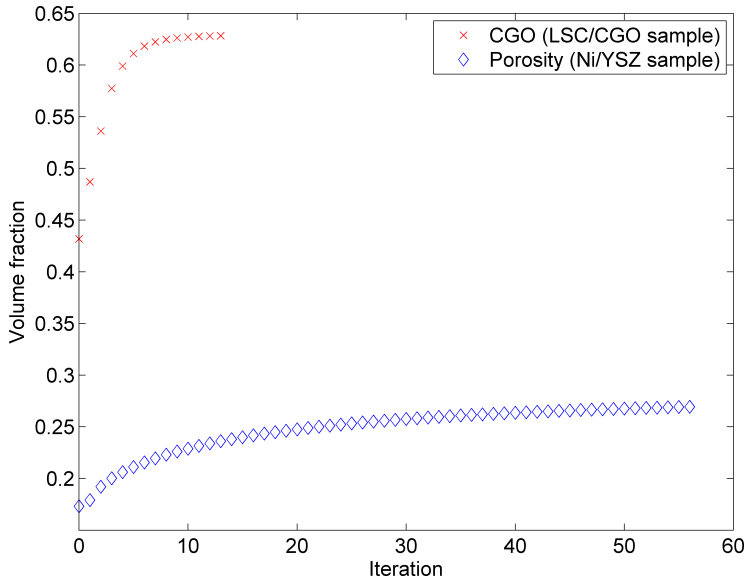

Figure 5.7: Segmented volume fraction as a function of iteration number for the LSC/CGO sample and the Ni/YSZ sample.

the already segmented porosity phase.

The CGO phase seed is created using a further constraint on seed voxels. To be part of the seed, a voxel must both have a high intensity value (above $\tau_{seed} = 0.8313$) but also a low gradient magnitude ($|g_{\mathbf{x}}| < 1/255$). This scheme ensures that only regions of relatively constant high intensity voxels will be included in the seed and makes it unlikely that small bright image artefacts are included in the seed. The function $\mathbf{a}(I_x) = 0.5\tanh(\ 20\,(I_x - 190/255\ )\ )$ is used to map intensities to values of $\mathbf{a}$. An edge attraction field with $c_{edge} = 0.1$ is applied as described earlier and the curvature penalizing scaling parameter is set as constant 0.1. The first and last iteration of the segmentation can be seen in figure 5.9. The volume fraction of the segmented phase as a function of iteration number can be seen in figure 5.7. The convergence is much quicker than seen with the porosity example. This is because the final resting place of the interface is mainly influenced by the $\mathbf{V}_{edge}$ vector field. The $\mathbf{V}_{edge}$ vector field has a strong contribution close to strong edges (see figure 5.4D). The interface will thus quickly be pulled towards the strong edge and remain there. Figure 5.7 indicates a volume fraction of CGO of 0.628 and it should be noted that this is a high value for the volume fraction of CGO for a 50/50 weight percent mixture of LSC/CGO. The distribution of CGO in the sample is however heterogeneous. The segmentation run on a larger $220 \times 300 \times 200$ voxel region of

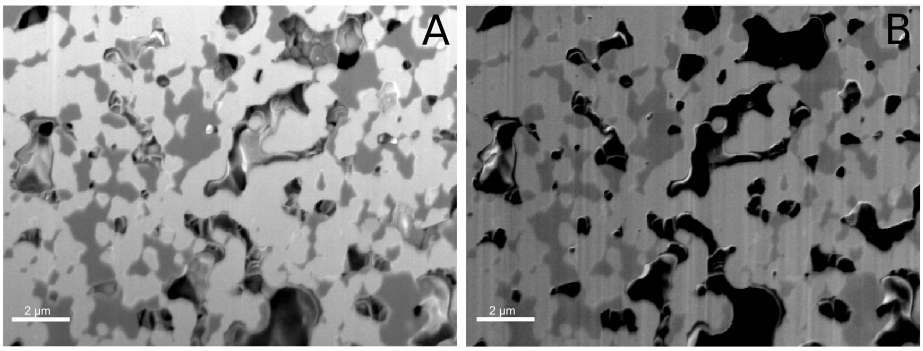

Figure 5.8: A dual detector image set of a LSC/CGO cathode. (A) The image acquired by the InLens detector. White is the CGO phase, grey is the LSC phase. The porosity is not easily distinguishable. (B) The image acquired by the SE2 detector. Light grey is the CGO phase, gray is the LSC phase and black is porosity.

the same sample with the same parameter settings resulted in a CGO fraction of 0.522. An even larger region would be preferable for reliable volume fraction measurements.

5.4 Discussion

The automatic segmentation scheme presented above works naturally in 3D. The final boundaries of the automatic segmentation will thus always form a coherent and smooth 3D surface when evolutions converge. This is a distinct advantage compared with an image by image segmentation. The 3D coherence enables the level set method to compensate for ambiguities in one slice of the data by enforcing a coherent smooth structure from the slices above and below it. This advantage can be exemplified by the task of porosity segmentation. In a 2D image, a region that is slightly darker than its surroundings could be interpreted as one of three possibilities: 1) it could be due to a pore just beneath the current image, 2) it could be due to a pore just above the current image, or 3) it could be an artefact or noise. The 3D segmentation handles this ambiguity by evolving the segmentation surface from a robust seed. If the dark area is due to a pore the segmentation boundary will naturally evolve to form a coherent pore structure from above or below. If the dark area is due to artefacts or signal noise it will not be part of the seed and will not be included in the segmentation. The seed will not evolve to include the area because the surrounding region will repel the evolving surface.

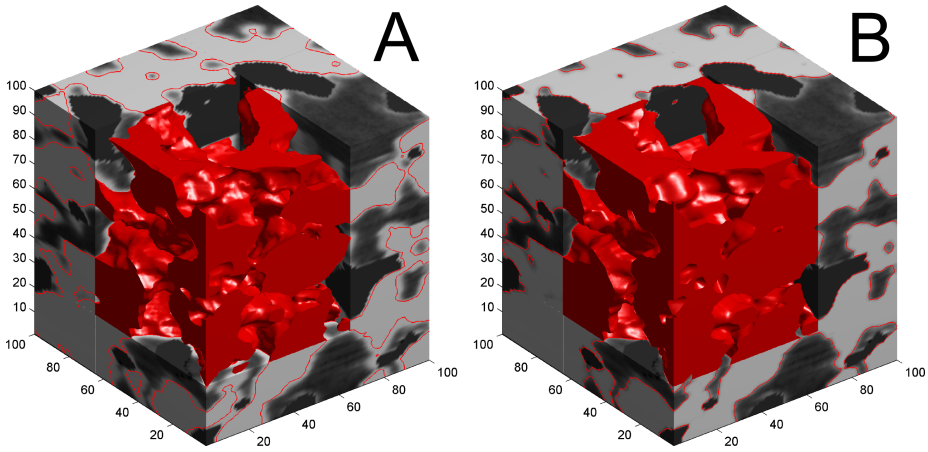

Figure 5.9: A 3D cube view of the segmentation on a subset of the LSC/CGO cathode data at the first and last iteration. The sides of the cube show the voxel intensities of the raw InLens image data. The red contour lines denote the intersection between the segmented surface and the boundary of the voxel cube. The red surface is a 3D rendering of the segmented CGO surface. The 2 cubes A and B correspond to iteration 1 and 13 (the final iteration). The cube has dimensions 100 × 100 × 100 voxels, corresponding to 5.86 × 5.86 × 4.45 μm.

These ambiguities can to some extent be resolved by a voxel growing technique [82] where the initial seed is grown iteratively to its neighbours if they are above or below a certain threshold. In essence, voxel growing is similar in functionality to the level set scheme but the level set scheme has 3 advantages. (1) The edge attraction field enables edge driven stopping of phase growth, where phase expansion can be stopped by an edge even though the intensity on the other side of the edge would dictate an expansion. (2) The curvature constraint will prevent the initial seed region from "spilling over" into thinly connected neighbouring regions of similar intensity. (3) The native sub-voxel precision of the level set method provides increased precision in the final surface.

5.4.1 Segmentation error types

In the level set approach, two main types of segmentation errors can be observed. The first type of error can be observed at sites 1 and 2 in figure 5.5A. These sites show errors in filling out the full extent of the pores. This error is caused by a combination of high intensity, sharp edges and the curvature constraint. The boundary is attracted to locations with a high intensity gradient and the curvature constraint keeps the surface from forming a thin spike into the dark area nearby. Note that the contour in figure 5.5 is part of the full 3D segmentation and thus forms coherent structures in the slices above and below as well. It is thus not possible to fully explain the location of the contours without looking at the neighbouring slices. Figure 5.6 gives an indication of this coherency.

The second type of error can be seen at site 3 in figure 5.5. This site shows the level set method including too much bright phase in the segmentation. This error is caused by a very slow sloping boundary between the pore and its surrounding phase. The slow slope causes a low contribution from the edge attraction velocity field and causes the upslope velocity field to dominate until the slope flattens out. Note that segmentation at site 4 is not erroneous, but rather the beginning of larger pores that can be seen by inspecting the neighbouring slices.

5.4.2 Parameter tuning

The level set method has more parameters than most segmentation schemes. As a consequence it does require experience to tune parameters correctly. To obtain quantitative results for the sensitivity of each parameter the level set segmentation was run numerous times with different parameter settings. The porosity segmentation example was chosen as the basis for the tests since this example is the most complex and has the largest amount of parameters. Since

no ground truth was available to compare against, the segmentation described in section 5.3.1 was used as the reference segmentation. The five parameters τ_a, b, c_{edge}, c_{slope} and τ_{seed} were changed one at a time by multiplying the value used in section 5.3.1 with a factor of 0.5 to 1.5. For each test the segmentation disparity was measured as the sum of false positives and false negatives in the segmentation divided by the total sum of porosity voxels in the reference segmentation. The results of the parameter sensitivity tests can be seen in figure 5.10.

To put the segmentation disparity values into perspective the value was measured for a systematic expansion of the reference segmentation. The result of expanding the reference segmentation by a certain distance in the direction normal to its surface can be calculated by choosing a different level set of the signed distance function that implicitly contains the surface. For instance, choosing the first level set instead of the zero level set corresponds to expanding the surface by 1 voxel in all directions. In this way, the segmentation disparity was measured to be 0.227 for an expansion of 1 voxel and 0.0241 for an expansion of 0.1 voxel.

Figure 5.10 shows that the τ_a parameter controlling the intensity value where the interface changes from expanding to contracting is the most sensitive to variation. The parameter has a disparity increase of 0.0587 and 0.0764 for a value variation of $\pm 10\%$. This parameter is expected to be the most sensitive since it has a similar effect as the threshold used in a threshold segmentation. The sensitivity is however lessened by the presence of the $\mathbf{V}_{edge}$ vector field since the $\mathbf{a}$ contribution needs to have high magnitude to overcome strong edges. This also explains the sudden large increase in disparity for τ_a past an increase of 1.3 of the reference value. At this value the $\mathbf{a}$ contribution becomes large enough to pass the majority of edges and the interface expands without hindrance. The value of $\mathbf{a}$ is often set by intensity histogram inspection using the same principles as one would use for deciding a threshold for a threshold segmentation.

The b parameter controls the curvature penalization. In figure 5.10 the b parameter can be seen to be the most robust to parameter changes. The parameter has a disparity increase of 0.0059 and 0.0060 for a value variation of $\pm 10\%$. The b parameter primarily has an influence on the texture of the interface and little influence on the overall location. The b parameter does in general not need to be changed when working with similar samples.

The c_{edge} parameter controlling the weighting of the $\mathbf{V}_{edge}$ vector field has a disparity increase of 0.0170 and 0.0136 for a value variation of $\pm 10\%$. The low disparity can be explained by the fact that the $\mathbf{V}_{edge}$ vector field directs the interface towards a specific intensity feature (an edge) and the $\mathbf{c}_{edge}$ parameter only changes the magnitude of the vector field not the direction. As was seen

above for the τ_a parameter it is important to set $\mathbf{c}_{edge}$ relative to τ_a. This is however only an issue if τ_a is set such that the $\mathbf{a}$ contribution would make the interface expand at the intensity level of the surrounding phases.

The c_{slope} parameter controlling the weighting of the $\mathbf{V}_{slope}$ vector field has a disparity increase of 0.0276 and 0.0303 for a value variation of $\pm 10\%$. For the porosity example the interface boundary at equilibrium will at some locations not coincide with an edge. These locations are characterized by having a small constant slope and having larger intensity. The $\mathbf{V}_{slope}$ contribution will be constant in these locations and the $\mathbf{a}$ contribution will be negative with increasing magnitude. The final equilibrium location will thus be determined by the parameters τ_a and c_{slope}. A change in c_{slope} in these locations will thus change the equilibrium location and the interface location.

The τ_{seed} parameter is the intensity threshold for the creation of the seed. The parameter can be seen in figure 5.10 to be robust to changes in value and has a disparity increase of 0.0170 and 0.0136 for a value variation of $\pm 10\%$. The exact size or shape of the seed is not decisive since a large group of different seeds eventually will evolve to the same final interface boundaries. Strictly speaking it is only necessary to have one correctly seeded voxel inside each phase grain. Having a larger seed will decrease the number of iterations needed to reach a converged surface, but the surface will likely be almost the same as if the seed had been only a single voxel.

5.4.3 The level set method as a framework for segmentation

The level set method can be used as an overall framework for segmentation. The formulation can be simple or more complex depending on the difficulty of the segmentation. If all the vector and scalar fields ($\mathbf{V}$, $\mathbf{a}$ and $\mathbf{b}$) are set to zero the seed is also the result, corresponding to traditional thresholding. Different combinations of $\mathbf{V}$, $\mathbf{a}$ and $\mathbf{b}$ can be included or set to zero depending on the segmentation task at hand. The level set method does not necessarily replace existing pixel classification schemes but rather extends upon them. In particular the normal expansion scalar field $\mathbf{a}$ can be constructed from many of the segmentation schemes currently used. A statistical method, such as the one described in [71], could be used to power the $\mathbf{a}$ field such that the surface would expand into voxels that were part of the statistical segmentation and be repelled by voxels that were not part of the statistical segmentation. This would preserve the general properties of the statistical segmentation scheme and add additional properties of the level set method such as surface smoothness and sub-voxel precision. A statistical pixel classification method ultimately assigns the

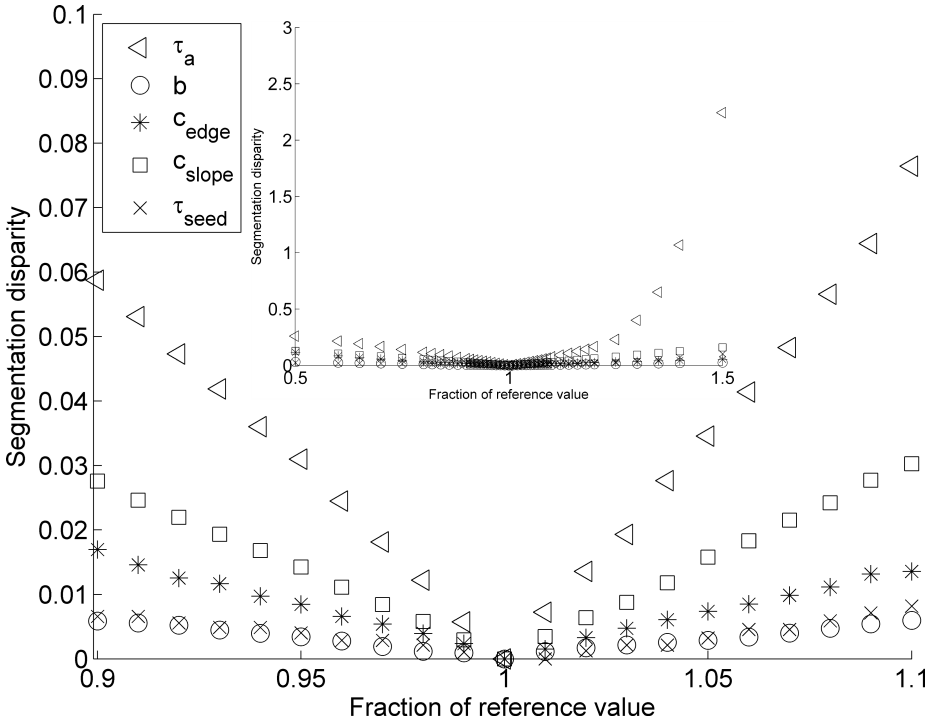

Figure 5.10: Parameter sensitivity analysis. The figure shows the amount of disparity between the reference segmentation and the segmentation done with a particular parameter variation. Disparity is measured as the sum of false positive and false negative voxels divided by the total number of voxels covered by the reference segmentation. The disparity is plotted against normalized parameter changes where the values correspond to fractions of the reference values. The inset shows the same data for a larger parameter value range. The reference values are:$\tau_a = 0.38$, $b = 0.1$, $c_{edge} = 0.04$, $c_{slope} = 0.45$ and $\tau_{seed} = 0.2917$.

class that has the largest probability of being correct to a given pixel. A natural combination of statistical methods and level set methods is to assign a value to **a** at each voxel that scales with the probability that the given voxel belongs to a given phase. The segmentation would then predominantly be controlled by **a** where there is a high probability of correctness. The vector and scalar fields **V** and **b** would dominate the evolution in regions with uncertainty. Examples 3.1 and 3.2 already use this scheme in a simplistic version. In the examples we make the assumption that the probability that a given voxel belongs to a specific phase is only dependent on the intensity of the voxel. Although intensity histograms are used to support the choice of the threshold where **a** changes sign, the scaling of **a** with intensity is ad hoc.

5.4.4 Interface precision

Level set segmentation is performed with sub-voxel accuracy. Because the segmented surface is represented as an iso-surface of the level set function it is not restricted to follow the discrete voxel locations of the voxel grid. The sub-voxel precision polygonal surface of the segmented region can be obtained by simple interpolation between neighbouring voxels with different signs [83]. Realistically smooth surfaces are particularly important when performing measurements of surface areas, where a voxelated structure will result in a general overestimation of measured surface areas.

The idea of smoothing the segmented surface to obtain realistically smooth surfaces with sub-voxel precision is not new [45]. The difference between the presented method and previous smoothing approaches is that in the previous work the smoothing has been performed as a post processing step to the already discretely segmented voxels. Post process smoothing will cause the smoothed surface to diverge from the segmentation and is likely to introduce artefacts independently of the underlying data. The level set method incorporates the smoothing as a regularization parameter during each iteration. The smoothing by curvature penalization is thus an integral part of the segmentation process itself. The sub-voxel precision surface of the level set method is thus directly supported by the underlying voxel grid intensities.

The level set segmentation consistently finds the phase boundary where the change in intensity is greatest due to the edge attracting velocity field term in the evolution equation. The location with the quickest change in intensity is often more consistent in defining the location of the physical phase boundary compared to defining the boundary purely based on intensity levels. The intensity level of all phases might change slightly from one region to another but the location where the change is most rapid will remain the same. The edge

attraction field thus increases the accuracy of the segmentation at the boundary compared to adjusting a threshold until the boundary visually appears optimal or performing the outlining of structures by hand. Manual segmentation will always be superior in marking up the rough outline of phases but when it comes to accurate outlining of boundaries automatic methods are more accurate and consistent [80].

5.4.5 Resource usage

The level set method needs considerable memory to run. Not only will the grid containing the signed distance values need to be in memory at all times but also the scalar and vector fields. However, with the recent increase in the memory size of desktop computers analysis of voxel grids of hundreds of millions of voxels is feasible and will increase as installed memory is increased. One option for decreasing memory demand is to avoid pre-computing the scalar and vector fields but calculating them from the intensity information as they are needed. Another option is to break up the volume into overlapping sub-volumes and run the method on each sub-volume in turn. In general, only the resolution of the microscope and the time available for serial sectioning limit the number of volume voxels that can be processed.

The level set method is slower than most traditional segmentation schemes but given its automatic nature it will not take up man hours during computations. The computations in the update rule (5.2) and the re-initialization scheme only depend on a neighbourhood of constant size. The computational cost of the algorithm thus scales linearly with the number of voxels in the volume. The update rule can be applied in any order only dependent on its neighbours, allowing parallel multi-core computation. The speed of the algorithm will depend on the implementation and the segmentation task. The single threaded implementation currently used by the authors uses approximately 2 seconds per iteration on a voxel grid of 6 million voxels on an Intel®Core^{TM}2 T7200 2GHz CPU.

5.4.6 Future work

An accurate segmentation of constituent phases as presented here is only one step towards precise 3D measurements. To exploit the sub-voxel precision of the method the further data analysis must be done in a way that preserves the precision. We are currently developing methods for extracting precise 3D measurements from multiphase data sets represented as implicit surfaces.

5.5 Conclusion

In this paper, we have described a framework for performing automatic segmentation of phase structures by level set methods. The framework is based on expanding iso-surfaces driven by vector fields derived from the intensity information in the raw image data. The proposed approach is suited for a multitude of problems where either an intensity difference or a gradient neighbourhood difference that distinguish the phases can be identified.

The method applies the segmentation iteratively in 3D independent of the orientation of the voxel grid. It naturally incorporates sub-voxel precision of segmented surface boundaries and allows curvature constraints to be added to enforce the smoothness of the segmented surface.

We show in particular how the method can be applied to the complex problem of porosity segmentation, where simple thresholding techniques are not feasible due to insufficient discrimination between porosity and the surrounding phases based only on voxel intensity.

Lastly, although the present work focuses on automatic segmentation of SEM images obtained by FIB tomography, the framework can equally be applied to other data volumes obtained by classic Micro CT and mechanical tomography such as serial grinding or microtome.

Acknowledgements

This work was supported financially by The Programme Commission on Sustainable Energy and Environment, The Danish Council for Strategic Research, via the Strategic Electrochemistry Research Center (SERC) (www.serc.dk), contract no. 2104-06-0011.

Appendix A

In the following the word *interface* will be used for the n-dimensional generalization of a surface.

Implicit functions: Let $\phi(\overrightarrow{x})$ be some function defined on $\overrightarrow{x} \in \mathbb{R}^n$. We call this the level set function. The isocontours of $\phi(\overrightarrow{x})$ have dimension $n-1$. We will choose the zero iso-contour for all further use. We can then define $\Omega^- = \{\overrightarrow{x} | \phi(\overrightarrow{x}) < 0\}$ as the inside, $\Omega^+ = \{\overrightarrow{x} | \phi(\overrightarrow{x}) > 0\}$ as the outside and $\partial\Omega = \{\overrightarrow{x} | \phi(\overrightarrow{x}) = 0\}$ as the interface itself.

The common choice of ϕ is the signed distance function. A *distance function* is defined as

$$d(\overrightarrow{x}) = min(|\overrightarrow{x} - \overrightarrow{x_I}|) \quad \text{for all} \quad \overrightarrow{x}_I \in \partial\Omega \tag{5.4}$$

A *signed distance function* is a function ϕ where $|\phi(\overrightarrow{x})| = d(\overrightarrow{x})$ for all $\overrightarrow{x}$. Thus, $\phi(\overrightarrow{x}) = d(\overrightarrow{x}) = 0$ for all $\overrightarrow{x} \in \partial\Omega$, $\phi(\overrightarrow{x}) = -d(\overrightarrow{x})$ for all $\overrightarrow{x} \in \Omega^-$ and $\phi(\overrightarrow{x}) = d(\overrightarrow{x})$ for all $\overrightarrow{x} \in \Omega^+$. That is, the level set function at any location is defined as the signed distance to the closest point on the interface. A signed distance function has the property $|\nabla\phi(\overrightarrow{x})| = 1$.

The level set method: The evolution of ϕ is done by solving a simple PDE. The basic form also known as the level set equation can be written as

$$\phi_t + \overrightarrow{V} \cdot \nabla\phi = 0, \tag{5.5}$$

where ϕ_t is the temporal derivative of ϕ in the time variable t, $\overrightarrow{V}$ is an externally generated vector field and $\nabla\phi$ is the spatial gradient of ϕ, such that

$$\overrightarrow{V} \cdot \nabla\phi = \mathrm{v}_1\phi_\mathrm{x} + \mathrm{v}_2\phi_\mathrm{y} + \mathrm{v}_3\phi_\mathrm{z} \tag{5.6}$$

in the 3-dimensional case. The basic level set equation can be extended to allow more control of the evolution by adding two more terms. If ϕ is a signed distance function the new evolution equation can be written as equation (5.1), where κ is the mean curvature defined as $\kappa = \Delta\phi = \phi_{xx} + \phi_{yy} + \phi_{zz}$. The **a** term of equation (5.1) describes evolution of the interface in the direction normal to the interface ($\overrightarrow{N}$). To see this, note that

$$a\overrightarrow{N} \cdot \nabla\phi = a\frac{\nabla\phi}{|\nabla\phi|} \cdot \nabla\phi = a\frac{|\nabla\phi|^2}{|\nabla\phi|} = a|\nabla\phi|. \tag{5.7}$$

Signed distance functions have the property $|\nabla\phi| = 1$, and the term reduces to a.

Cartesian grid implementation: Let $\phi^n = \phi(t^n)$ denote the level set function at time t^n, the time at the n^{th} time step. By utilizing a simple first order accurate method for the time discretization of equation (5.1) we obtain

$$\frac{\phi^{n+1} - \phi^n}{\Delta t} + \overrightarrow{V} \cdot \nabla\phi^n + a = b\kappa^n \tag{5.8}$$

where ϕ^{n+1} is the level set function at time $t^{n+1} = t^n + \Delta t$, Δt is the time step. The superscript n denotes evaluation at time t^n. By reordering equation (5.8) we obtain the update rule in equation (5.2).

Note that great care must be taken when discretizing partial differential equations. The finite difference scheme used must be chosen in accordance with *upwind differencing* when discretizing the remaining partial derivatives in equation (5.2). Reference [78] provides a more thorough coverage.

In order to obtain the relative simple expression for the complete level set equation (5.1) we have assumed that ϕ is a signed distance function. There is however no guarantee that ϕ will remain a signed distance function after applying the update rule. This means that it is necessary at regular intervals to reinitialize ϕ to a signed distance function, without changing the interface itself. There are several approaches to this. The authors' implementation uses the method described in [84] at each iteration, but there are many successful alternatives.

Voxel grids can have hundreds of millions of voxels. To apply the update rule at each grid position each iteration and subsequently also reinitializing the value to be a signed distance function can be computationally demanding. We are however not interested in the entire domain but only the interface itself. An efficient method is to only update voxel grid values in a narrow band around the interface. The narrow band needs only be as thick as to contain enough valid voxels to compute finite differences and to ensure that the interface itself does not move outside the band in a single iteration.

The narrow banding and re-initialization to a signed distance function are incorporated in the level set algorithm as follows: The update rule is applied for all voxels in the narrow band around the interface only. The re-initialization method is applied in the narrow band to reinitialize the level set function to a signed distance function. The two steps of updating followed by re-initialization

to a signed distance function are repeated until there is no change between iterations.

Appendix B

The samples corresponding to figure 5.5 and figure 5.8 were prepared by mounting in Cold Mount^{TM}resin supplied by Structure Probe Inc. at atmospheric pressure, thus the internal pores were not filled with any substance compared to mounting under vacuum. The samples were subsequently mechanically ground and diamond polished to 1μm and then extracted from the thermoplastic resin by melting the Cold Mount prior to data collection. Data sets where acquired using a Zeiss XB 1540 CrossBeam^{TM}.

Two detectors where used for imaging. The InLens secondary electron detector is positioned close to the beam path and detects almost purely secondary electrons (SE_1 and SE_2) resulting in images containing little topographic information as they are sucked vertically up by the external electric of the final lens. This has the effect that the bottom of the pores often appears bright. Furthermore, due to the Gemeni column optics BSE electrons are not directed intentionally into the InLens detector due to their high energy and thus form a minimal fraction of the signal. The Everhardt Thornley SE2 detector is positioned on the wall of the microscope specimen chamber and detects a combination of secondary electrons and backscattered electrons (including secondary electrons induced by backscattered electrons scattering from the microscope itself (SE_3). Low energy SE_2 electrons are captured by the electric field generated by the bias voltage applied to the external detector cage, whereas BSE's that are captured by the SE2 detector have a direct line of sight to the detector. The SE2 detector features good topographic contrast and gives a low yield of detected electrons from the pores of the sample due to shadowing effects. In interpreting the signals from both detectors it is important to note that the acquisition surface is tilted 36° away from the detectors. Therefore, this reduces the fraction of BSE's with a line of sight to the InLens and SE2 detectors due to forward scattering of the primary electrons as well as shadowing the SE2 detector from high energy BSE's. The microscopes option for automatic tilt correction was used on to obtain isotropic pixel dimensions in the image plane.

For the Ni/YSZ sample a 120 μm aperture, high current mode and an accelerating voltage of 10 kV were used for SEM imaging. The FIB milling current was set at 200 pA. Images were acquired from the SE2 detector asynchronously with the FIB milling. Images where acquired during live milling. The voxel dimension in the data is 41.9 x 41.9 x 47.1 nm.

For the LSC/CGO sample a 120μm aperture, high current mode and an accelerating voltage of 2 kV were used for SEM imaging. The FIB milling current was set at 500 pA. Images were acquired simultaneously from the InLens and SE2 detectors synchronously with the FIB slicing. Images were thus acquired between milling operations. It should also be noted here that the BSE detection efficiency of the SE2 detector at 2 kV is very low. The voxel dimension in the data is 58.6 x 58.6 x 44.6 nm.

All the intensity data were acquired with intensities in the range [0,255]. The data were scaled to the range [0, 1] before any further processing was performed.

Chapter 6

High accuracy interface characterization of three phase material systems in three dimensions

P. S. Jørgensen[a], K. V. Hansen[a], R. Larsen[b], J. R. Bowen[a]

(a) Fuel Cells and Solid State Chemistry Division, Risø National Laboratory for Sustainable Energy, Technical University of Denmark, Frederiksborgvej 399, 4000 Roskilde, Denmark.

(b) Department of Informatics and Mathematical modelling, Technical University of Denmark, Richard Petersens Plads, Building 321, Denmark.

Manuscript submitted for publication 2009.

Abstract

Quantification of interface properties such as two phase boundary area and triple phase boundary length is important in the characterization of many material microstructures, in particular for solid oxide fuel cell electrodes. These microstructures can be obtained by tomography schemes such as focused ion beam serial sectioning or micro computed tomography. We present a high accuracy method of calculating two phase surface areas and triple phase length of triple phase systems from sub-voxel accuracy segmentations of constituent phases. The method performs a three phase polygonization of the interface boundaries which results in a non-manifold mesh of connected faces. We show how the triple phase boundaries can be extracted as connected curve loops without branches. The accuracy of the method is analyzed by calculations on geometrical primitives.

6.1 Introduction

Three dimensional reconstructions of microstructural data have become an important tool in analysing and gaining understanding of multiphase material systems on the microstructure level [46, 45, 49]. These material systems include solid oxide fuel cell (SOFC) electrodes where the distribution and interconnectivity of constituent phases are of critical importance for the electrochemical performance. The volume specific Triple Phase Boundary (TPB) length is one of the most critical performance attributes of SOFCs. The TPB length is the total length of the curves within the microstructure that bound the interface between each of the three constituent phases. The TPBs are important since they are the locations where the electrochemical reactions within the cell take place.

We present here a high accuracy method for quantifying the TPB length and interface area of triple phase systems from a prior segmentation of constituent phases. This paper can be seen as a direct continuation of our previous work [1] that describes a computational framework for sub-voxel accuracy segmentation of microstructural tomography data. The methodology described here enables the extraction of accurate measurements from already obtained segmentations.

TPB length calculations have been carried out before [73, 47]. Common to

these calculations are that TPB sites are identified by systematically searching through the voxel structure for locations that contain all 3 types of phases in a small neighbourhood. The result is a grid of TPB sites. The TPB length has then been quantified either by counting the number of TPB sites and weighting them by the voxel dimension or the points are connected and the total length of the line segments is summed. Simply connecting the TPB sites has been reported to cause problems of branching curves [73]. These approaches lead to an overestimation of the TPB curve length because the TPB line segments do not form a smooth curve. A correction factor can be used to account for the overestimation. However, determination of the correction factor is challenging since the factor will be a function of both the sample and the resolution of the voxel grid. Two-dimensional stereological approaches such as [85] essentially fall into the same category since the TPB estimate is based on a count of TPB sites and a correction factor derived from an assumption of microstructure shape.

Reference [51] employs a method where each phase is expanded slightly and the TPB curves are extracted as the centrelines of the overlapping tube-like volumes. This method will theoretically calculate the TPB length of the segmentation accurately for an infinitesimally small expansion. The segmentation will however need to accurately represent the physical sample interfaces. For a conventional voxel labelling the interfaces between phases can not be represented as smooth surfaces. This means that an infinitesimal expansion of a voxel accuracy segmentation will result in line segments that lie exactly on the sides of the voxels. This in turn leads to an overestimation of the TPB curve length as the TPB line segments do not form a smooth curve.

The method presented here uses a different approach to identifying TPB curves. The method reconstructs all the two phase interfaces in the structure as connected polygons. In this representation of the phase interfaces each polygon shares edges and vertices with adjacent polygons. After the data structure is constructed the TPB curves can be identified as the edges that are shared by 3 polygons of different phase. This is a conceptually simplified explanation; further details are given in section 6.2.5.

In this approach the data structure is in essence the result. The accuracy of the method thus ultimately depends on the accuracy of the interface polygonization. This approach results in TPB curves that all form connected curves without branches. The connectivity of the TPB line segments makes it possible to perform smoothing operations on the TPB line segments to lessen the overestimation problem.

The result of the framework presented in our previous work [1] is phase segmentations represented by the zero iso-level of a signed distance function. This representation of the segmentation provides increased accuracy compared to

voxel classification. This is because smooth surfaces can be represented since the interface is not limited to lie exactly halfway between two voxels. The proposed methodology is developed specifically to take advantage of the increased accuracy of sub-voxel segmentations. Conventional voxel accuracy segmentations can however be processed using the same methodology by converting the segmentation to an implicit surface representation.

6.2 Methodology

An iso-surface polygonization scheme converts a specific iso-level in a volumetric representation into a surface representation defined by a mesh of polygons. Good polygonization techniques such as the Marching Cubes algorithm [83] and its extensions [86] have been around for decades. These techniques are an integral part of many software visualization packages. However, most of these techniques are only designed for two phase structures. Polygonization techniques have been developed for the n-phase case [87]. The polygonization methodology presented below is similar to this technique and in many cases simpler. The methodologies do diverge in the problem they attempt to solve; mainly due to the need to collapse two representations of the same interface into one as is described in section 6.2.1. Since the polygonization scheme is based primarily on existing techniques the methodology described here will be on a largely conceptual level, hence the reader is referred to the general literature [88, 78]. Focus is placed instead upon the areas of the presented approach that diverge from existing literature.

All figures in this section are constructed as two dimensional examples for illustration purposes only. Therefore phase boundaries are thus represented as curves and TPBs as points. However, the algorithm itself works in three dimensions where phase boundaries are calculated as surfaces (polygons) and TPBs as curves (polygon edges).

6.2.1 Segmentation representation

Let us start by introducing the notion of a signed distance function in the context of a three dimensional surface. The signed distance function of an interface at a location (X, Y, Z) in the segmented volume is defined as the distance to the closest point on the interface, negative on the inside of the surface and positive on the outside of the surface.

The polygonization method described below takes as input two sub-voxel accuracy segmentations where the interfaces are described implicitly as the zero iso-level of a signed distance function. The concept is exemplified in figure 6.1 (A) and (B). The signed distance function is discretely evaluated on a regular grid of voxels. Such a segmentation can be created as described in [1]. For a three phase system it is required that a segmentation of two of the constituent phases are provided. Let the three phases of the system be denoted P_1, P_2 and P_3 respectively and let the functions that implicitly contain the interface of two of the phases be denoted ϕ_1 and ϕ_2. A function ϕ_3 is not required since the P_3 interface information can be derived from ϕ_1 and ϕ_2. Figure 6.1 shows the segmentation representation. The segmentation representation leads to a double representation of the interface between P_1 and P_2. A polygonal interface representation is sought where each polygon separates two specific phases and where only one polygon represents a specific interface section. A way to collapse the double interface representation to a single interface is thus needed. Since the segmentation of P_1 and P_2 can be performed independently it can not be assumed that the two representations of the same interface are congruent. The two segmented phases represented by the zero level set of $\phi_1$and ϕ_2 may overlap in some sections or have small gaps between them in other areas. This means that a definition of when the two interfaces are considered to coincide is needed.

In the locations where the zero level set of $\phi_1$and ϕ_2 overlap the optimal location of the collapsed surface is halfway in between the two interfaces. This surface can be expressed as the zero level set of the function $\phi_{12} = (\phi_1 - \phi_2)/2$. Recall that $\phi_1$and ϕ_2 are signed distance functions which means that the value at any location corresponds to the distance to the interface. The function ϕ_{12} will thus evaluate to zero at locations that have the same distance to the ϕ_1 and ϕ_2 zero level sets. The division by 2 ensures that ϕ_{12} is a good approximation to a signed distance function whenever the nearest ϕ_1 and ϕ_2 zero iso-levels are close to parallel.

We define coincidence of the interface between P_1 and P_2 as the locations $\mathbf{x}$ on the zero level set of ϕ_{12} for which the distance to both the P_1 and P_2 interface is smaller than a merge threshold τ_m ($\phi_1(\mathbf{x}) < \tau_m \wedge \phi_2(\mathbf{x}) < \tau_m \wedge \phi_{12}(\mathbf{x}) = 0$). At these locations the coinciding interfaces will be collapsed to the nearest point on the ϕ_{12} zero level set. The following sections will describe step by step how the interfaces are polygonized and stitched together at the triple phase boundaries and how the zero level set locations are obtained when $\mathbf{x}$ no longer represents a continuous location but a discrete voxel location.

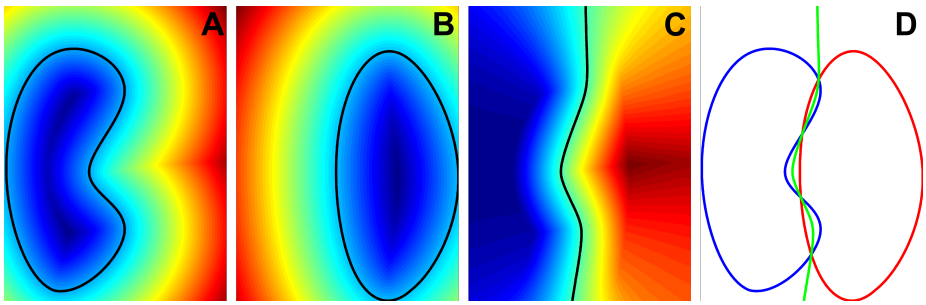

Figure 6.1: An example of an implicit three-phase segmentation representation. (A-C) The functions ϕ_1, ϕ_2 and ϕ_{12} evaluated on the same domain. Blue shades denote negative values, red shades denote positive values and the black curves show the zero level set of the functions. (D) The zero level sets of the functions ϕ_1, ϕ_2 and ϕ_{12}.

6.2.2 Phase map

A phase map is created from the two implicit functions ϕ_1 and ϕ_2. This phase map corresponds to a conventional voxel accuracy segmentation where each voxel location $\mathbf{x}$ has a label $L(\mathbf{x})$ describing what phase it belongs to**.** This phase map defines the topology of the interface structures. The phase map is constructed at each location $\mathbf{x}$ based on a set of basic rules:

1. $\phi_1(\mathbf{x}) \leq 0 \ \wedge \ \phi_2(\mathbf{x}) > 0$. The location is only inside the ϕ_1 zero level set. Label as phase 1.
2. $\phi_1(\mathbf{x}) > 0 \ \wedge \ \phi_2(\mathbf{x}) \leq 0$. The location is only inside the ϕ_2 zero level set. Label as phase 2.
3. $\phi_1(\mathbf{x}) \leq 0 \ \wedge \ \phi_2(\mathbf{x}) \leq 0$. The location is inside both the ϕ_1 and ϕ_2 zero level sets (overlapping interfaces). Label as the phase with the lowest $\phi(\mathbf{x})$ value.
4. $\phi_1(\mathbf{x}) > 0 \ \wedge \ \phi_2(\mathbf{x}) > 0$. The location is outside both the ϕ_1 and ϕ_2 zero level sets. If $\phi_1(\mathbf{x}) > \tau_m \ \wedge \ \phi_2(\mathbf{x}) > \tau_m$ label as phase 3 otherwise label as the phase with the lowest $\phi(\mathbf{x})$ value.

Rules 1-3 are straight forward and intuitive however the 4th might seem counter intuitive since these locations are clearly outside both the ϕ_1 and ϕ_2 zero level set. The threshold τ_m is used to prevent problems in locations where the ϕ_1 and ϕ_2 zero level sets almost coincide but still have a small gap between them. At

these locations a thin layer of phase 3 voxels can form between P_1 and P_2. Note that this threshold typically is set to a very low value depending on the quality of the segmentation, typically between 0 and 1 voxels. A final optional step is to collapse isolated phase regions in the phase map with an unrealistic small volume as these are more likely to be segmentation errors than actual particles. The phase map creation is illustrated in figure 6.2 (A).

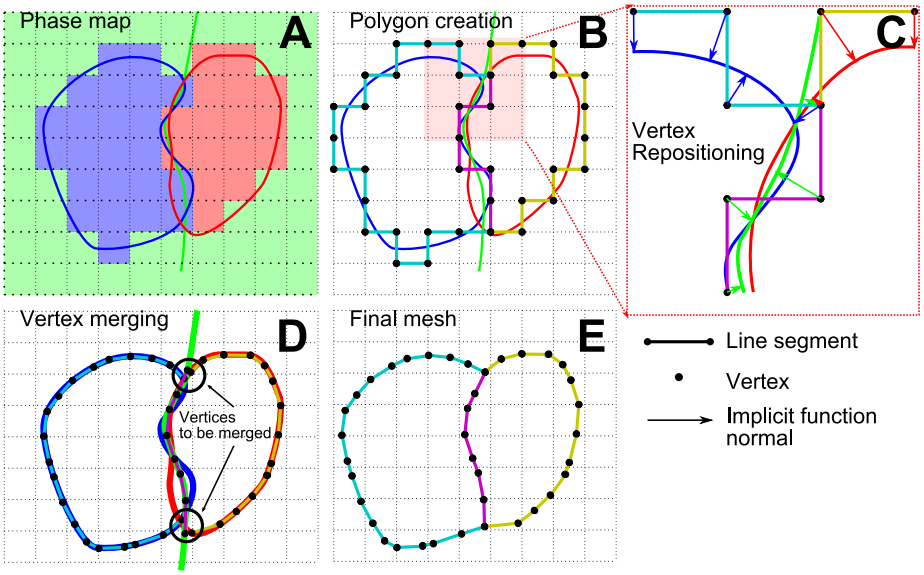

Figure 6.2: 2D illustration of the mesh creation and vertex repositioning.

6.2.3 Voxel accuracy polygonization

The phase map is processed through a very simple initial polygonization scheme. The scheme runs through all voxels in the phase map. Three of the voxels six non-diagonal neighbours are checked and if the neighbour has a different class than the current voxel a quadrilateral is created on the side of the voxel. Only three of the six sides are tested in order to avoid duplicated polygons of the same interface. Quadrilaterals consist of 4 edges that connect pairs of vertices. The concept is illustrated in figure 6.2 (B). At this initial stage each polygon is unaware of its position in the mesh relative to other polygons and no edges or vertices are shared between polygons. To be able to resolve the connectivity of the polygons a reference to the polygon is inserted into a connectivity matrix. A connectivity matrix exists for each voxel corner in the voxel grid that is used by at least one polygon. The connectivity matrix makes it possible to merge

vertices that coincide according to the topology of the phase map.

6.2.4 Sub-voxel accuracy polygonization

The voxel accuracy polygon structure is ill suited for surface or TPB calculations due to its cubic nature as it overestimates area and length calculations. If no sub-voxel accuracy information was available this would typically be remedied by smoothing the polygon mesh after the connectivity of the polygons had been resolved. Smooth sub-voxel accuracy surface information is however available from the three functions ϕ_1, ϕ_2 and ϕ_{12}. The implicit surfaces described by those functions are supported directly by the intensity information in the raw intensity data [1] and are thus preferred over post segmentation smoothing.

To obtain polygon vertex coordinates with sub-voxel accuracy a vertex movement scheme is applied to all vertices. The positioning scheme takes advantage of two important properties of signed distance functions:

- The absolute value of the signed distance function evaluated at a certain location corresponds to the distance to the closest point on the zero level set.
- The gradient of a signed distance function evaluated at a certain location corresponds to the direction in which the closest point on the zero level set is found.

These two properties enable us to design a vertex movement scheme that moves each vertex onto the interface described by the zero level set of the signed distance function.

For each vertex position $\mathbf{x_v}$ we interpolate the value $\phi(\mathbf{x_v})$ and the gradient $\nabla\phi(\mathbf{x_v})$ of the appropriate signed distance function from the values at the surrounding voxels. The new vertex position is then found at $\mathbf{x_{new}} = \mathbf{x_v} - \phi(\mathbf{x_v})\nabla\phi(\mathbf{x_v})$. Given a continuous signed distance function this movement scheme will move the vertex exactly on to the interface represented by the zero level set of the signed distance function. The appropriate function to evaluate is based purely on which set of two phases the polygon that the vertex is part of separates.

- For an interface between P_1 and P_3, ϕ_1 is used.
- For an interface between P_2 and P_3, ϕ_2 is used.

- For an interface between P_1 and P_2, ϕ_{12} is used.

The sub-voxel accuracy repositioning scheme is illustrated in figure 6.2 (C).

At this point the set of polygons do not constitute surfaces but rather a collection of individual polygons. To form a data structure from which calculations can be extracted the vertices from adjacent polygons must be merged and neighbour connection information must be stored for each polygon. This connectivity information is constructed by processing each connectivity matrix in turn. The connectivity matrix contains a reference to all polygons that according to the topology of the phase map share the same vertex. This vertex is denoted the centre vertex. The index of the polygon in the connectivity matrix corresponds to the polygons relative position to the centre vertex. Using these matrix indices it is possible to efficiently connect all pairs of polygons that share two vertices (an edge). Compared to polygons that separate two phases only, the polygons located at the TPB use three different functions (ϕ_1, ϕ_2 and ϕ_{12}) for vertex repositioning. It is therefore not expected that these vertices are repositioned to the same position. The position of the centre vertex is thus calculated as the average position of the vertices that are to be merged. The vertex merging is illustrated in figure 6.2 (D) and the final mesh can be seen in figure 6.2 (E).

At the initial polygon creation step a reference to the two voxels that the polygon separates is stored together with the polygon. A powerful dual representation of the phase structure is thus available. A phase map representation that represents the interior of each phase and a polygonal mesh that accurately describes the interface between the phases. The references between polygons and the two voxels they separate facilitate a simple transition between the two representations.

6.2.5 TPB identification

With the interface topology/morphology of the phases stored in a data structure as described above it is straight forward to locate the edges in the mesh that correspond to a triple phase curve segment. The TPB edges are identified as the outline of the interface surface between phase 1 and phase 2. In other words, edges that are part of only one polygon that separates a phase 1-2 and both a 1-3 and a 2-3 polygon simultaneously. This definition intuitively provides us with the TPB edges since a TPB must exist where a two phase boundary meets the third phase. This definition also guarantees that the TPB curves will form non-branching curves that either form a loop inside the voxel cube or that intersect the side of the voxel cube.

In practice this scheme is implemented by running through all polygons checking the criteria described above. The qualifying edges are stored in a separate list and linked to their neighbours by identifying when two edges use the same vertex. To obtain accurate TPB measurements it is important that the extracted TPB curves are smooth. However, this will often not be the case. The connectivity information of the edges can however be used to smooth the edges.

A simple smoothing approach has been applied here where each vertex of the edge curve is repositioned by calculating a weighted average of the position of itself and its immediate neighbours. A weight of 0.5 is used for the vertex undergoing repositioning and a weight of 0.25 is used for each of its neighbours. This smoothing approach can be repeated, each time resulting in a smoother curve of connected edges. However, a higher number of smoothing iterations is not strictly an improvement in accuracy. An increase in smoothing iterations corresponds to an increase in the number of neighbour vertices with non-zero weights. This will result in a shrinkage effect of the TPB curve loops. It is thus necessary to weigh the benefit of the smoothing against the shrinkage effect. A quantitative analysis of this dilemma is available in section 6.2.5. Note that the smoothing operations reposition vertices that are still part of the interface mesh. The mesh structure will thus contain the smoothed TPB edges which remain congruent with the surface mesh for visualization purposes.

The finished mesh structure allows very simple calculations of both two phase boundaries and TPBs. Area calculations are performed by simply summing the area of polygons that separate specific phase pairs. In this way the two phase boundary areas between two specific phases can be calculated along with the total interface area of each phase. The TPB length can be calculated by summing the length of the edges that make up the TPB curves. The link between the mesh representation and voxel representation can be used to further characterize the TPB locations. As an example the percolation [47] of each TPB edge can be identified by performing the percolation identification on the voxel grid and then utilizing the reference between the polygons and the voxel grid to check for percolation of the TPB edges.

6.3 Results

One approach to assessing the accuracy of proposed TPB measurement methods is to compare measured values with physical measurements that are believed to be governed mainly by the TPB length [46]. It is however difficult to determine the cause of the disparities by comparing a TPB length measure based on microstructure calculations to a measure obtained by a model of a directly

measurable physical property. The disparity could be due to several factors.

- Errors or insufficient resolution in the segmentation.
- Errors in the computation method used to calculate the TPB length from the segmentation.
- The model that connects the physical measurement to the TPB length could be too simple.

The accuracy analysis presented here takes a different approach. To asses the accuracy of the TPB length and surface area measurement method alone, tests are constructed that seek to eliminate segmentation errors and model errors. This is achieved by applying the method to spheres of varying diameter in a setup similar to [51]. The advantage of using spheres for the analysis is three fold.

1. Segmentation error can be eliminated since a sphere can be constructed as an implicit surface of a signed distance function without errors. This is achieved by calculating the exact Euclidian distance to the centre of the sphere at each voxel location. To obtain a sphere with a specific radius the radius is simply subtracted from all voxels. This representation is error free in the sense that it is the most accurate signed distance function representation that can be achieved at a given resolution.
2. The surface area of a sphere is defined as $A_s = 4\pi r$, where r is the radius of the sphere. Since this is an exact formula and not a model the calculations of surface area can be directly compared. For the TPB length calculations a similar exact formula can be obtained for the intersection of two spheres of radius r, at the distance r from eachother. The length of the intersection curve between the two spheres, and thus the TPB length, is defined as $L_{int} = 2\pi\sqrt{r^2 - \left(\frac{r}{2}\right)^2}$.
3. The segmentation sampling resolution error can be assessed by varying the resolution of the sphere representation on the voxel grid. Let the radius of the sphere be measured in voxel widths. Increasing the radius will then result in a subsequent increase in sampling resolution. Recall that the basis of the surface area and TPB length calculations is a voxel grid with finite resolution.

To show the accuracy of the proposed method on voxel accuracy segmentation the following changes were made to the implicit surfaces used as input for the

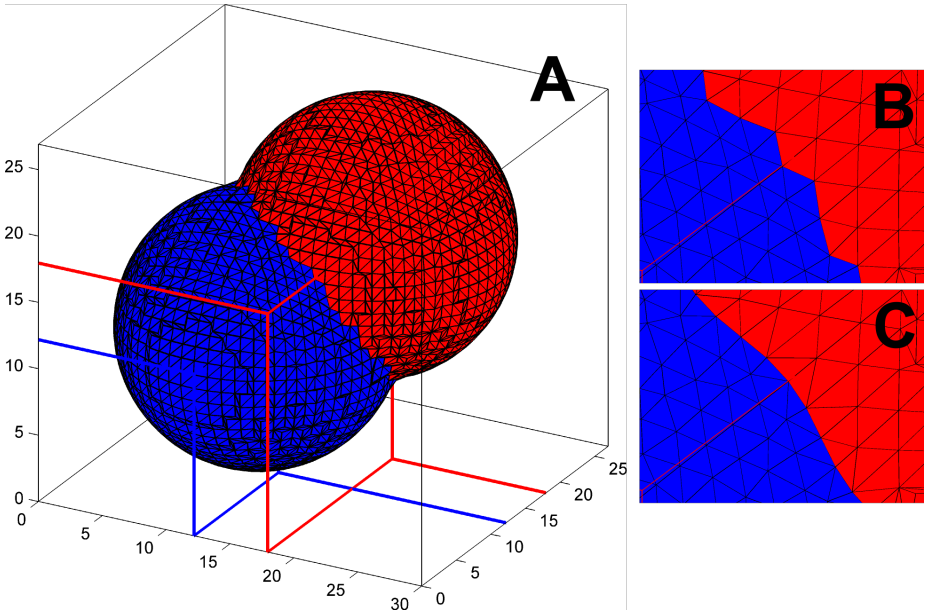

Figure 6.3: The intersecting sphere test setup. (A) An example of the mesh reconstruction without smoothing. (B) Enlargement of a TPB section of A. (C) The same enlargement with two smoothing iterations applied.

method. The signed distance map was first constructed as outlined above for the implicit surface representation of the spheres. The voxel grid was then discretized into 3 classes based on the voxel value of both implicit functions using the same methodology as described in section 6.2.2. This segmentation represents the best achievable voxel accuracy segmentation of the sphere. The voxel accuracy segmentation was then converted back to a set of signed distance function implicit surfaces but loosing the advantage of the sub-voxel accuracy in the process. Both a sub-voxel accuracy and voxel accuracy setup were used in the tests. All tests were performed with radii ranging from 1 to 30 voxels with increments of 0.25 voxel.

Two different setups where used for the sphere centres, an aligned setup and a skewed setup:

- The aligned setup sets the centre of the spheres in the centre of a voxel. This setup creates a setup with many symmetric axes that coincide with the three major directions of the voxel grid. The TPB length calculations place the second sphere the distance r away in the direction of the vector $V_{align} = [1\ 0\ 0]^T$. This setup represents a best case scenario where the most accurate reconstruction result can be expected.

- The skewed setup represents a more challenging case where the centres of the spheres are not in the centre of voxels. The offset from the centre of the voxel was chosen to be 0.1817. This offset was selected to not be a simple fraction between 0 and 1 but the value could have been any other number. The TPB length calculations place the second sphere the distance r away from the first in the direction of the vector $V_{skew} = [1\ 1\ 1]^T$. This setup represents a worst case scenario where the least accurate results can be expected. The setup can be seen in figure 6.3.

It is reasonable to assume that the accuracy of a real microstructure with a similar radius of curvature of its interface boundaries would be somewhere between the accuracy of these worst case and best case setups. In total four setups were tested: Aligned spheres and skewed spheres both with and without sub-voxel accuracy. The sphere results are reported as relative errors calculated as $err = (m_{exact} - m_{measured})/m_{exact}$, where m_{exact} is the exact analytical quantity.

6.3.1 Surface area

The surface area of the sphere in each of the 4 setups was calculated as described in section 6.2.5. For the surface area calculations only one sphere was used in each of the setups. The results of the surface area calculations can be seen in figure 6.4. The general trend for all four setups is the same. The error is large for spheres with small radii and has an asymptotic convergence as the radii increases. To increase the readability of the figures they only show the relative error for radii larger than 3.

As can be seen the results for the aligned and skewed setup are almost identical in the case of sub-voxel accuracy. This is because the signed distance function evaluated on the voxel grid can represent the skewed case just as well as the aligned case. The voxel accuracy tests perform worse than the sub-voxel accuracy tests for large radii. This is because the discrete voxel classification can not accurately represent the sphere; the surface of the polygonized sphere will thus have artefacts caused by the voxelated surface structure of the voxel accuracy representation. The noise in the curves is caused by some radii resulting in surfaces that are more aligned with the voxel grid than others.

By observing figure 6.2 the consequence of two kinds of errors can be observed. The two types of error are illustrated in figure 6.5.

1. The first error is sampling frequency. The sphere is represented by a finite number of polygons and the surface area is calculated as the combined area of these. The total measured area will thus always be smaller than the theoretical value assuming that the polygon vertices are placed exactly on the surface of the sphere. This error is clearly visible in figure 6.4 where all the tests have a large under estimation for small radii and a decreasing under estimation for larger radii. The limit of the sampling frequency error is zero as the radius approaches infinity.

2. The second error is sampling accuracy. This is an error in the location of the vertices of the polygons. An extreme case would be to use the sides of the voxel segmentation as polygons directly. This scheme would cause a severe over estimation of the surface area. This error can be made very low but it will always be present to some degree and the error will always be an overestimation assuming that the sampling accuracy error has no preferred direction.

Any measurement will thus always be the result of two errors that influence the measurement in different directions. The surface area measurements are heavily

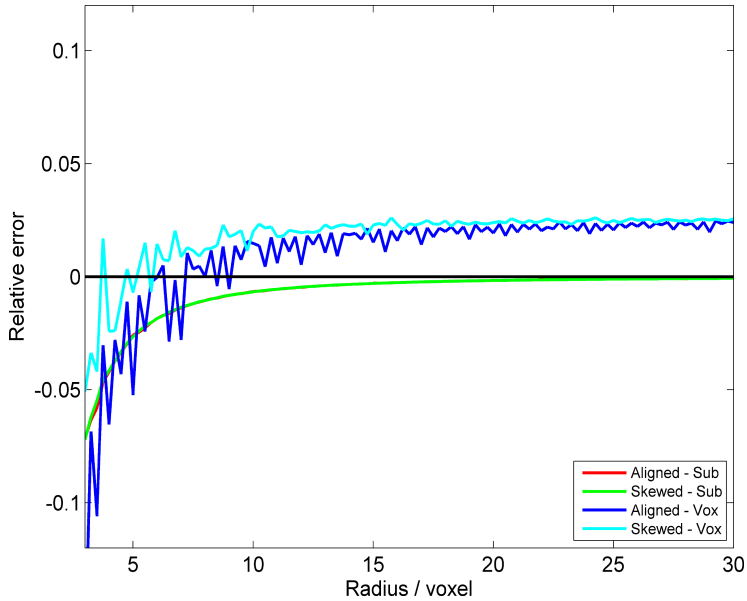

Figure 6.4: Surface area results. The relative error of the surface area calculations as a function of sphere radius for the 4 different setups. Note that results of the aligned and skewed sub-voxel precision setups are very similar and thus indistinguishable from each other in the figure.

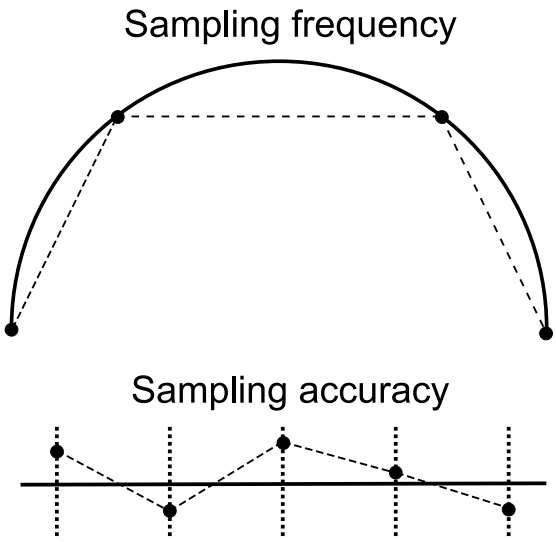

Figure 6.5: The two types of errors in the reconstruction. Sampling frequency error is the error introduced by approximating smooth curves with line segments. Sampling accuracy error is the error introduced by inaccurate vertex positioning.

influenced by sampling frequency error for small radii thus giving a large under estimation. The sub-voxel accuracy measurements generally have a very small sampling accuracy error. This can be seen by the fact that the underestimation decreases towards zero as the radii becomes larger. If a significant sampling accuracy error was present we would expect the estimation error to cross zero at some point. The voxel accuracy tests in figure 6.4 both cross zero, indicating that they are affected by sampling accuracy error. This also explains why the error does not tend to zero for large radii.

The sub-voxel test becomes very accurate for large radii with a relative error of -0.17% for a radius of 20, where the aligned and skewed voxel accuracy tests have a relative error of approx. 2.0% and 2.5% respectively for the same radius.

The radius of curvature of interfaces in a real sample will be a distribution of values rather than a single value. These interfaces will have a random alignment with respect to the three major axes of the voxel grid. The relative error of a random interface orientation is expected to be bounded by the aligned and skewed setups since they correspond to a worst case and best case scenario. For this purpose sub-voxel and voxel precision results should be treated separately. In general the plots of relative error versus sphere radius can be used in two ways:

1. To asses the accuracy and precision of a calculation on a physical microstructure. The worst case accuracy can be assessed as the largest absolute relative error of the aligned and skewed setup results. Only the values in the expected radius of curvature range of the sample should be used. The precision can be assessed based on the difference between the signed highest and lowest value in the same range of the aligned and skewed setup. A high precision is characterized by relative errors that are close to constant (a smooth flat relative error curve in figure 6.4) and a small difference between the relative errors of the aligned and skewed setups. In such a case there is a systematic under or over estimation that can be systematically corrected in the interface calculations.

2. To choose image resolution and slice frequency before data acquisition. Based on the assessed curvature of radius of the sample the expected accuracy of the results can be assessed before data requisition begins. As an example supposed an estimate of the surface area of interfaces in a sample is wanted with maximum +-1.5% error. From inspection of two dimensional micrographs the typical radius of curvature of the grains is assessed to range between 200 nm and 450 nm. From figure 6.4 it is observed that a radius of curvature of between 4 and 9 voxels has relative errors between -4% and -1% for a sub-voxel precision segmentation. Estimates in this range can thus be systematically corrected for the 2.5% underestimation we would expect when centring the range on zero. The maximum error with this correction in place would be $\pm 1.5\%$. The maximum voxel size that would ensure this accuracy is 50 nm (200 nm/4 voxels). Note that this maximum error is only a measure of the error introduced by the sampling resolution and the calculation method. Errors introduced by inaccurate segmentation need to be assessed separately.

6.3.2 TPB length

Figure 6.6 shows the TPB length calculations with zero and two smoothing iterations for the four different setups. Note that for the TPB figures, the relative errors are plotted versus the radius of the TPB circle at the intersection rather than the radius of the spheres. This is done to be able to asses the accuracy versus the radius of curvature of the TPB curve loop.

Figure 6.6 (A) shows the accuracy for the four setups with no smoothing. For TPB curve radii larger than 5 voxels an overestimation of approx. 5-25% of the TPB length is observed for the four setups. Note that there is a large difference between the results for the aligned and skewed setups and each curve has significant noise. This indicates that the calculated result is very dependent

on the orientation of the surface relative to the voxel grid. As a result, calculations performed without any smoothing iterations would have a large confidence interval.

The overestimation is more severe for the skewed setup and for voxel precision surface representations. This is as expected since the reconstructed TPB curve loop of the skewed setup by construction is not parallel with the major axes of the voxel grid tessellation, thus resulting in a rough TPB curve loop. A visualization of the rough reconstruction of the TPB curve can be seen in figure 6.3. Clear artefacts from the original voxel structure can be seen in the curve. Even though these artefacts are much smaller than if the vertex positioning scheme had not been applied they still cause a substantial lengthening of the TPB curve.

If the shape of the phase structures where known it would be possible to obtain a much better initial TPB curve reconstruction. The algorithm described here however works on any phase structure. The algorithm thus reconstructs the spheres as it would any other structure.

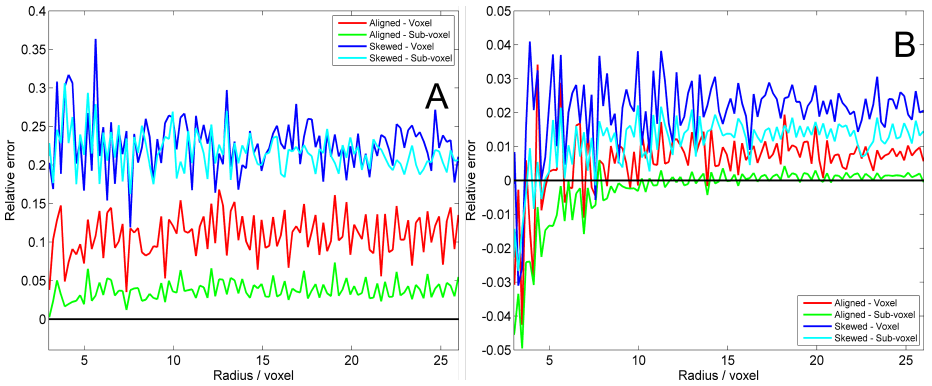

Figure 6.6: TPB length accuracy for the 4 setups for no smoothing and two smoothing iterations. (A) The relative error of the TPB length calculations without smoothing. (B) The relative error of the TPB length calculations with 2 smoothing iterations.

Figure 6.6 (B) shows the accuracy for the four setups with two smoothing iterations applied. The two smoothing iterations increase the accuracy across the four setups compared to not applying smoothing. For radii larger than 5 voxels a relative error of between -1.1% and 3.8% is observed for the voxel accuracy test and a relative error of between -1.6% and 2.2% is observed for the sub-voxel accuracy tests. The figures can be used to give a rough confidence interval of the TPB estimates analogous to the discussion in section 6.3.1. Similarly the

figures can be used to select the required voxel size for a given TPB estimation accuracy. The typical radius of curvature range of the TPB curves is not easily extracted from two-dimensional slices and a three-dimensional reconstruction of a similar sample is thus likely required to select the required voxel size of future samples.

The relative errors of the four setups appear to reach an asymptotic value with added noise for radii larger than 5 in figure 6.6 (A) and for radii larger than 10 in figure 6.6 (B). These radii indicate the minimum radius of curvature of the TPB curve that is sufficient to obtain the maximum accuracy of the reconstruction algorithm using the given smoothing scheme.

Figure 6.6 (A) and (B) show how applying a few iterations of smoothing dramatically reduces the relative error of the reconstruction for radii larger than 5 voxels. The number of smoothing iterations is a parameter of the TPB curve reconstruction and can in theory be set at arbitrarily large values. Increasing the number of smoothing iterations can be seen as including a larger neighbourhood of vertices in the weighted sum calculation of the new vertex position (section 6.2.5). Since the TPB curves form loops this means that for the number of iterations approaching infinity the TPB curves will collapse to a single point that is the centre of gravity of the loop. In practice the amount of smoothing iterations will have to be chosen based on the radius of curvature of the TPB curves. Figure 6.7 shows the skewed sub-voxel accuracy setup for smoothing iterations between zero and ten. It is seen that a high number of smoothing operations result in slower convergence but increases the asymptotic accuracy and precision. For 10 smoothing iterations the calculated TPB length estimate is within $\pm 0.4\%$ of the theoretical value but a radius of curvature of 15 or higher is required to obtain this accuracy. For smaller radii of curvature 10 smoothing iterations can result in a severe under estimation of the TPB length. This shows the significance of choosing a smoothing scheme that is appropriate for the chosen voxel dimension and sample structure size. Figure 6.7 also gives a hint of the accuracy that can be achieved as tomography resolution and quality improves.

6.3.3 Interface characterization of a SOFC electrode

The proposed characterization calculations were applied to a segmentation of a SOFC sample. The sample is a composite SOFC cathode consisting of 50/50 weight-percent strontium-substituted lanthanum cobaltite (LSC) and gadolinia-substituted ceria (CGO), similar to [81]. The calculations were performed on a grid of 180 x 180 x 190 voxels with a voxel dimension of 58.6 x 58.6 x 44.5 nm. The sub-voxel segmentation was carried out as described in [1].The collapse

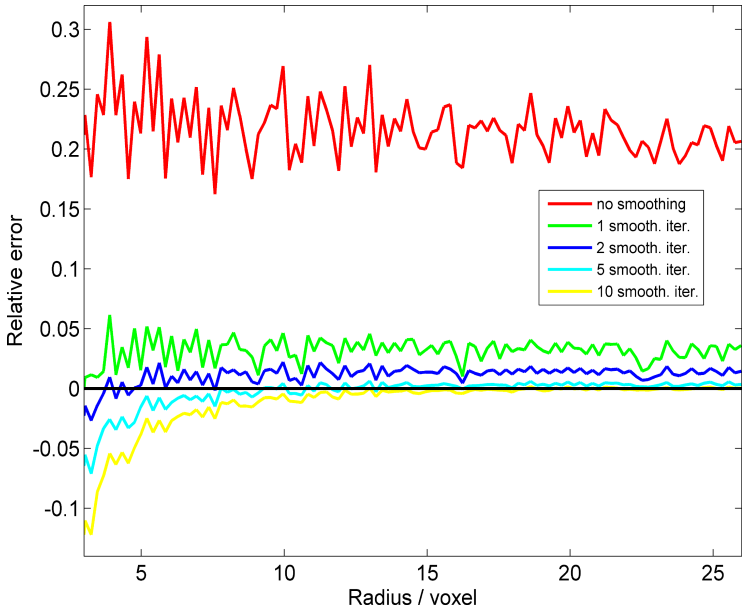

Figure 6.7: The effect of smoothing. The relative error of the TPB length calculations for varying amounts of smoothing iterations on the skewed sub-voxel setup.

L_v (TPB)	1.808 μm^{-2}
L_v (Percolating TPB)	1.413 μm^{-2}
S_v (Pore surface area)	1.074 μm^{-1}
S_v (CGO surface area)	1.896 μm^{-1}
S_v (LSC surface area)	1.822 μm^{-1}
S_v (Pore/CGO surface area)	0.575 μm^{-1}
S_v (Pore/LSC surface area)	0.500 μm^{-1}
S_v (CGO/LSC surface area)	1.322 μm^{-1}

Table 6.1: Interface measurements on a LSC/CGO cathode. All measurements are volume specific.

range was set to 0.5 voxels and connected material grains smaller than 10 voxels were collapsed before the polygonization. These two parameters were described in section 6.2.2. Two smoothing iterations were applied to the TPB curves. Figure 6.8 (A) shows a visualization of the pore/LSC two phase boundaries and figure 6.8 (B) shows the extracted TPB curves. Table 6.1 shows the calculated interface measurements for the sample. The percolating TPB length was calculated by only summing TPB line segments that have a percolating path in each phase to their respective source side: the LSC phase to the side facing the electrolyte and the CGO and pore phases to the opposite side.

6.3.4 Future work

The algorithm as described above has a direct correspondence between the resolution of the voxel grid and the resolution of the constructed polygon mesh. Many algorithms exist for polygonal subdivision [89] such as Catmull-Clark subdivision [90]. The resulting increase in polygons could increase the accuracy of both the area and TPB estimates by reducing the influence of sampling frequency error without increasing voxel resolution.

More advanced smoothing techniques could be employed for the TPB vertex smoothing, for instance by fitting splines to the TPB vertex positions. The use of splines to represent the TPB curves would reduce the error introduced by sampling frequency since the spline representation can be evaluated continuously.

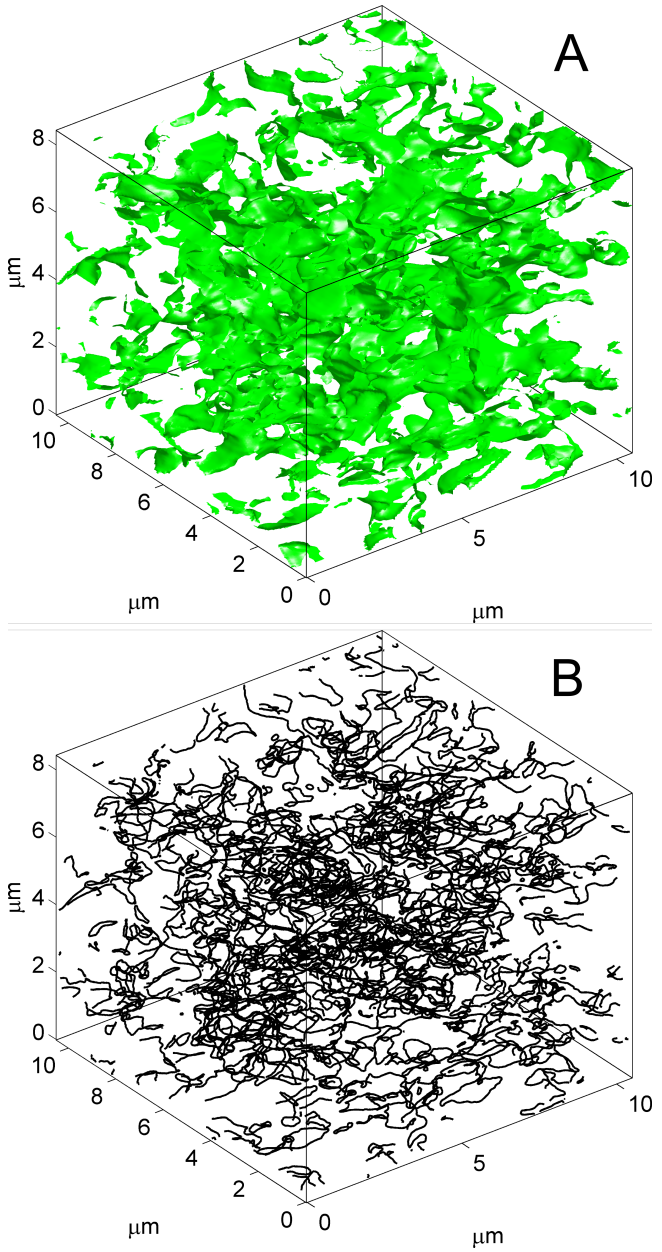

Figure 6.8: LSC/CGO electrode interface visualization. (A) The pore/LSC interface. (B) The TPB curves.

6.4 Conclusion

In this paper a method for accurate calculation of two phase interface area and TPB curve length is presented. The method reconstructs the phase interfaces as a non-manifold polygonal mesh of the structure based on sub-voxel accuracy segmentations of two of the phases. After the creation of the mesh the two phase interface areas are extracted as the sum of the area of polygons that separate sets of two phases. The TPB curve is extracted as the line segments in the mesh that join three polygons of different phase. This TPB curve definition ensures that the extracted TPB curve is a single connected loop of line segments without branches.

The accuracy of the method is analyzed based on calculations performed on perfect spheres of varying radius. The accuracy of the method is dependent on the radius of curvature of the phase structures in the sample. For radii of curvature larger than 5 voxels a relative error between -1.6% and 2.2% was achieved as a worst case estimate for TPB calculations on sub-voxel accuracy setups with two smoothing iterations. The TPB estimate is shown to have asymptotical very high accuracy for future increases in voxel resolution.

It is shown how graphs of the relative error as a function of radius can be utilized to assess the accuracy of the interface calculations based on knowledge of typical radii of curvature of the interfaces and the TPBs.

The reconstruction algorithm is used to reconstruct the phase interfaces of a SOFC electrode. The two phase interface areas and TPBs are calculated and a qualitative overview of the extracted TPB curves is presented.

Acknowledgements

This work was supported financially by The Programme Commission on Sustainable Energy and Environment, The Danish Council for Strategic Research, via the Strategic Electrochemistry Research Center (SERC) (www.serc.dk), contract no. 2104-06-0011.

Chapter 7

Geometrical characterization of interconnected phase networks in three dimensions

P. S. Jørgensen[a], **K. V. Hansen**[a], **R. Larsen**[b], **J. R. Bowen**[a]

(a) Fuel Cells and Solid State Chemistry Division, Risø National Laboratory for Sustainable Energy, Technical University of Denmark, Frederiksborgvej 399, 4000 Roskilde, Denmark.

(b) Department of Informatics and Mathematical modelling, Technical University of Denmark, Richard Petersens Plads, Building 321, Denmark.

Manuscript submitted for publication 2010.

Abstract

In electrochemical devices such as fuel cells or batteries the microstructure is a determining factor for the performance of the device. To be able to optimize the microstructure it is important to be able to quantitatively measure key structural parameters, such that systematic studies can be made. We present several general methods for quantitative characterization of network structures without prior assumptions of shape or application. The characterization is performed by extracting distributions of values rather than single value descriptions, thus allowing more detailed comparisons between samples to be made. The methods characterize tortuosity, path diameters, the novel dead ends property and a particle shape independent alternative to a *particle size distribution*. The parameters are calculated by the computation of arrival time maps by the Fast Marching Method. The methods are applied to the analysis of each of the three phases in a solid oxide fuel cell sample.

7.1 Introduction

Many devices such as fuel cells or batteries consist of microstructures whose 3D morphology has a large impact on their performance. The acquisition of 3D image data of the microstructure of a sample can be obtained through methods such as focused ion beam (FIB) tomography, micro computed tomography (CT) or serial mechanical sectioning for coarse structures. These data acquisition techniques have facilitated the analysis of sample properties that were previously not feasible to directly quantify. In particular solid oxide fuel cell (SOFC) electrodes have benefited from the progress in the analysis of 3D sample data [46]. Properties such as triple phase boundary length [51], phase percolation [73] and interface areas [49] have been successfully quantified.

This paper focuses on calculating quantitative properties relating to the network of a specific phase in a sample. This could for instance be the pore phase of an SOFC but the techniques described here are applicable to any phase of any sample in general. As a result the extracted quantitative values do not seek to calculate a specific physical value such as effective tortuosity. Rather, it seeks to quantitatively characterize the phase networks in terms of geometrical quantities such as the width of network paths or the degree of dead ends in the structure.

Software based on finite element modeling allows the modeling of for instance gas flow through a network. However, such calculations will not necessarily offer any insight into why the gas flow is good or bad. The goal of the methods presented here is to offer structural parameters that quantitatively characterize the shape and connectivity of phase networks without any assumptions on sample type or network shape. The proposed parameters can thus be used in a direct comparison between samples to establish how the networks in the samples differ from each other. These comparisons can in turn provide valuable information about why a particular microstructure is superior compared to another when related to measured properties.

The parameters described in this paper can be divided into four main areas.

- **Geometrical tortuosity**. The word geometrical was chosen to emphasize that the measured quantity only quantifies the degree of network paths being tortuous in a geometrical sense. It will thus not necessarily be directly usable in, for instance, gas diffusion calculations. It quantifies the extra length of travel through the structure due to the path being tortuous.

  Quantification of tortuosity has been described by several contributors. Iwai et al. [51] calculated the tortuosity factor of a phase network using both a random walk method and the Lattice-Boltzmann method. Wilson et al. [46] quantified the tortuosity from the 3D data set using a steady-state finite difference calculation. Gostovic et al. [49] measured the length of the path between the opposite sides of the voxel grid through the center of mass locations of the network.

- **Network cavities/agglomerates**. The quantification of the size and distribution of different parts of the network. It indicates whether the structure is made up of few large structures or many smaller structures.

  Al-Roush et al. [91, 92] performed a thorough investigation of 3D particle characterization, including: particle size distributions, particle contact network analysis and packing efficiency. The analysis assumes a network constructed from a packing of spheres. It is our observation that real experimental microstructure (SOFC electrodes in particular) often deviates substantially from one that would be expected from a packing of spheres.

  Wilson [53] calculated a characteristic diameter of network tubes based on the volume to surface area ratio. Münch [41] performed a robust calculation of a particle size distribution on continuous structures without any assumptions of shape.

- **Characteristic network diameter**. Quantification of the thickness of the pathways that make up the phase network. It quantifies the abundance

and severity of venturis.

Al-Roush & Willson [91] divide the network into pore-bodies and pore-throats and calculate the diameter inside each discrete pore-body and pore-throat.

- **Dead ends**. Quantification of the parts of the network that are not part of the shortest path pathways. A dead end is thus defined as a part of the structure that when entered brings one further away from the goal rather than closer to it (see section 7.3.4 for details). To the authors knowledge this property has not previously been quantified.

For all the calculations it is assumed that a 3D voxel grid of segmented [1] phase data is already available. Such a voxel grid consists of a structured 3D grid of voxels, where each voxel has been assigned one of a limited number of integer numbers. One number is assigned for each phase in the voxel grid. This paper deals with extraction of geometrical properties from phase networks in general and as such only one phase at a time is utilized in the calculations. This means that the voxel grid containing information about all the phases in the structure can be reduced to a voxel grid where 1 denotes the phase of interest and 0 denotes any other phase. This binary voxel grid is the basic input for any further calculations.

7.2 Theory

The objective of this paper is to extract quantitative parameters that describe the phase network of a specific network. To facilitate this, a method for calculating distances within the structure is needed. The fast marching method [93, 94, 95] (FMM) is well suited for this task, as it allows rapid calculation of distances between any two voxels.

7.2.1 The Fast Marching Method

The FMM is a computational technique for tracking the evolution of an expanding front over time. The FMM approximates the solution to the non-linear Eikonal equation

$$\|\nabla T(\mathbf{x})\| F(\mathbf{x}) = 1,\ F(\mathbf{x}) > 0 \tag{7.1}$$

where $T(\mathbf{x})$ is the arrival time at the location $\mathbf{x}$ and $F(\mathbf{x})$ is the speed of the front. The front can only expand and the arrival time is thus a scalar value for each location. By solving the Eikonal equation (7.1) for arrival time T on a regular grid of locations an arrival time map is obtained. The value of this map at a location $\mathbf{x}$ corresponds to the time the front arrives at that location.

The FMM is a computational rapid method of solving the Eikonal equation. The FMM gains its speed by essentially solving equation (7.1) in a single pass. The central idea is to build the solution outwards from the smallest time value T.

The voxels in the domain are classified into three types:

- *Accepted*: The voxels for which the arrival time has already been calculated.
- *Trial*: The voxels for which a tentative arrival time has been calculated.
- *Far away*: The voxels who are neither *accepted* or *trial*.

At initialization a set of voxels V_S are chosen as the source and classified as *Accepted* with the value $T = 0$. These voxels are the source from which the front expands. The arrival time is calculated for each of the neighbors to the initial voxels using only information from *Accepted* voxels. These voxels are then classified as *Trial*.

An essential part of the algorithm is a Binary Heap data structure [96] that efficiently retrieves the *Trial* voxel with the smallest tentative arrival time. A sorted list of tentative arrival times are maintained in the Binary Heap at all times which means that the lowest valued voxel can be retrieved by removing the voxel at the top of the heap.

The main loop is as follows:

1. Find the *Trial* voxel v with the smallest tentative arrival time.
2. Change classification of v to *Accepted*.
3. Calculate the arrival time of the non *Accepted* neighbors of v and update their classification to *Trial*.
4. If the heap of *Trial* voxels is empty then stop, otherwise go to 1.

Using the above algorithm the FMM solves the Eikonal equation with loglinear computational complexity, $O(NlogN)$, where N is the number of voxels in the voxel grid. The reader is referred to [95] for calculation of arrival times and implementation in detail.

If $F = 1$ across the domain, the arrival time also corresponds to the distance to the point. The FMM can thus be used to calculate the shortest path distance map between one set of voxels and all other voxels. As we shall see in section 7.3, F can be changed to give different weights to different parts of the domain, thus extending its functionality. Section 7.3 provides details on how the FMM can be utilized to calculate a number of geometrical properties from a voxel grid of segmented phase data.

7.3 Experimental methods

7.3.1 Geometrical tortuosity

We define geometrical tortuosity as the extra length an entity must travel from one location to another due to the transport network being twisted or tortuous.

In 2D, tortuosity is well defined for a curve between its two end points. Let L_e denote the effective length of the curve and L the length of the straight line between the end points of the curve. Tortuosity of 2D curves can then be defined as

$$\tau = \frac{L_e}{L}. \tag{7.2}$$

When extending the tortuosity definition to branching 3D networks the 2D curve formulation becomes ambiguous since pathways no longer are infinitesimally thin and no longer have a single start and end point. Additionally, because the network can branch and merge the 3D tortuosity is no longer a function of a single pathway but a measure that describes a multitude of pathways in the network.

The proposed calculation of geometrical tortuosity attempts to make a direct extension of the 2D tortuosity definition in equation (7.2) to 3D. A tortuosity parameter is sought that characterizes the network of a specific phase in a volume, henceforth known as the network.

Designate one plane as the source plane and another as the destination plane. A tortuosity value is then defined for each location $\mathbf{x}$ on the destination plane where it intersects the network

$$\tau(\mathbf{x}) = \frac{L_e(\mathbf{x})}{L(\mathbf{x})} \tag{7.3}$$

The term $L_e(\mathbf{x})$ denotes the length of the shortest path through the investigated phase from the source plane to the location $\mathbf{x}$. The term $L(\mathbf{x})$ denotes the shortest path through any phase from the source to the destination. By defining L and L_e for each location on the destination plane this definition can thus be used for calculations on non parallel source and destination planes. Note that since a value τ is obtained for each network location on the destination plane the result will be a distribution rather than a single value. Statistical moments or percentiles are then applied to provide single values that characterize the distribution.

To calculate the location specific L and L_e of equation (7.3) the FMM as described in section 7.2.1 is used. Instead of calculating the distances for one location at a time a distance map is created that contains distances to every location.

Let the network be represented as a binary 3D voxel grid. The source voxels V_s for the FMM algorithm are chosen as those network voxels that belong to the tortuosity source plane. The $\mathbf{F}$ scalar field controlling the speed is set to 1 for voxels that belong to the network and 0 for the rest of the voxel grid. This definition of the F field means that the calculated arrival time has the same value as the distance inside the network. The zero speed means that the arrival time outside the network is undefined.

To better show how the calculations are performed a 2D test network has been created, see figure 7.1. This network has been created to both contain shapes that are circular and tubular and contain more than one path between the bottom row and top row of the image. Note that although examples are given in 2D the methodology extends trivially to 3D.

Figure 7.2 shows the result of the FMM calculations on the setup described above. For this figure the source plane was set as the bottom row of pixels in the image and the colors in the figure correspond to the calculated distance to that location. Note that the distance is undefined outside the network but has been set to zero only for contrast purposes.

After the creation of the distance map the $L_e(\mathbf{x})$ values are obtained as the

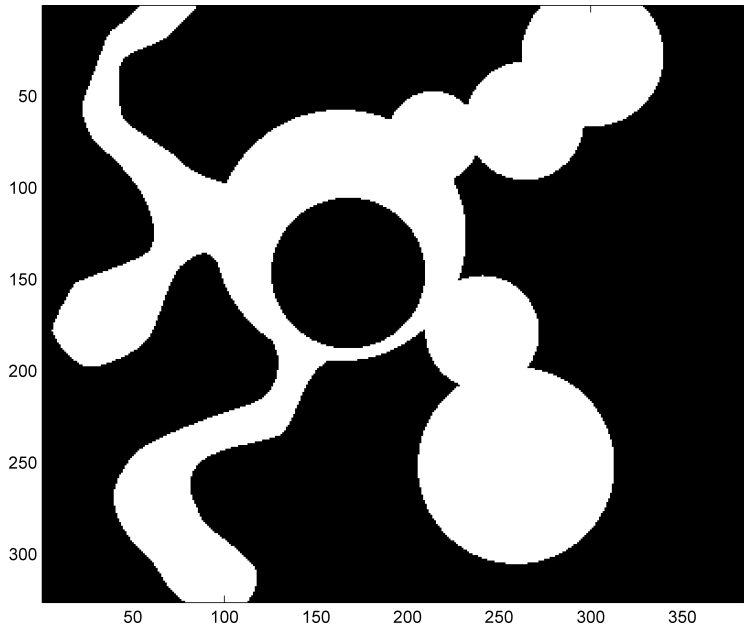

Figure 7.1: The 2D test network. White denotes the network of interest.

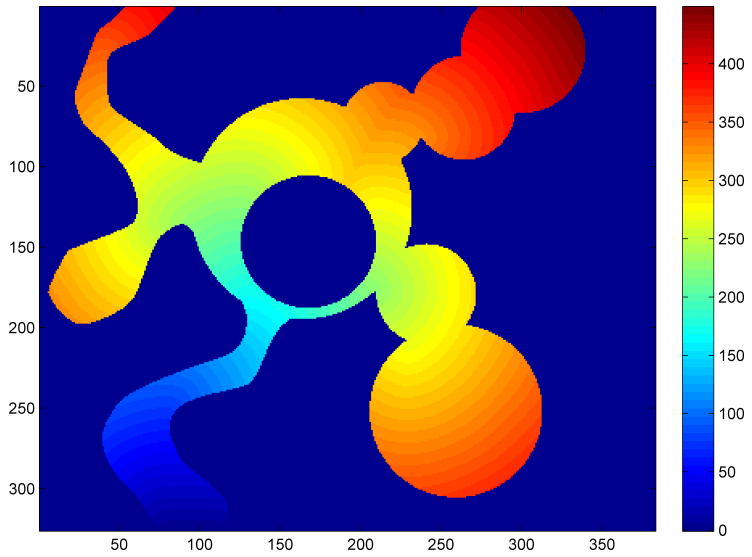

Figure 7.2: The distance map calculated by the FMM. The distance values outside of the network have been set to zero to increase contrast between the network and its surroundings.

distance map values at the destination plane where the location $\mathbf{x}$ is part of the network. Note that these values are the shortest distance from $\mathbf{x}$ to any network location on the source plane.

The values $L(\mathbf{x})$ can be obtained similar to $L_e(\mathbf{x})$. The only change is that the source and destination voxels are not restricted to be part of the network and F is set to 1 for all voxels. If the source and destination planes are parallel, L will simply be the distance between the planes and be constant for all $\mathbf{x}$. However, this method of calculating L allows us to calculate geometric tortuosity for a more general group of source and destination definitions. The source and destination voxels are essentially just two sets of voxels and as such they are not restricted to be parallel or even planes but can have any shape.

With $L_e(\mathbf{x})$ and $L(\mathbf{x})$ calculated the values of the geometric tortuosity distribution are calculated by equation (7.3).

7.3.2 Interface distance distribution

The interface distance distribution (IDD) (see figure 7.3) is similar to the more commonly used particle size distribution. Like the particle size distribution, it seeks to characterize how the sizes of different parts of the structure vary. However, instead of characterizing the structure as a distribution of particles it characterizes it as a distribution of distances from the interface between the network and other phases. That is, for each location $\mathbf{x}$ inside the network the shortest distance to an interface is calculated. The IDD is then defined as the distribution of these distance values.

As in section 7.3.1 the calculations consist of calculating a distance map from a set of source voxels V_s. All non-network voxels are included in V_s and the speed F is set as constant 1 for all voxels. Running the FMM with this setup results in a map that has the value zero outside the network and the value of the distance to the interface inside the network.

The result of the FMM calculations on the test network of figure 7.1 can be seen in figure 7.3 (The dark curves through the structure are explained in section 7.3.3). The IDD is simply the distribution of the distance values at each network voxel, which also corresponds to all non-zero valued voxels.

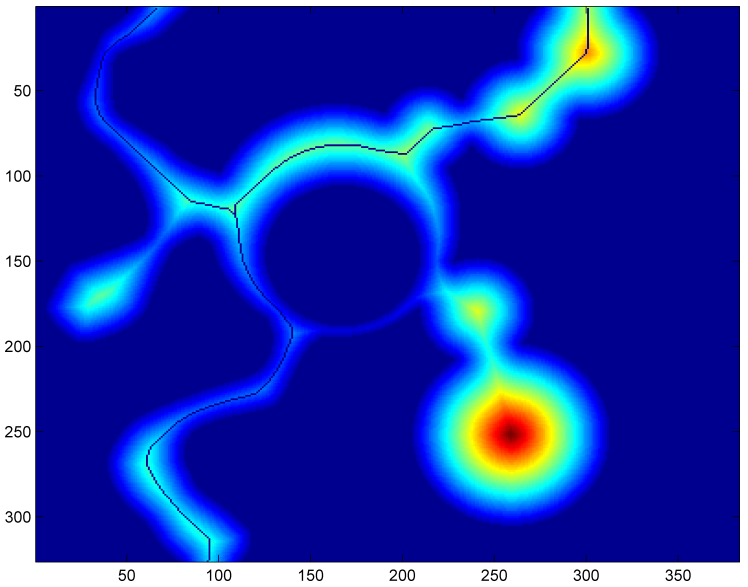

Figure 7.3: The interface distance map calculated by the FMM. The marked path through the network marks the location of the thickness sampling locations.

7.3.3 Characteristic path diameter

The characteristic path diameter seeks to characterize the local diameter of the network path. For any section of path in the network we define the characteristic path diameter as the diameter of the largest sphere (or circle in 2D) that can pass through that section of the network.

The value of the radii of these spheres is previously calculated in section 7.3.2 as part of the interface distance map. The radii can be obtained at those network locations that are equidistant and furthest away from the network boundaries. However, a method of sampling only these locations is needed.

To find the sampling locations the FMM is utilized again. The idea is to perform a new shortest path calculation similar to the one in section 7.3.1. However, the F scalar field is weighted based on the distance to the network boundary rather than uniformly having the value 1 inside the network. This change of $\mathbf{F}$ means that the result of the FMM no longer can be interpreted as distances but must be seen as arrival times. However, the value of the shortest arrival time is not of interest in itself but rather the fastest path that results in this arrival time. The locations that make up the fastest path can then be used directly as the

sampling locations.

Let $d_n(\mathbf{x})$ denote the distance to the network boundary at location $\mathbf{x}$ as previously computed in section 7.3.2 (figure 7.3). The speed scalar field $\mathbf{F}$ is set to 0 at locations outside of the network and to the value

$$F(\mathbf{x}) = F_{max}^{(\frac{d_n(\mathbf{x})}{d_{max}})} \tag{7.4}$$

inside the network. The constant F_{max} is an upper limit on the value that $F(\mathbf{x})$ can attain and d_{max} is the maximum value of $d_n(\mathbf{x})$ for all $\mathbf{x}$. The constants are introduced to make the calculations independent of the scale of the structure. The value of F_{max} is chosen from numerical considerations to allow summation of very large and very small numbers. The effect of F_{max} is discussed in section 7.4. A value of $F_{max} = 10^{10}$ has been used for all calculations both in 2D and 3D.

The FMM is run with the same source voxels as defined in section 7.3.1 but with the new speed field $\mathbf{F}$. The result is an arrival time map where each location corresponds to the time the front takes to reach that location.

The exponential scaling of $F(\mathbf{x})$ with $d_n(\mathbf{x})$ acts as a soft constraint on the fastest paths through the network. Since $\mathbf{F}$ controls the speed of the expansion, the front will expand much faster in the center of pathways than anywhere else since d_n has the largest value here. The fastest paths through the network will thus be the ones that follow the center of pathways and use the widest paths.

The fastest path through the network is extracted by backtracking through the arrival time map starting at the lowest arrival time voxel on the destination plane. However, this only finds a single path through the structure and would be a poor representation of large networks with many branching or separate pathways. To sample the thickness of these pathways as well, the set of starting voxels for which to backtrack from is expanded. The larger set of backtracked voxels consists of the lowest arrival time voxel within each connected region on the destination plane. This is exemplified in figure 7.3 where the fastest pathways through the structure has been marked, one pathway starting within each connected network region on the top row of the image (the destination plane). Note how the paths from the two distinct starting locations eventually merge to one path.

The backtracking algorithm itself is straightforward. Let $\mathbf{V_{start}}$ denote the set of starting voxels from which to backtrack and $\mathbf{V_{path}}$ denote the set of voxels that make up the optimal path (initially empty). The pseudo code of the algorithm

is as follows:

for each voxel in $\mathbf{V_{start}}$ **do**
  Set the current voxel v_c to the starting voxel.
  while v_c is not in the destination plane **do**
    if v_c is not in $\mathbf{V_{path}}$ **then**
      Add v_c to $\mathbf{V_{path}}$
    end if
    Set v_c to the lowest valued neighbor of v_c in the arrival time map.
  end while
end for

A neighbor is defined as one of the 26 voxels adjacent to the voxel in a 3D voxel grid. The result of running the algorithm on the 2D example can be seen in figure 7.3. The values of the characteristic path diameter distribution can now be obtained as twice the interface distance of the voxels that make up the paths, $d_a = 2d_n(\mathbf{V_{path}})$

7.3.4 Dead ends

The analysis of dead ends is an attempt to quantify the amount of dead end branches in the network. These branches are called dead ends since traversing them can only result in lengthening the distance traveled between two locations. One could say that they are considered detours.

An analysis of the percolation of a network quantifies how much of a given phase is part of the percolating network. The analysis of dead ends further characterizes the percolating network by quantifying how much of the network consists of dead ends. The dead end analysis is thus a further tool in characterizing how a given phase is distributed in the sample.

It is important to point out that dead end branches only are dead ends in a geometrical context. The dead ends might contribute to important physical phenomena dependent on which phase is being investigated and what context the physical sample is used in.

To analyze dead ends a clear definition of what constitutes a dead end is needed. A definition is chosen that defines dead ends in terms of the shortest distance between two locations. We define a dead end as a part of the network that when traversed does not bring us closer to the destination. This definition is better explained by the algorithm from which it is calculated.

The basis for the calculation is a shortest distance map as was constructed in section 7.3.1. Designate one set of voxels as source voxels V_s and another set as destination voxels V_d. Define an initial set of active voxels V_a as all the destination voxels. Let $\mathbf{M_{status}}$ be a map of the status of all voxels in the voxel grid, with possible states of *dead* or *alive*. Set the state of all initial active voxels to *alive* and all other voxels to *dead*. Let $d_n(\mathbf{x})$ denote the shortest distance from the source to the voxel $\mathbf{x}$. The partition of the network into *dead* and *alive* voxels can then be performed by the following algorithm.

while V_a is not empty **do**
  Remove the voxel v_c from V_a
  for each neighbor v_n of v_c **do**
    if v_n is *alive* **then**
      Skip this neighbor
    end if
    if $d_n(v_n) < d_n(v_c)$ **then**
      Add v_n to V_a and change the state of $\mathbf{M_{status}}(v_n)$ to *alive*
    end if
  end for
end while

A neighbor is again defined as one of the 26 voxels adjacent to the voxel in a 3D voxel grid. The algorithm is thus a simple region growing scheme where a region is iteratively expanded into neighbors with lower distance values.

The result of running the algorithm on the 2D test example structure can be seen in figure 7.4. The figure shows the result of running the algorithm on distance maps calculated from two opposite directions, 7.4(A-B and C-D). Note that the algorithm gives different results depending on the direction it is run. However, the calculations from the two directions can be combined by a set intersection of the dead ends obtained from each direction (7.4E), thus yielding a direction invariant result. The dead ends in the network can be quantified by the ratio of *dead* voxels to *alive* voxels in $\mathbf{M_{status}}$.

A further analysis of the depth or severeness of the dead ends can be performed by using the FMM to calculate a distance map from the *alive* part of the structure (7.4F). The distribution of distances can then be analyzed similarly to the boundary distance in section 7.3.2.

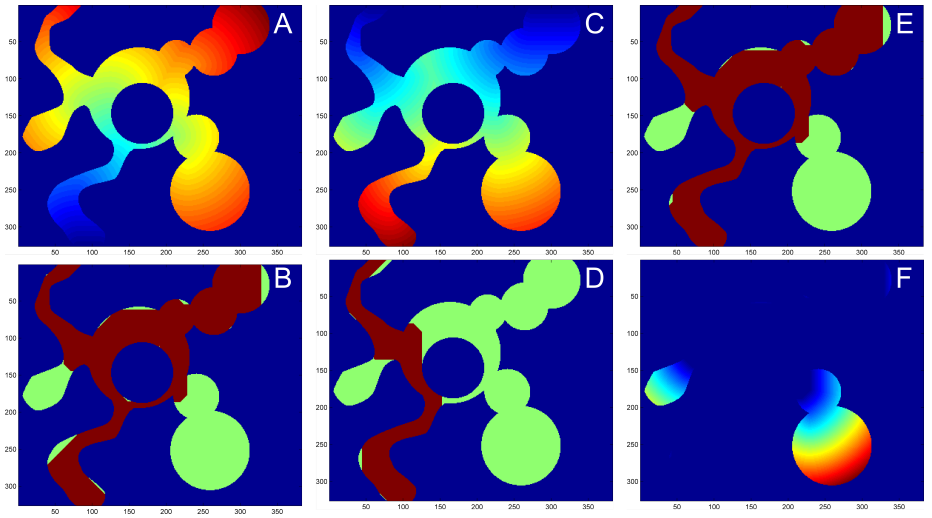

Figure 7.4: The calculation of dead ends. (A) The distance map calculated from the bottom row. (B) The dead ends calculated from A (dead ends in green and alive network in red). (C) The distance map calculated from the top row. (D) The dead ends calculated from C. (E) The dead ends obtained by the intersection of the results from B and D. (F) The distance map calculated from the alive part of E.

7.3.5 SOFC sample background

To show how the methods described in this section work in 3D the methods have been applied to an SOFC electrode sample. The investigated SOFC sample is a composite SOFC cathode consisting of 50/50 weight-percent strontium-substituted lanthanum cobaltite (LSC) and gadolinia-substituted ceria (CGO), similar to [81]. The calculations were performed on a grid of $150 \times 150 \times 150$ voxels with a voxel dimension of $58.6 \times 58.6 \times 44.5\,nm$. The segmentation was carried out as described in [1].

All the calculations were performed between the bottom and top planes of the voxel grid using only parts of the network that formed a percolating network between these planes. The phase volume fractions of the percolating networks were 0.1593 for the pore phase, 0.2913 for LSC phase and 0.5494 for the CGO phase.

7.4 Results

The methods described in section 7.3 are used to calculate parameters that do not necessarily correspond to a property that can be measured by physical devices. The calculated characterization values thus have little value on their own. To provide insight the parameters must be compared to the same calculation performed on a different network. To form a basis for comparison the networks of each of the three phases: pore, LSC and CGO are compared to each other. This comparison serves as an example of the methods. A comparison between the same phases of different samples would typically be of more interest in larger experimental studies of microstructure variation. The results are presented here to show how the methods can be utilized to analyze the network structure of a real sample. Emphasis will thus be put on illustrating the methodology rather than analyzing the sample specific implications of the results.

All histograms in this section use normalized counts such that the value for all bins sum to 1 for each phase network. All histograms were created using 20 bins equally distributed between the lowest and the highest value of all 3 phase networks combined.

7.4.1 Geometrical tortuosity

The results for geometrical tortuosity can be seen in figure 7.5. The figure contains the pooled results for calculations performed from the bottom and top layers of the voxel grid. The distribution of geometrical tortuosity for each phase is quite distinct. The CGO phase network generally has low geometrical tortuosity values. This is expected since the CGO phase is significantly the most abundant and will thus have more and wider path options. The LSC phase has the highest geometrical tortuosity values with the bulk of the values forming a peak at 1.5. The distribution contains a significant number of high valued observations between 1.7 and 2.1. By 3D inspection of the voxel grid it was observed that the high values are due to a number of narrow and twisted paths in the LSC phase. These interesting details would not have been as easily observable had the geometrical tortuosity been a single mean value rather than a distribution.

It is interesting to note that the pore phase distribution is bimodal. The pore phase has a small number of samples (4312 samples compared to 10774 for LSC and 24298 for CGO). Although this is a sufficient number for a histogram in general, keep in mind that the samples are strongly correlated with their connected neighboring voxels. By inspection of the voxel grid the pore distribution value samples were found to be taken from 5-6 distinct network paths with significantly different values. In contrast the LSC and CGO networks have better connected backbone like structures resulting in more uniform values. The number of pore network paths is too low to draw any conclusions about whether the two peaks are characteristic for the cell or a sample volume artifact.

The geometrical tortuosity calculations are sensitive to the size of the voxel grid. The dimensions of the voxel grid should be considered carefully to obtain robust results. If the distance between the source and destination planes (the thickness) is too short the majority of the paths will pass straight through without the need to turn or twist. The resulting geometrical tortuosity values will thus be close to 1. The smallest thickness that provides robust results is network dependent. However, a test for sufficient thickness can be performed by observing how the results change for varying source and destination plane spacing. The thickness is sufficient when the geometrical tortuosity does not decrease with decreasing thickness.

The two other voxel grid dimensions should be large enough to allow the calculations to be based on a sufficiently large number of distinct paths to avoid the problems observed for the pore network. Similar to the thickness a sensitivity analysis can be performed to determine if the dimensions are sufficient. Tests performed on the SOFC sample indicated that the voxel grid dimensions where

sufficient for the LSC and CGO networks but insufficient for the pore phase. The geometrical tortuosity parameter has previously been applied to the analysis of a graphite negative electrode from a Li-ion battery [97] where the effect of varying grid dimensions was analyzed.

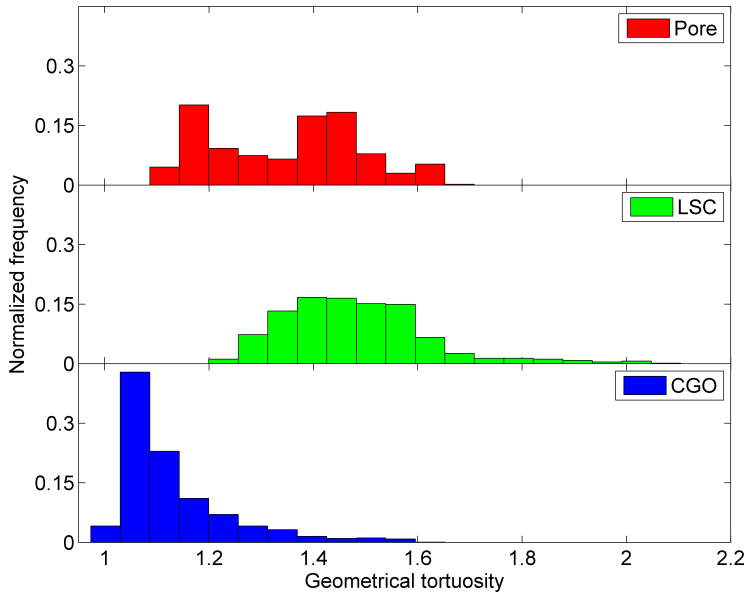

Figure 7.5: The geometrical tortuosity distributions for the SOFC sample

7.4.2 Interface distance distribution

The results for the IDD can be seen in figure 7.6. It is interesting to note that in figure 7.6 the pore and LSC networks both have their peaks at the lowest valued bin while the CGO network does not. This is due to the backbone nature of the CGO phase. For minor phase networks, consisting of pathways through an abundant phase, the surface area is small compared to the volume of the phase. This means that the IDD consists of many observations with a small interface distance and fewer with a larger interface distance. For a backbone structure this relation is opposite resulting in relatively more high valued observations. Another interesting observation is that the LSC network contains the largest agglomerate even though the CGO phase is the most abundant.

The IDD should be seen as an alternative or supplement to a conventional particle size distribution (PSD) [41]. The main motivation behind the IDD is to provide a description of network cavity/agglomerate sizes without any

assumptions on the shape of the network. The IDD essentially contains the same information as a PSD although it is presented differently. A PSD makes little sense for networks that often consist of essentially one large connected particle. However, the PSD performs well when the network consists of tubular or spherical cavities/agglomerates since it provides an intuitive description of these structures. However, for other types of network structures the results can be uninformative or even misleading. One example is a backbone structure exemplified by the CGO network of the SOFC sample. Although a particle size distribution can be calculated for this structure an interpretation of the results as a distribution of spheres would be misleading.

One direct application of the IDD is in optimization of surface area. Even though the interface area can be measured directly [48] the IDD provides additional information by quantifying how much of the volume is occupied by large cavities/agglomerates. Take as an example the CGO network of the SOFC sample where 16.2 % of the structure is taken up by material further than 300 nm away from the interface. This might tell us that the structure could be further improved by better utilization of the material to form more abundant but thinner pathways. In other words, this information makes it possible to evaluate to what extent the structure can be further optimized.

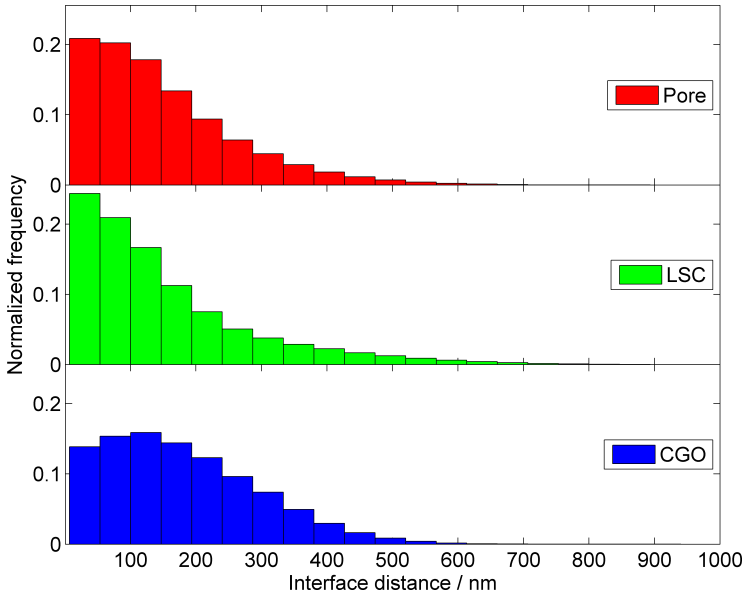

Figure 7.6: The interface distance distributions for the SOFC sample

7.4.3 Characteristic path diameter

A visualization of the characteristic diameter of the sampled pathways for the SOFC sample can be seen in figure 7.7. The figure was constructed by plotting a circle at each of the diameter sampling locations with the diameter of the circle proportional (for visualization purposes) to the characteristic diameter of the pathway. To reduce the number of pathways in the figure the calculations were only performed from the bottom to the top of the voxel cube. This type of plot allows a visual identification of where the main bottlenecks in the structure are located. The corresponding characteristic diameter distribution can be seen in figure 7.8. As would be expected, the CGO phase consists of pathways with larger diameters than the pore phase pathways. However, it is interesting to note that the CGO phase still has a significant fraction of pathways with a small diameter which could be critical bottlenecks. The LSC distribution is bimodal as was previously seen in the tortuosity distribution in figure 7.5. As was discussed earlier the two peaks are caused by the LSC phase consisting of both wide and narrow pathways. Although these results are interesting a larger part of the network should be analyzed to obtain robust results that are representative for the sample.

An alternative approach to finding sampling locations for the characteristic path diameter is to perform a topological skeletonization or medial axis transform [98, 99, 100] of the network. The medial axis transform finds the set of points in the network that have more than one closest point on the network's boundary. This approach gives a good sampling of characteristic path diameter in the entire network. The suggested approach should thus not be seen as a replacement for this methodology but rather as an addition to it.

The suggested approach performs a more selective sampling of the path diameters than the medial axis approach. Diameter sampling is only performed on the major pathways rather than on every pathway due to the scaling of the $\mathbf{F}$ speed field with the distance to the interface. This effectively means that if a network contains more than one path to a given location only the widest of the paths will be sampled. The selective sampling results in a diameter distribution that describes the widest possible pathways that can be used to traverse the network. Network bottlenecks are thus easy to identify since a small diameter in the distribution means that no alternative wider pathways exist. However, the selective sampling also means that a larger part of the network must be analyzed to obtain robust results because fewer pathways are sampled.

There is a downside to using a soft constraint for the calculations because the optimal pathways with the shortest arrival times (as calculated in section 7.3.3) are dependent on F_{max}. This means that the calculated pathways can be rather

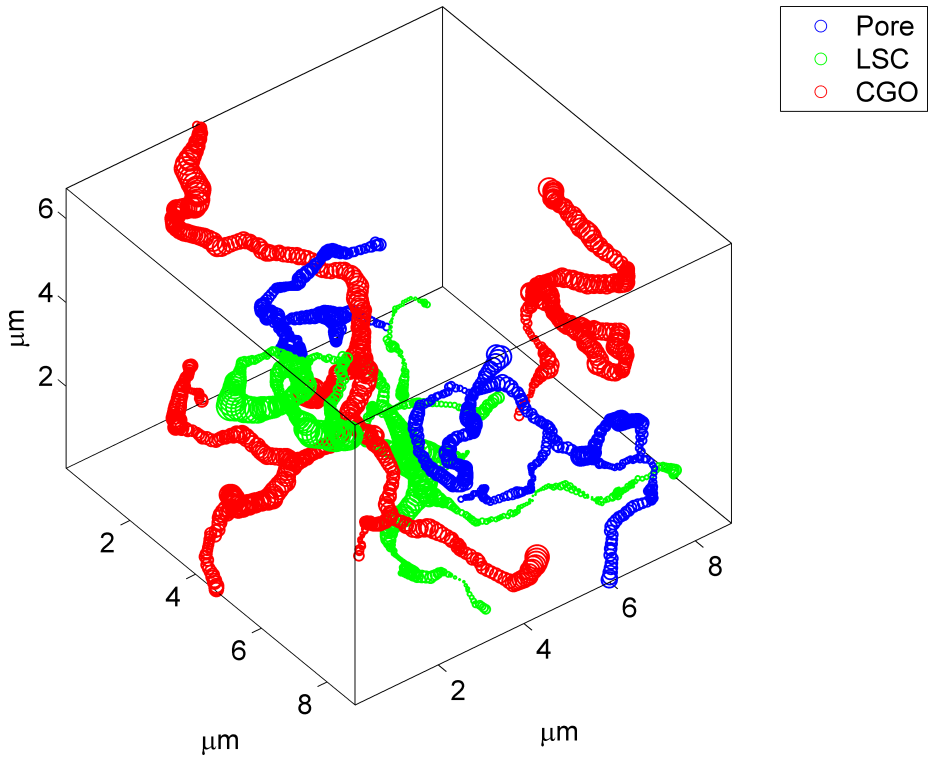

Figure 7.7: Visualization of characteristic path diameter. The diameter of each circle is proportional to the diameter of the pathway at that location. The calculations were performed from the bottom to the top of the voxel cube. Note that only percolating pathways are represented.

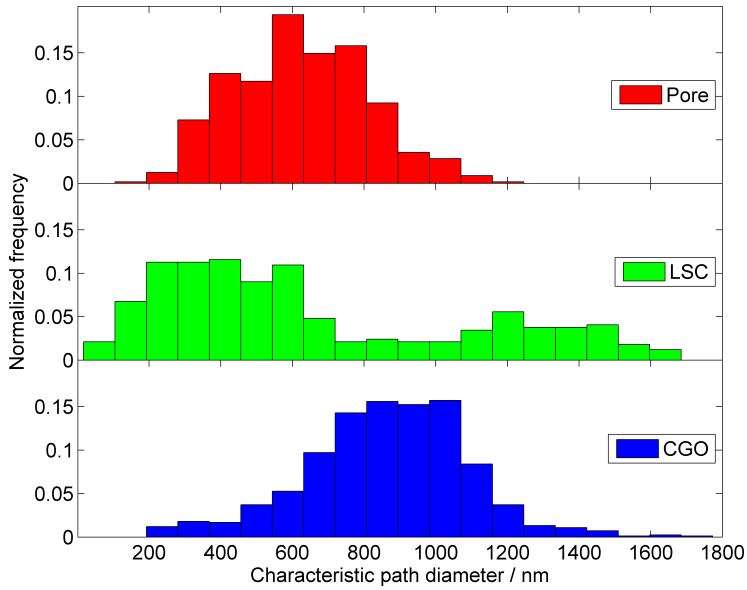

Figure 7.8: Distributions of characteristic path diameters for the SOFC sample

arbitrary in situations where the network contains several paths of similar diameter. To illustrate this, consider a network that contains two alternative pathways. One pathway being short and one being long but slightly wider. The pathway used in the diameter sampling will be the one with the shortest crossing time. The long path will have the largest speed values, but due to it being longer the arrival time could be greater than through the short but slightly narrower path. The F_{max} parameter determines which path is ultimately chosen. However, for large values of F_{max} the two alternative pathways will still have similar diameters. For fixed F_{max} the method can be seen as a random sampling of similar pathways. The accuracy will thus increase with increasing network size.

7.4.4 Dead ends

A visualization of the calculated dead ends in the pore phase can be seen in figure 7.9. The dead ends were calculated between the top and bottom sides of the voxel grid for visualization purposes. This results in more dead ends compared to calculating between all six sides. Only the percolating part of the structure is shown. The network volume fraction of dead ends was 0.278 for the pore phase, 0.245 for the LSC phase and 0.047 for the CGO phase.

To further analyze the dead ends the shortest distance to the alive part of the network was calculated for each dead end voxel (as shown in figure 7.4F). We refer to this distance as the dead end depth. The cumulative distribution of dead end depths for each of the phases can be seen in figure 7.10. It is seen that the pore phase consists of deep dead ends that, as is evident on figure 7.9, form narrow twisting pathways, whereas the few dead ends present in the CGO phase are shallow.

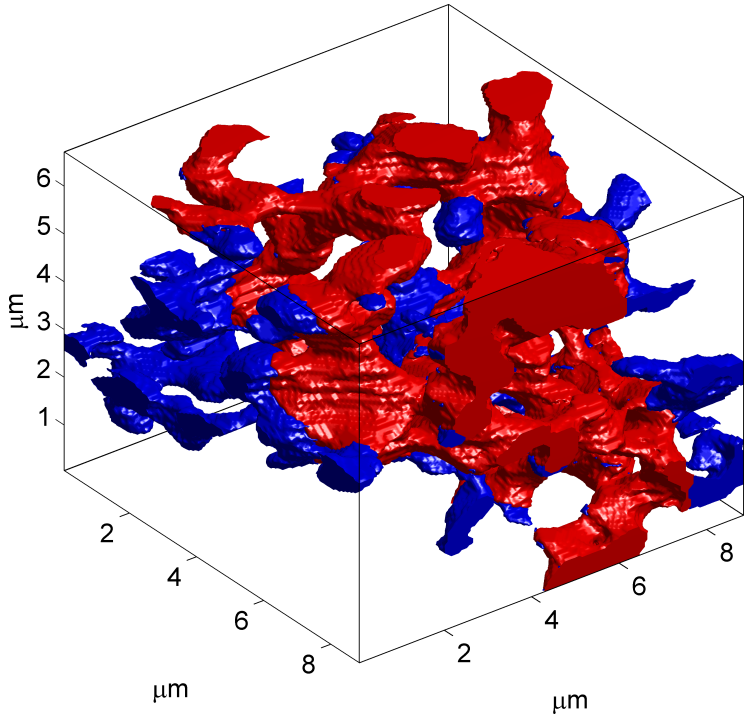

Figure 7.9: Visualization of the dead ends in the pore phase. Blue denotes the *dead* ends in the network and red denotes the *alive* part of the network.

7.4.5 Computational cost

The computational cost of the presented methods is taken up almost exclusively by the FMM calculations, where the time taken to run one FMM setup is dependent on the network. For the SOFC sample containing 3.4 million voxels the FMM took 2.3 s, 3.7 s and 8.2 s for the pore, LSC and CGO phases respectively on an Intel®Core™2 T7200 2GHz CPU. The FMM is run once for each direction analyzed, thus six times for a full six-way analysis.

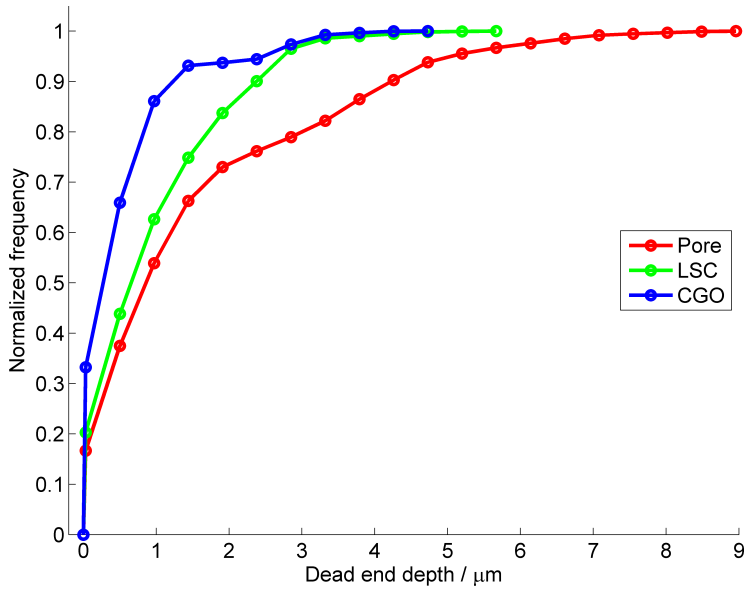

Figure 7.10: The cumulative dead end depth distributions for the SOFC sample

7.4.6 Future work

The methods presented here have focused on quantitatively characterizing each network structure separately. However, the performance properties of the investigated sample will often depend on the interaction between all the constituent phases. These interactions are, however, application dependent. One promising extension for SOFC characterization is to combine the results of TPB calculations with the methodology introduced here. The TPB sites could be used as the source or destination voxels for the calculation of characteristic path diameter and dead ends, thus yielding application specific results.

7.5 Conclusion

Four general methods are presented for characterizing network structure morphology, based on FMM calculations. The methods result in distributions of values rather than single value descriptors, thus allowing more detailed comparisons of network microstructure to be made between various samples.

The geometrical tortuosity parameter provides a general quantification of how

tortuous or twisted the network is without any assumptions of phase or application and can be calculated between any two sets of voxels.

An alternative to the particle size distribution was introduced in the form of the more general interface distance distribution. For non-spherical particle network structures the interface distance distribution can provide a more intuitive interpretation.

The characteristic diameter parameters were shown to allow easy identification of bottlenecks in the network structure. The novel dead end analysis allows quantification of how large a fraction of the network is occupied by dead ends and provides information about the severity or depth of dead ends.

The above methods are applied to the characterization of an SOFC electrode to illustrate the value to real microstructures. The methods reach their greatest potential when used to analyze how microstructure varies between groups of samples.

Acknowledgements

This work was supported financially by The Programme Commission on Sustainable Energy and Environment, The Danish Council for Strategic Research, via the Strategic Electrochemistry Research Center (SERC) (www.serc.dk), contract no. 2104-06-0011.

Chapter 8

Perspective and outlook

This chapter provides perspectives and outlook based upon the results of the previous chapters. An example of the use of the calculated parameters in electrochemical electrode modeling is given in the first section. The rest of the chapter discusses the general perspectives of the results, possible issues, and outlook on future work.

8.1 Electrochemical modeling perspectives

This section gives an example of how the measured values may be used in SOFC modeling by supplying parameters that are not otherwise easily obtainable. All experimental work and calculations not related to the 3D reconstruction were carried out by Søgaard et al. [101]. The model results are included here as perspective on the usability of the extracted 3D geometrical parameters for modeling purposes.

8.1.1 The ALS model

The Adler-Lane-Steel (ALS) model [102] describes the impedance spectrum of a porous mixed ionic and electronic conducting oxygen electrode. In this context a slightly modified ALS-model is used to model the impedance spectrum of an LSC-CGO composite SOFC cathode.

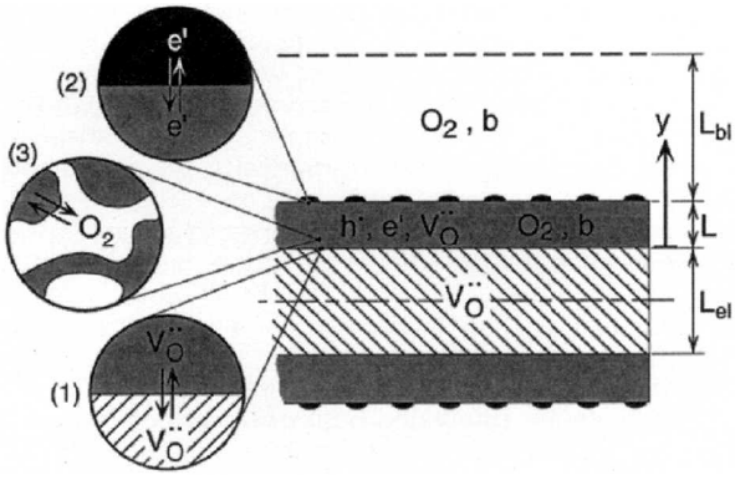

Figure 8.1: The cell geometry considered in the ALS model.

The model used here considers a symmetric cell (see figure 8.1) with an electrolyte of thickness L_{el}, two identical porous electrodes containing a mixed conductor and a purely ionic conductor (typically CGO) with thickness L and a boundary layer at each side with the thickness L_{bl}. The boundary layer represents the transition from perfectly mixed gas outside the boundary layer to a reaction dependent gas mixture towards the electrode surface. The overall reactions in this setup occur by the three interface reactions denoted by numbers 1-3 in figure 8.1. (1) Charge transfer of oxygen ion vacancies across the electrolyte/mixed conductor interface. (3) Chemical exchange of oxygen at the gas phase/mixed conductor interface. (2) Charge exchange of electrons across the mixed conductor/current collector interface.

The ALS model consists of a system of coupled partial differential equations:

$$\frac{\partial c_v^*}{\partial t} = AD_v \frac{\partial^2 c_v^*}{\partial y^2} - ar \tag{8.1}$$

$$c\epsilon \frac{x^*}{\partial t} = c\epsilon \frac{\chi_p}{\tau_p} D_{AB} \frac{\partial^2 c_v^*}{\partial y^2} - \frac{a(1 - x^0 \chi_p)}{2} r \tag{8.2}$$

$$\frac{\partial x_{bl}^*}{\partial t} = D_{AB} \frac{\partial^2 x_{bl}^*}{\partial y^2} \tag{8.3}$$

Equation 8.1 models the displacement of the oxygen vacancy concentration, c_v^*, in the mixed conductor from its equilibrium value. The parameter A is a thermodynamic factor, D_v is an effective transport parameter dependent on the volume fractions, tortuosities and transport properties of the mixed conductor and the CGO, a is the volume specific mixed conductor/gas phase interface area and r is the reaction rate of the oxygen-exchange occurring over the internal surface area of the mixed conductor.

Equation 8.2 models the displacement of the oxygen mole fraction, x^*, in the gas phase from its equilibrium value. The parameter c is the total gas phase concentration, ϵ is the porosity, χ_p is a correction factor for Knudsen diffusion, τ_p is the gas phase tortuosity and x^0 is the oxygen mole fraction at equilibrium.

Equation 8.3 models the displacement of the oxygen mole fraction in the boundary layer, x_{bl}^*, from its equilibrium value. Note that the full ALS model is significantly more complicated since several of the parameters in equation 8.1, 8.2 and 8.3 cover composite terms.

A detailed derivation of the ALS model is beyond the scope of this example and the reader is referred to [102]. Note that the system of equations 8.1-8.3 represents a slightly modified ALS model compared to [102] due to the presence of the CGO phase.

Six boundary conditions are defined from considerations of charge and mass conservation at the boundaries between regions of the cell (see [102]). The equations are solved by separation of variables to obtain analytical expressions for c_v^*, x^* and x_{bl}^*. The impedance of the half cell can then be expressed in terms of c_v^*.

Impedance spectroscopy is a powerful tool for the analysis of electrochemical devices. Impedance spectroscopy makes it possible to estimate the resistance contribution of the individual cell components by fitting an equivalent circuit model [103] to the impedance spectrum.

If all parameters are known the ALS model can be used to model the resulting impedance spectrum. Alternatively, the ALS model can be used to estimate unknown parameters given a measured impedance spectrum.

8.1.2 Modeling results

The modified ALS model was fitted to experimentally acquired impedance spectra of an LSC-CGO composite SOFC cathode measured at 42 different combinations of oxygen partial pressure and temperature. Selected impedance spectra and the model fits are shown in figure 8.2. The LSC/pore surface area (a) and the porosity phase tortuosity (τ_p) of equation 8.1 and 8.2 were calculated by the methods described in chapters 6 and 7. The average tortuosity of the tortuosity distributions was used. The porosity, ϵ, was easily obtained from the segmented image data by a simple summation of the voxels of each phase divide by the total number of voxels. The fit performed with the calculated geometrical parameters enabled estimates of the surface exchange coefficient and the vacancy diffusion coefficient to be determined. The values extracted using the model compared well with literature data and showed the same magnitude and tendencies as expected. That is, a vacancy diffusion coefficient which is constant with oxygen partial pressure, whereas the reaction rate decreased significantly with decreasing oxygen partial pressure. The reader is referred to [101] for detailed results and discussion.

8.1.3 Modeling parameters

This modeling example shows the perspectives of the calculated 3D parameters for modeling purposes. The porosity parameter can be estimated from 2D images. However, the 3D porosity estimate is based on a large amount of data and contains the effect of possible porosity gradients in all three dimensions.

From Equation 8.1 and 8.2 it is seen that a and r are only found as a product. This illustrates that it is not possible to distinguish between a larger available area or enhanced kinetics. Thus, it is imperative to have a precise measure of the volume specific area, if the aim of the analysis is to deduce precise reaction kinetics. The volume specific surface area can be estimated from 2D images by stereology. However, the 3D measurements are based on a larger amount of data and include the effects of possible spatial variation of surface area in all three dimensions. Tortuosity is not easily estimated from 2D image data and the 3D calculations benefit fully from being able to consider the connectivity of the phases in 3D.

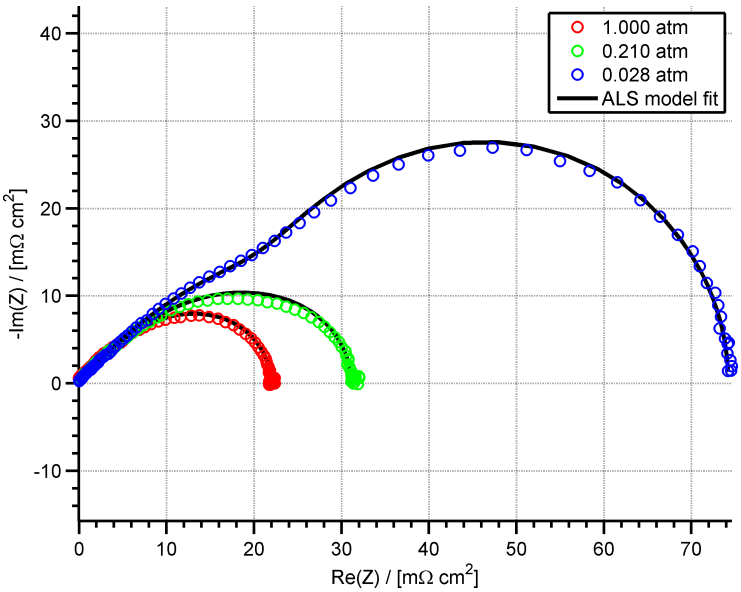

Figure 8.2: The ALS model fitted to impedance spectra acquired at 752 $°C$ for three oxygen partial pressures.

The modeling presented here is one dimensional but uses parameters determined by calculations on 3D reconstructions. The next step will be to model electrodes in full 3D. To do this, the voxel and mesh data from a segmented 3D reconstruction of an electrode microstructure should be exported to a finite element program. By setting up the differential equations and boundary conditions that govern the individual transport processes, a detailed model of for instance the impedance spectrum could be constructed based on the reconstructed 3D microstructure of a specific cell.

8.2 Scope of the calculated parameters

The calculated parameters in chapter 6 provide important characteristics of microstructure. The TPB density is one of the most important SOFC microstructure parameters and has a precise geometrical and physical definition. This means that, the TPB density parameter can be seen as a direct measurement of a physical property. The current computational approach by FIB tomography is currently the closest we get to directly measuring the TPB length. As with any physical quantity it is advantageous to be able to perform the measurement

with high precision and accuracy. The high accuracy of the TPB calculation methodology presented in chapter 6 makes the method attractive for studies that attempt to correlate TPB density with SOFC performance.

The network structure parameters in chapter 7 are calculated to supply quantitative parameters for a systematic optimization of parameters. These parameters are almost purely geometrical in nature and do not characterize a simple well defined physical property. This means that they will not necessarily be attractive to use for performance modeling purposes. The strengths of the parameters are in characterizing microstructure such that it can be directly and quantitatively compared to other microstructures. The parameters are designed for being attractive in studies of microstructure variation. This could be spatial variation in microstructure within the same sample or variation between groups of samples made with different recipes or production methods.

The goal of the network and interface parameter calculations is to be able to provide more detailed microstructure related answers to why specific cells are either high or low performing. As an example, suppose that a new cell recipe has been designed to improve the TPB density in the cell, however the cell shows poor performance. The parameters extracted from the 3D microstructure analysis reveals that the microstructure does indeed have a high TPB density but other microstructural properties are detrimental to the performance. The changes in the recipe might have caused large parts of the network structure to become non-percolating. Alternatively, the network structure of the pore phase might have changed to now consist of a combination of very wide and very narrow pathways, resulting in poor effective diffusivity and phase utilization. By comparing the calculated microstructure parameters between the new cell and the current state of the art, such microstructural problems would be identifiable.

The calculated parameters are not tied to the analysis of SOC. In fact, most of the parameters are applicable to two and three phase structures that exhibit similar stochastic phase distribution behavior. For instance, some of the parameters have previously been used for the analysis of the microstructure of a Li-ion battery electrode [97].

8.3 Importance of automation

Studies of microstructure variation or correlation of microstructure with performance requires many individual microstructures to be analyzed. To make large studies of microstructure variation practically feasible the analyses must be rapid and easy to perform.

The ambitious end goal of 3D microstructure analysis is a characterization tool that completely automates the process: from having identified the section of the sample to be analyzed to the parameter values are calculated. The goal requires both additional automation in the FIB/SEM system but also fully automatic computational methods and error handling.

This work has shown several steps towards achieving a higher degree of automation. This has been done through the automation of the image alignment and non-uniform illumination correction routines. The level set segmentation method described in 5 can be used to perform a segmentation of the 3D image data by a computer. The method still requires several parameters to be tuned manually. However, it has many desirable properties and can be designed to segment image data with non-consistent phase intensities. To fully automate the segmentation step the parameters must be automatically determined from the data as well.

8.4 Issues and outlook

Human bias

The analysis of the accuracy of the calculated parameters in this work assumes that there is an accurate correspondence between the phase structure in the physical sample and the phase structure of the segmented voxel grid. The accuracy and reproducibility of the calculated measures are however ultimately dependent on the accuracy of each of the earlier steps of data acquisition and segmentation.

The reproducibility and accuracy of the calculated parameters are potentially compromised by the current need for the tuning of several parameters by visual feedback. The main issue is in the combination of manual control of contrast and brightness when acquiring the images and manual tuning of optimal segmentation parameters. The manual control of contrast and brightness means that different users will acquire image data with varying levels of intensity. The varying intensities could cause variations in segmentation parameter tuning. This is because the human visual system is more sensitive to some intensity levels than others. Even with identical intensity levels different human operators will tune parameters differently. Additionally, the need for human decisions in the analysis introduces the possibility of human error.

The current need for human interaction in the analysis process introduces bias

that is detrimental for the overall reproducibility. This can be seen as a further motivation for the complete automation of the analysis process. The challenges will be to introduce standardized intensity settings combined with automatic parameter tuning for the segmentation method.

These reproducibility issues do not mean that the accuracy of the calculation methods is not important since one source of error certainly is better than two. It does however mean that the human factor must be taken into account when assessing any calculated parameter. The different parameter's sensitivity to the errors should be further analyzed.

Representative microstructure

A different issue is that of representative microstructure. A typical microstructure cuboid has dimensions that only cover a small fraction of the entire electrode. A valid question is whether the analyzed volume is representative for the entire electrode. The microstructure might vary between different ends of the cell and it might also vary locally. An in depth study of 3D microstructure variation would be interesting to be able to investigate how much microstructure varies. This study would be interesting both for the sake of the relevance of locally calculated parameters but also for SOC development in general.

Chapter 9

Conclusion

The main objective of this work has been to develop automated methods for quantitative 3D characterization of microstructures. The main results of the work can be divided into two main parts: data processing and parameter extraction.

Data processing

- Automatic methods for aligning the image data into a 3D rectangular cuboid was described and applied to the alignment of SOFC microstructures. The methodology is based on detecting reference marks in each image and aligning each image to these invariant image features. The alignment is performed in two passes such that information about neighboring slices can be used to increase detection accuracy and robustness. The methodology allows identification of changes in sectioning spacing.
- An automatic method for non-uniform intensity correction of the aligned image stack was developed. The methodology fits a 3D intensity variation field to randomly sampled voxels from the image stack. The fitting is performed iteratively by removing outliers and refitting to the remaining samples. The method was shown to work well for SOFC microstructure data.

- A level set segmentation method was developed and tested on SOFC samples. The technique is based on an iterative method that uses vector fields derived from the image data to drive the evolution of the segmentation phase boundaries. The method was shown to have a number of desirable properties for microstructure segmentation: sub-voxel accuracy representation of phase boundaries, three dimensionally coherent segmented structures and integrated surface smoothing.

Parameter extraction

- A method for accurate calculation of surface areas and TPBs was developed. The method extracts the TPBs as connected curve loops without branches. The polygonization routine includes a technique for merging a double representation of an implicit surface into one. The polygonization results in a data structure that easily facilitates the combination of mesh data and voxel data. The method was shown to have high accuracy by validating the calculated values by comparisons to exact theoretical values of geometrical primitives.
- Several methods for characterizing 3D network structures were developed. The calculated parameter distributions describe network shape and connectivity characteristics. The methods work without any assumptions of microstructure shape and structure and describe aspects of: tortuosity, pathway radius, cavity/agglomerate size and degree of dead ended network branches.

Bibliography

[1] Jørgensen, P. S., Hansen, K. V., Larsen, R., Bowen, J. R. A framework for automatic segmentation in three dimensions of microstructural tomography data. *Ultramicroscopy*, **110**, 216–228, 2010.

[2] Lutz, A. E., Larson, R. S., Keller, J. O. Thermodynamic comparison of fuel cells to the Carnot cycle. *International Journal of Hydrogen Energy*, **27**, 1103–1111, 2002.

[3] Minh, N. Q. Ceramic fuel cells. *Journal of the American Ceramic Society*, **76**, 563–588, 1993.

[4] Singhal, S. C. Advances in solid oxide fuel cell technology. *Solid State Ionics*, **135**, 305–313, 2000.

[5] Stambouli, A. B., Traversa, E. Solid oxide fuel cells (SOFCs): a review of an environmentally clean and efficient source of energy. *Renewable and Sustainable Energy Reviews*, **6**, 433–455, 2002.

[6] Kapdan, I. K., Kargi, F. Bio-hydrogen production from waste materials. *Enzyme and microbial technology*, **38**, 569–582, 2006.

[7] Ni, M., Leung, M. K. H., Leung, D. Y. C. Technological development of hydrogen production by solid oxide electrolyzer cell (SOEC). *International Journal of Hydrogen Energy*, **33**, 2337–2354, 2008.

[8] Turner, J. A. A realizable renewable energy future. *Science*, **285**, 687–689, 1999.

[9] Leah, R. T., Brandon, N. P., Aguiar, P. Modelling of cells, stacks and systems based around metal-supported planar IT-SOFC cells with CGO

electrolytes operating at 500-600 °C. *Journal of Power Sources*, **145**, 336–352, 2005.

[10] O'Hayre, R., Cha, S. W., Colella, W., Prinz, F. B. *Fuel cell fundamentals*. Wiley New York, 2005.

[11] Wang, Z., Qian, J., Cao, J., Wang, S., Wen, T. A study of multilayer tape casting method for anode-supported planar type solid oxide fuel cells (SOFCs). *Journal of Alloys and Compounds*, **437**, 264–268, 2007.

[12] Müller, A. C., Herbstritt, D., Ivers-Tiffée, E. Development of a multilayer anode for solid oxide fuel cells. *Solid State Ionics*, **152**, 537–542, 2002.

[13] Dollen, P. V., Barnett, S. A study of screen printed yttria-stabilized zirconia layers for solid oxide fuel cells. *Journal of the American Ceramic Society*, **88**, 3361–3368, 2005.

[14] Itoh, H., Yamamoto, T., Mori, M., Horita, T., Sakai, N., Yokokawa, H., Dokiya, M. Configurational and Electrical Behavior of Ni-YSZ Cermet with Novel Microstructure for Solid Oxide Fuel Cell Anodes. *Journal of the Electrochemical Society*, **144**, 641–646, 1997.

[15] Iwata, T. Characterization of Ni-YSZ Anode Degradation for Substrate-Type Solid Oxide Fuel Cells. *Journal of the Electrochemical Society*, **143**, 1521–1525, 1996.

[16] Jiang, S. P., Badwal, S. P. S. An electrode kinetics study of H_2 oxidation on Ni/Y_2O_3-ZrO_2 cermet electrode of the solid oxide fuel cell. *Solid State Ionics*, **123**, 209–224, 1999.

[17] Haanappel, V. A. C., Mertens, J., Rutenbeck, D., Tropartz, C., Herzhof, W., Sebold, D., Tietz, F. Optimisation of processing and microstructural parameters of LSM cathodes to improve the electrochemical performance of anode-supported SOFCs. *Journal of Power Sources*, **141**, 216–226, 2005.

[18] Kawada, T., Sakai, N., Yokokawa, H., Dokiya, M., Mori, M., Iwata, T. Characteristics of Slurry-Coated Nickel Zirconia Cermet Anodes for Solid Oxide Fuel Cells. *Journal of the Electrochemical Society*, **137**, 3042–3047, 1990.

[19] Lee, J. H., Heo, J. W., Lee, D. S., Kim, J., Kim, G. H., Lee, H. W., Song, H. S., Moon, J. H. The impact of anode microstructure on the power generating characteristics of SOFC. *Solid State Ionics*, **158**, 225–232, 2003.

[20] Spowart, J. E. Automated serial sectioning for 3-D analysis of microstructures. *Scripta Materialia*, **55**, 5–10, 2006.

[21] Alkemper, J., Voorhees, P. W. Quantitative serial sectioning analysis. *Journal of microscopy*, **201**, 388–394, 2001.

[22] Kral, M. V., Mangan, M. A., Spanos, G., Rosenberg, R. O. Three-dimensional analysis of microstructures. *Materials Characterization*, **45**, 17–23, 2000.

[23] Stampanoni, M., Borchert, G., Wyss, P., Abela, R., Patterson, B., Hunt, S., Vermeulen, D., Rüegsegger, P. High resolution X-ray detector for synchrotron-based microtomography. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, **491**, 291–301, 2002.

[24] Donoghue, P. C. J., Bengtson, S., Dong, X., Gostling, N. J., Huldtgren, T., Cunningham, J. A., Yin, C., Yue, Z., Peng, F., Stampanoni, M. Synchrotron X-ray tomographic microscopy of fossil embryos. *Nature*, **442**, 680–683, 2006.

[25] Izzo, J. R., Joshi, A. S., Grew, K. N., Chiu, W. K. S., Tkachuk, A., Wang, S. H., Yun, W. Nondestructive reconstruction and analysis of SOFC anodes using X-ray computed tomography at sub-50 nm resolution. *Journal of the Electrochemical Society*, **155**, B504–B508, 2008.

[26] Weyland, M. Electron tomography of catalysts. *Topics in Catalysis*, **21**, 175–183, 2002.

[27] Velichko, A., Holzapfel, C., Mücklich, F. 3D characterization of graphite morphologies in cast iron. *Advanced Engineering Materials*, **9**, 39–45, 2007.

[28] Seliger, R. L., Fleming, W. P. Focused ion beams in microfabrication. *Journal of Applied Physics*, **45**, 1416–1422, 1974.

[29] Krohn, V. E., Ringo, G. R. Ion source of high brightness using liquid metal. *Applied Physics Letters*, **27**, 479–481, 1975.

[30] Melngailis, J. Focused ion beam technology and applications. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, **5**, 469–495, 1987.

[31] Kirk, E. C. G., Williams, D. A., Ahmed, H. Cross-sectional transmission electron microscopy of precisely selected regions from semiconductor devices. *Institute of Physics Conference Series*, **100**, 501–506, 1989.

[32] Overwijk, M. H. F., Van den Heuvel, F. C., Bulle-Lieuwma, C. W. T. Novel scheme for the preparation of transmission electron microscopy specimens with a focused ion beam. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, **11**, 2021–2024, 1993.

[33] Giannuzzi, L. A., Stevie, F. A. A review of focused ion beam milling techniques for TEM specimen preparation. *Micron*, **30**, 197–204, 1999.

[34] Sugiyama, M., Sigesato, G. A review of focused ion beam technology and its applications in transmission electron microscopy. *Journal of electron microscopy*, **53**, 527–536, 2004.

[35] Kometani, R., Hoshino, T., Kanda, K., Haruyama, Y., Kaito, T., Fujita, J., Ishida, M., Ochiai, Y., Matsui, S. Three-dimensional high-performance nano-tools fabricated using focused-ion-beam chemical-vapor-deposition. *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, **232**, 362–366, 2005.

[36] Akita, S., Nakayama, Y., Mizooka, S., Takano, Y., Okawa, T., Miyatake, Y., Yamanaka, S., Tsuji, M., Nosaka, T. Nanotweezers consisting of carbon nanotubes operating in an atomic force microscope. *Applied Physics Letters*, **79**, 1691–1693, 2001.

[37] Groeber, M. A., Haley, B. K., Uchic, M. D., Dimiduk, D. M., Ghosh, S. 3D reconstruction and characterization of polycrystalline microstructures using a FIB-SEM system. *Materials Characterization*, **57**, 259–273, 2006.

[38] Inkson, B. J., Steer, T., Mobus, G., Wagner, T. Subsurface nanoindentation deformation of Cu-Al multilayers mapped in 3D by focused ion beam microscopy. *Journal of microscopy*, **201**, 256–269, 2001.

[39] Inkson, B. J., Mulvihill, M., Möbus, G. 3D determination of grain shape in a FeAl-based nanocomposite by 3D FIB tomography. *Scripta Materialia*, **45**, 753–758, 2001.

[40] Holzer, L., Indutnyi, F., Gasser, P. H., Münch, B., Wegmann, M. Three-dimensional analysis of porous $BaTiO_3$ ceramics using FIB nanotomography. *Journal of microscopy*, **216**, 84–95, 2004.

[41] Münch, B., Holzer, L. Contradicting Geometrical Concepts in Pore Size Analysis Attained with Electron Microscopy and Mercury Intrusion. *Journal of the American Ceramic Society*, **91**, 4059–4067, 2008.

[42] Holzer, L., Münch, B., Wegmann, M., Gasser, P., Flatt, R. J. FIB-nanotomography of particulate systems-Part I: Particle shape and topology of interfaces. *Journal of the American Ceramic Society*, **89**, 2577–2585, 2006.

[43] Münch, B., Gasser, P., Holzer, L., Flatt, R. FIB-nanotomography of particulate systems-Part II: Particle recognition and effect of boundary truncation. *Journal of the American Ceramic Society*, **89**, 2586–2595, 2006.

[44] Holzer, L., Münch, B. Toward Reproducible Three-Dimensional Microstructure Analysis of Granular Materials and Complex Suspensions. *Microscopy and Microanalysis*, **15**, 130–146, 2009.

[45] Lasagni, F., Lasagni, A., Engstler, M., Degischer, H. P., Mücklich, F. Nano-characterization of Cast Structures by FIB-Tomography. *Advanced Engineering Materials*, **10**, 62–66, 2008.

[46] Wilson, J. R., Kobsiriphat, W., Mendoza, R., Chen, H.-Y., Hiller, J. M., Miller, D. J., Thornton, K., Voorhees, P. W., Adler, S. B., Barnett, S. A. Three-dimensional reconstruction of a solid-oxide fuel-cell anode. *Nature materials*, **5**, 541–544, 2006.

[47] Wilson, J. R., Gameiro, M., Mischaikow, K., Kalies, W., Voorhees, P. W., Barnett, S. A. Three-Dimensional Analysis of Solid Oxide Fuel Cell Ni-YSZ Anode Interconnectivity. *Microscopy & Microanalysis*, **15**, 71–77, 2009.

[48] Wilson, J. R., Duong, A. T., Gameiro, M., Chen, H.-Y., Thornton, K., Mumm, D. R., Barnett, S. A. Quantitative three-dimensional microstructure of a solid oxide fuel cell cathode. *Electrochemistry Communications*, **11**, 1052–1056, 2009.

[49] Gostovic, D., Smith, J. R., Kundinger, D. P., Jones, K. S., Wachsman, E. D. Three-dimensional reconstruction of porous LSCF cathodes. *Electrochemical and Solid-State Letters*, **10**, B214–B217, 2007.

[50] Shearing, P. R., Golbert, J., Chater, R. J., Brandon, N. P. 3D reconstruction of SOFC anodes using a focused ion beam lift-out technique. *Chemical Engineering Science*, **64**, 3928–3933, 2009.

[51] Iwai, H., Shikazono, N., Matsui, T., Teshima, H., Kishimoto, M., Kishida, R., Hayashi, D., Matsuzaki, K., Kanno, D., Saito, M., Muroyama, H., Eguchi, K., Kasagi, N., Yoshida, H. Quantification of SOFC anode microstructure based on dual beam FIB-SEM technique. *Journal of Power Sources*, **195**, 955–961, 2010.

[52] Kenney, B., Valdmanis, M., Baker, C., Pharoah, J. G., Karan, K. Computation of TPB length, surface area and pore size from numerical reconstruction of composite solid oxide fuel cell electrodes. *Journal of Power Sources*, **189**, 1051–1059, 2009.

[53] Wilson, J. R., Cronin, J. S., Duong, A. T., Rukes, S., Chen, H. Y., Thornton, K., Mumm, D. R., Barnett, S. Effect of composition of ($La_{0.8}Sr_{0.2}MnO_3$-Y_2O_3-stabilized ZrO_2) cathodes: Correlating three-dimensional microstructure and polarization resistance. *Journal of Power Sources*, **195**, 1829–1840, 2010.

[54] Goldstein, J., Newbury, D. E., Echlin, P., Lyman, C. E., Joy, D. C., Romig, A. D., Lyman, C. E., Fiori, C., Lifshin, E. *Scanning electron microscopy and X-ray microanalysis*. Plenum Press, 1992.

[55] Amelinckx, S., Dyck, D. V., Landuyt, J. V., Tendeloo, G. V. *Electron microscopy: principles and fundamentals*. Wiley-VCH, 1997.

[56] Giannuzzi, L. A., Stevie, F. A. *Introduction to focused ion beams: instrumentation, theory, techniques, and practice*. Springer Verlag, 2005.

[57] Mangan, M. A., Shiflet, G. J. Three dimensional investigation of Cu-Ti discontinuous precipitation. *Scripta Materialia*, **37**, 517–522, 1997.

[58] Mangan, M. A., Lauren, P. D., Shiflet, G. J. Three-dimensional reconstruction of Widmanstätten plates in Fe-12·3Mn-0·8C. *Journal of microscopy*, **188**, 36–41, 1997.

[59] Sutton, M. A., Cheng, M., Peters, W. H., Chao, Y. J., McNeill, S. R. Application of an optimized digital correlation method to planar deformation analysis. *Image and Vision Computing*, **4**, 143–150, 1986.

[60] Schreier, H. W., Braasch, J. R., Sutton, M. A. Systematic errors in digital image correlation caused by intensity interpolation. *Optical Engineering*, **39**, 2915–2921, 2000.

[61] Kubis, A. J., Shiflet, G. J., Hull, R., Dunn, D. N. Focused ion-beam tomography. *Metallurgical and Materials Transactions A*, **35**, 1935–1943, 2004.

[62] Duda, R. O., Hart, P. E. Use of the Hough transformation to detect lines and curves in pictures. *Communications of the Association of Computing Machinery*, **15**, 11–15, 1972.

[63] Lassen, N. C. K. Automatic high-precision measurements of the location and width of Kikuchi bands in electron backscatter diffraction patterns. *Journal of microscopy*, **190**, 375–391, 1998.

[64] Schaffer, M., Wagner, J. Block lift-out sample preparation for 3D experiments in a dual beam focused ion beam microscope. *Microchimica Acta*, **161**, 421–425, 2008.

[65] Jørgensen, P. S., Larsen, R., Wraae, K. Unsupervised Assessment of Subcutaneous and Visceral Fat by MRI. *Lecture Notes in Computer Science*, **5575**, 179–188, 2009.

[66] Phaneuf, M. W. Applications of focused ion beam microscopy to materials science specimens. *Micron*, **30**, 277–288, 1999.

[67] Langford, R. M., Dale, G., Hopkins, P. J., Ewen, P. J. S., Petford-Long, A. K. Focused ion beam micromachining of three-dimensional structures and three-dimensional reconstruction to assess their shape. *Journal of Micromechanics and Microengineering*, **12**, 111–114, 2002.

[68] Ma, L. W., Cairney, J. M., McGrouther, D., Hoffman, M., Munroe, P. R. Three dimensional imaging of deformation modes in TiN-based thin film coatings. *Thin Solid Films*, **515**, 3190–3195, 2007.

[69] Uchic, M. D., Holzer, L., Inkson, B. J., Principe, E. L., Munroe, P. Three-Dimensional Microstructural Characterization Using Focus Ion Beam Tomography. *MRS Bulletin*, **32**, 408–416, 2007.

[70] Ribeiro, E., Shah, M. Computer Vision for Nanoscale Imaging. *Machine Vision and Applications*, **17**, 147–162, 2006.

[71] Simmons, J. P., Chuang, P., Comer, M., Spowart, J. E., Uchic, M. D., Graef, M. D. Application and further development of advanced image processing algorithms for automated analysis of serial section image data. *Modelling and Simulation in Materials Science and Engineering*, **17**, 025002 (22pp), 2009.

[72] Thydén, K. *Microstructural Degradation of Ni-YSZ Anodes for Solid Oxide Fuel Cells*. Ph.D. thesis, Niels Bohr Institute, University of Copenhagen, DK, 2008.

[73] Golbert, J., Adjiman, C. S., Brandon, N. P. Microstructural modeling of solid oxide fuel cell anodes. *Industrial & Engineering Chemistry Research*, **47**, 7693–7699, 2008.

[74] Russ, J. R. *Image Processing Handbook*. CRC Press, 2006.

[75] Kass, M., Witkin, A., Terzopoulos, D. Snakes: Active Contour Models. *International Journal of Computer Vision*, **1**, 321–331, 1998.

[76] Osher, S., Sethian, J. Fronts Propagating with Curvature dependent Speed: Algorithms Based on Hamilton-Jacobi Formulations. *Journal of Computational Physics*, **79**, 12–49, 1988.

[77] Baillard, C., Hellier, P., Barillot, C. Segmentation of brain 3D MR images using level sets and dense registration. *Medical image analysis*, **5**, 185–194, 2001.

[78] Osher, S., Fedkiw, R. *Level Set Methods and Dynamic Implicit Surfaces*. Springer, 2003.

[79] Bagger, C., Linderoth, S., Mogensen, M., Hendriksen, P. V., Kindl, B., Primdahl, S., Larsen, P. H., Poulsen, F. W., Bonanos, N., Jørgensen, M. J. Status of the danish Solid Oxide Fuel Cell R&D. Singhal, S. C., Dokia, M. (eds.), *Proceedings of Sixth International Symposium on Solid Oxide Fuel Cells (SOFC-VI)*, pp. 28–35, The Electrochemical Society, 1999.

[80] Jørgensen, P. S., Bowen, J. R. Automatic Quantitative Image Analysis of Micrographs. Mogensen, M. (ed.), *European fuel cell forum 2008*, Lucerne, CH, A0308, July 3, p. 8, 2008.

[81] Tietz, F., Mai, A., Stöver, D. From powder properties to fuel cell performance - A holistic approach for SOFC cathode development. *Solid State Ionics*, **179**, 1509–1515, 2008.

[82] Adams, R., Bischof, L. Seeded region growing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **16**, 641–647, 1994.

[83] Lorensen, W. E., Cline, H. E. Marching cubes: A high resolution 3D surface construction algorithm. Stone, M. C. (ed.), *SIGGRAPH '87: Proceedings of the 14th annual conference on Computer graphics and interactive techniques*, New York, USA, July, vol. 1, pp. 163–169, ACM, 1987.

[84] Krissian, K., Westin, C.-F. Fast Sub-Voxel Reinitialization of the Distance Map for Level Set Methods. *Pattern Recognition Letters*, **26**, 1532–1542, 2005.

[85] Zhang, S., Lynch, M., Gokhale, A. M., Liu, M. Unbiased characterization of three-phase microstructure of porous lanthanum doped strontium manganite/yttria-stabilized zirconia composite cathodes for solid oxide fuel cells using atomic force microscopy and stereology. *Journal of Power Sources*, **192**, 367–371, 2009.

[86] Newman, T. S., Yi, H. A survey of the marching cubes algorithm. *Computers & Graphics*, **30**, 854–879, 2006.

[87] Wang, C. C. L. Direct extraction of surface meshes from implicitly represented heterogeneous volumes. *Computer-Aided Design*, **39**, 35–50, 2007.

[88] Bloomenthal, J., Wyvill, B. *Introduction to Implicit Surfaces*. Morgan Kaufmann Publishers Inc., 1997.

[89] Warren, J., Weimer, H. *Subdivision Methods for Geometric Design*. Morgan Kaufmann Publishers Inc., 2001.

[90] Catmull, E., Clark, J. Recursively generated B-spline surfaces on arbitrary topological meshes. *Computer-Aided Design*, **10**, 350–355, 1978.

[91] Al-Raoush, R. I., Willson, C. S. Extraction of physically realistic pore network properties from three-dimensional synchrotron X-ray microtomography images of unconsolidated porous media systems. *Journal of hydrology*, **300**, 44–64, 2005.

[92] Al-Raoush, R. I. Microstructure characterization of granular materials. *Physica A: Statistical Mechanics and its Applications*, **377**, 545–558, 2007.

[93] Sethian, J. A. A Fast Marching Level Set Method for Monotonically Advancing Fronts. *Proceedings of the National Academy of Sciences of the United States of America*, **93**, 1591–1595, 1996.

[94] Sethian, J. A. Fast Marching Methods. *SIAM Review*, **41**, 199–235, 1999.

[95] Popovici, A. M., Sethian, J. A. Three-dimensional travel-time computation using the Fast marching method. *Mathematical Methods in Geophysical Imaging V*, **3453**, 82–93, 1998.

[96] Williams, J. W. J. Algorithm 232: heapsort. *Communications of the ACM*, **7**, 347–348, 1964.

[97] Shearing, P. R., Howard, L. E., Jørgensen, P. S., Brandon, N. P., Harris, S. J. Characterization of the 3-dimensional microstructure of a graphite negative electrode from a Li-ion battery. *Electrochemistry Communications*, **12**, 374–377, 2010.

[98] Blum, H. A transformation for extracting new descriptors of shape. *Models for the perception of speech and visual form*, **19**, 362–380, 1967.

[99] Delerue, J. F., Perrier, E., Yu, Z. Y., Velde, B. New algorithms in 3D image analysis and their application to the measurement of a spatialized pore size distribution in soils. *Physics and Chemistry of the Earth, Part A*, **24**, 639–644, 1999.

[100] Pudney, C. Distance-ordered homotopic thinning: a skeletonization algorithm for 3D digital images. *Computer Vision and Image Understanding*, **72**, 404–413, 1998.

[101] Søgaard, M., Mortensen, J., Jacobsen, T. Impedance Modeling of Mixed Ionic Electronic Conducting Cathodes for Solid Oxide Fuel Cells. *217th ECS Meeting (Accepted)*, Vancouver, CA, Symposium B11, Abstract 668, April 26, 2010, p. 1, 2010.

[102] Adler, S. B., Lane, J. A., Steele, B. C. H. Electrode Kinetics of Porous Mixed-Conducting Oxygen Electrodes. *Journal of the Electrochemical Society*, **143**, 3554–3564, 1996.

[103] Barsoukov, E., Macdonald, J. R. *Impedance spectroscopy: theory, experiment, and applications*. John Wiley and Sons, 2005.