



HAL
open science

Kaleidoscopic Imaging

Ilya Reshetouski

► **To cite this version:**

Ilya Reshetouski. Kaleidoscopic Imaging. Graphics [cs.GR]. Faculty of Natural Sciences and Technology I of Saarland University, 2014. English. NNT: . tel-01091397

HAL Id: tel-01091397

<https://inria.hal.science/tel-01091397>

Submitted on 5 Dec 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Kaleidoscopic Imaging

Ilya Reshetouski

Thesis for obtaining the title of
Doctor of Engineering Science (Dr.-Ing.)
of the Faculty of Natural Sciences and Technology I
of Saarland University

Saarbrücken
2014

Dekan – Dean

Prof. Dr. Markus Bläser

Betreuender Hochschullehrer – Supervisor

Dr. habil. Ivo Ihrke

Gutachter – Reviewers

Dr. habil. Ivo Ihrke

Prof. Dr. Hans-Peter Seidel

Prof. Dr. Martin Vetterli

Kolloquium – Defense

Datum – Date:

November 6, 2013, in Saarbrücken

Vorsitzender – Head of Colloquium:

Prof. Dr. Christoph Weidenbach	MPI Informatik	Saarbrücken, Germany
--------------------------------	----------------	----------------------

Prüfer – Examiners:

Dr. habil. Ivo Ihrke

Prof. Dr. Hans-Peter Seidel

Prof. Dr. Martin Vetterli

Protokoll – Reporter:

Dr. Levi Valgaerts

INRIA Sud-Ouest

Bordeaux, France

MPI Informatik

Saarbrücken, Germany

EPFL

Lausanne, Switzerland

MPI Informatik

Saarbrücken, Germany

Abstract

Kaleidoscopes have a great potential in computational photography as a tool for redistributing light rays. In time-of-flight imaging the concept of the kaleidoscope is also useful when dealing with the reconstruction of the geometry that causes multiple reflections. This work is a step towards opening new possibilities for the use of mirror systems as well as towards making their use more practical. The focus of this work is the analysis of planar kaleidoscope systems to enable their practical applicability in *3D* imaging tasks.

We analyse important practical properties of mirror systems and develop a theoretical toolbox for dealing with planar kaleidoscopes. Based on this theoretical toolbox we explore the use of planar kaleidoscopes for multi-view imaging and for the acquisition of *3D* objects. The knowledge of the mirrors positions is crucial for these multi-view applications. On the other hand, the reconstruction of the geometry of a mirror room from time-of-flight measurements is also an important problem. We therefore employ the developed tools for solving this problem using multiple observations of a single scene point.

Kurzfassung

Kaleidoskope haben in der rechnergestützten Fotografie ein großes Anwendungspotenzial, da sie flexibel zur Umverteilung von Lichtstrahlen genutzt werden können. Diese Arbeit ist ein Schritt auf dem Weg zu neuen Einsatzmöglichkeiten von Spiegelsystemen und zu ihrer praktischen Anwendung. Das Hauptaugenmerk der Arbeit liegt dabei auf der Analyse planarer Spiegelsysteme mit dem Ziel, sie für Aufgaben in der 3D-Bilderzeugung praktisch nutzbar zu machen. Auch für die Time-of-flight-Technologie ist das Konzept des Kaleidoskops, wie in der Arbeit gezeigt wird, bei der Rekonstruktion von Mehrfachreflektionen erzeugender Geometrie von Nutzen.

In der Arbeit wird ein theoretischer Ansatz entwickelt der die Analyse planarer Kaleidoskope stark vereinfacht. Mithilfe dieses Ansatzes wird der Einsatz planarer Spiegelsysteme im Multiview Imaging und bei der Erfassung von 3-D-Objekten untersucht. Das Wissen um die Spiegelpositionen innerhalb des Systems ist für diese Anwendungen entscheidend und erfordert die Entwicklung geeigneter Methoden zur Kalibrierung dieser Positionen. Ein ähnliches Problem tritt in Time-of-Flight Anwendungen bei der, oft unerwünschten, Aufnahme von Mehrfachreflektionen auf. Beide Problemstellungen lassen sich auf die Rekonstruktion der Geometrie eines Spiegelraums zurückführen, das mit Hilfe des entwickelten Ansatzes in allgemeinerer Weise als bisher gelöst werden kann.

Summary

Mirrors have been used in a number of vision applications in the past. Examples using curved mirror surfaces include catadioptric imaging, reflectance, and texture measurement systems. Planar mirror systems have been used for multi-view imaging of flat and extended depth samples and for confocal imaging. Most of these systems are designed to work with single bounce reflections. On the other hand, multiple bounce systems are rarely used even though, as we show in this thesis, they have a great potential, especially for multi-view imaging applications.

The initial simplicity of the reflection operator converts into something non-trivial when multiple reflections are considered. Even for configurations with only planar mirrors, the light propagation map is hard to predict in most of the cases. This complicates the design of mirror systems with required properties. Another complicating factor is the laboriousness of geometrical calibration routines for such systems.

The target of the thesis is to open new horizons for the use of multiple-bounce planar mirror systems, known as kaleidoscopes, in computer vision and other related areas by enabling a better understanding of the properties of such systems and by providing practical ways for the interpretation of mirror interactions with light. In addition, we develop a flexible method to recover the structure of planar mirror systems which can be, e.g., used for an automatic calibration of kaleidoscopes.

Theory of Mirror Systems Part II of the thesis is concerned with theoretical aspects of planar mirror systems. We introduce the concept of a "space partitioning" to describe the behavior of idealized light rays in multiple-bounce planar mirror systems.

This concept allows us to derive a novel continuity result for the modification of an observer position in a planar mirror system which, in turn, enables viewpoint sampling, e.g., in order to optimize the observer position. Another application of the observer continuity is a simplification of the mirror system interpretation when dealing with non-central camera/projector

devices.

Labeling of Kaleidoscopic Images Part III of the thesis presents the solution of the *labeling* problem. Planar kaleidoscopes are capable of generating multiple views of the same scene in a single camera image. For a successful interpretation of such images it is important to distinguish different viewpoints, which may overlap in the recorded image. Image labeling is the assignment of viewpoints to image pixels. Successful labeling enables the extraction of single view point images from the recorded multiple-bounce mirror image, opening the way for the application of standard multi-view computer vision algorithms.

Additionally, we show how the labeling can be produced for projectors. That enables controlled illumination of an object inside a kaleidoscope from multiple directions by a single projector. Thus, a camera-projector pair inside a multiple-bounce mirror system can be successfully used for reflectance acquisition.

Recovering the Structure of Planar Mirror Systems Most of the vision applications involving mirrors require the geometrical calibration of the mirrors. In the last part of the thesis we address this problem and develop an approach for recovering the structure of a planar mirror system from multiple observations of a single scene point.

By formulating the problem in terms of apparent distance of the observed point via multiply reflected paths, our method applies to time-of-flight measurement systems like RADAR, active SONAR, and LIDAR, but also to acoustics where room geometry is to be inferred from the travel time of a pulse emitted by a speaker.

Our approach enables the recovery of the geometry of a mirror room from such data in more general settings than previously available techniques.

The work described in this thesis has been published at different international conferences. The core of Chapter 2 is published in [Reshetouski13a]. Chapter 5 is based on the publications [Reshetouski11, Ihrke12a, Klehm12]. Section 6.3 has its origins in [Ihrke12b], while Chapter 7 describes material from [Reshetouski13b]. The theoretical Part II is under submission.

Zusammenfassung

Spiegel werden für eine Reihe von Bildverarbeitungsanwendungen genutzt. Beispiele für den Gebrauch gekrümmter Spiegeloberflächen sind die kadioptrische Bilderzeugung, und die Erfassung von Reflektanz- und Texturparametern von strukturierten Oberflächen. Planare Spiegelsysteme werden auch für Multiview Imaging und zur konfokalen Bildgebung verwendet. Die Mehrheit dieser Systeme ist auf Einzelspiegelungen ausgelegt. Systeme in denen Mehrfachspiegelungen auftreten werden dagegen selten genutzt, obwohl sie – wie in dieser Arbeit aufgezeigt wird – ein großes Potenzial, insbesondere für Multiview Imaging-Anwendungen, besitzen.

Die anfängliche Einfachheit der Einfachreflektion wird bei der Betrachtung von Mehrfachspiegelungen zu etwas nicht trivialem. Die Lichtausbreitung ist in den meisten Fällen, selbst bei Systemen von lediglich planaren Spiegeln, nur schwer vorherzusehen. Dies erschwert den Entwurf von Spiegelsystemen mit gewünschten Eigenschaften. Ein weitere Schwierigkeit besteht hinsichtlich des Arbeitsaufwands der für die geometrische Kalibrierung solcher Systeme aufgewandt werden muss.

Das Ziel dieser Arbeit ist es, mittels eines verbesserten Verständnisses der Eigenschaften solcher Systeme, Hilfsmittel für die Interpretation von Spiegel-Licht-Interaktionen zu schaffen. Hierdurch ergeben sich neue Wege zur Nutzung von planaren Mehrfachreflexionssystemen – auch Kaleidoskope genannt – in der rechnergestützten Bildanalyse und verwandten Bereichen. Darüber hinaus wird eine flexible Methode zur automatischen Kalibrierung planarer Spiegelsysteme entwickelt.

Theorie planarer Spiegelsysteme Teil II der Arbeit befasst sich mit theoretischen Aspekten planarer Spiegelsysteme. Als Kernkonzept wird eine Zerlegung des virtuellen Spiegelraumes eingeführt um das Verhalten idealisierter Lichtstrahlen in planaren Spiegelsystemen, die durch einen perspektivischen Beobachter betrachtet werden, einfach und anschaulich beschreiben zu können.

Dieses Konzept erlaubt die Herleitung eines neuen Kontinuitätsresultats.

tats: Die Zerlegung des Spiegelraumes erfährt, bei Veränderung der Beobachterposition, eine gleichmässig stetige Änderung. Dieses Ergebnis ermöglicht, z.B., ein Sampling der Beobachterposition und somit deren globale Optimierung. Eine weitere Anwendung der Beobachterkontinuität ist eine Vereinfachung der Interpretation von planaren Spiegelsystemen bei nicht-zentralperspektivischen Kameras/Projektoren.

Die Kennzeichnung von kaleidoskopischen Bildern Teil III der Arbeit zeigt die Lösung des Kennzeichnungsproblems auf. Planare Kaleidoskope sind in der Lage, Aufnahmen derselben Szene aus verschiedenen Blickwinkeln in einem einzigen Kamerabild darzustellen. Für eine erfolgreiche Interpretation solcher Bilder ist die Unterscheidung der verschiedenen Blickpunkte wichtig da diese sich im aufgenommenen Bild überschneiden können. Das Kennzeichnungsproblem von kaleidoskopischen Bildern bezieht sich hierbei auf die Zuweisung von Blickpunkten auf Bildelemente (Pixel). Die erfolgreiche Kennzeichnung eines kaleidoskopischen Bildes erlaubt es einzelne Blickpunkt-Bilder aus dem aufgenommenen Mehrfachspiegelbild zu extrahieren, was den Grundstein für die Anwendbarkeit von bekannten Multiview-Algorithmen im Bereich Computer Vision legt.

Außerdem wird aufgezeigt wie die Kennzeichnung von Projektoren in einem kaleidoskopischen System durchgeführt werden kann. Dies erlaubt die kontrollierte Beleuchtung eines Objekts aus mehreren Richtungen mithilfe eines einzigen Projektors. So kann ein Kamera-Projektor-Paar innerhalb eines Mehrfachspiegelsystems erfolgreich zur Reflektanzerfassung von komplexen drei-dimensionalen Objekten genutzt werden.

Geometrierekonstruktion planarer Spiegelsysteme Die meisten Bildverarbeitungsanwendungen in Spiegelsystemen erfordern die geometrische Kalibrierung dieser Spiegel. Der letzte Teil der Arbeit beschäftigt sich daher mit diesem Problem. Es wird ein Ansatz zur Rekonstruktion der Geometrie planarer Spiegelsysteme mit Hilfe der Aufnahme eines einzigen Szenenpunktes demonstriert. Das Problem wird dabei auf die Messung des scheinbaren Abstands des beobachteten Punktes über verschiedene mehrfach reflektierte Pfade zurückgeführt und ist damit potentiell für Time-of-Flight-Messsysteme wie RADAR, aktives SONAR und LIDAR anwendbar. Eine zusätzliche Verbindung besteht zur Akustik, in der ein wichtiges Problem die Aufnahme einer Raumgeometrie, z.B. einer Konzerthalle, durch die Echomessung von Lautsprecherimpulsen ist. Der in der Arbeit vorgestellte Ansatz erlaubt eine allgemeinere Lösung des Problems als bislang verfügbare Techniken.

Die in dieser Doktorarbeit vorgestellten Arbeiten wurde schon im Rahmen mehrerer internationaler Konferenzen veröffentlicht. Die Kernaussa-

gen des zweiten Kapitels finden sich in [Reshetouski13a]. Kapitel 5 basiert auf den Veröffentlichungen [Reshetouski11, Ihrke12a, Klehm12]. Abschnitt 6.3 hat seinen Ursprung in [Ihrke12b], während Kapitel 7 Material aus [Reshetouski13b] beschreibt. Der theoretische Teil II wird zur Veröffentlichung eingereicht.

Acknowledgements

I am very grateful for the perfect research conditions I was provided with at Saarland University, MPI Informatik, and at INRIA Bordeaux. Especially, I would like to express my sincere gratitude to my supervisor, Ivo Ihrke, for the support and guidance in my research. Thanks also to my family, friends, colleagues for your help and fun.

Contents

I	Introduction	1
1	Kaleidoscopes in Multi-View and Time-of-Flight Imaging	3
1.1	Motivation	3
1.2	Thesis Outline	4
2	Related Work	5
2.1	Overview	5
2.2	Passive Imaging Systems Utilizing Planar Mirrors	6
2.2.1	Single-Mirror, Single-Bounce	7
2.2.2	Multi-Mirror, Single-Bounce per Mirror	8
2.2.3	Multi-Mirror, Multi-Bounce	10
2.3	Active Imaging Systems	10
2.3.1	Single-Mirror, Single-Bounce	11
2.3.2	Multi-Mirror, Single-Bounce per Mirror	11
2.3.3	Multi-Mirror, Multi-Bounce	12
2.4	Planar Mirror Calibration	12
2.5	Connection to Time-of-Flight Imaging and the Multi-Bounce Problem	13
2.6	Curved Mirrors	14
3	Basic Concepts, Experimental Setting, Calibration, and Pre-Processing	17
3.1	Unfolding - A Convenient Way for Interpreting Image Formation in Planar Mirror Systems	17
3.2	Background Concepts	23
3.2.1	The Visual Hull	23
3.2.2	BRDF	23
3.3	Experimental Setting, Calibration, and Pre-Processing	24
3.3.1	Geometric Calibration	25

3.3.2	Radiometric Calibration	26
II	Theory of Planar Mirror Systems	29
4	A Theory of Projective Imaging within Mirror Systems	31
4.1	Mirror Unfolding and Space Partitioning	31
4.1.1	Basic Definitions	31
4.1.2	Ray Unfolding	41
4.1.3	Space Partitioning	46
4.2	Continuity of The Space Partitioning	49
4.2.1	Continuity of the Space Partitioning with Respect to Changes in Observation Point	49
4.2.2	Observer-Continuity For Point Sets	58
	Conclusions	60
III	Three-Dimensional Kaleidoscopic Imaging	63
	Introduction	65
5	The Labeling Problem and its Solution	67
5.1	Planar Mirror System Containing Objects	67
5.2	Visual Hull as Sufficient Approximation of the Object for the Determination of the Labeling Function	71
5.3	Visual Hull Estimation	72
5.4	Applicability of the Kaleidoscopic Visual Hull for Image La- beling	74
5.5	Projector Image Labeling	75
5.6	Non-central Projection Camera and Projector Image Labeling	76
5.7	Segmentation	77
5.8	Experiments	77
5.8.1	Setup Details	77
5.8.2	Space Carving Implementation	77
5.8.3	Results	79
5.9	Limitations	81
6	Application Examples	83
6.1	Visual Hull Approximation	83
6.2	Multiview Stereo Reconstruction	84
6.3	Surround Geometry and Reflectance Acquisition	86
6.3.1	Operational Principle	87
6.3.2	Hardware design	89
6.3.3	Implementation	92

6.3.4 Results	93
Conclusions	97
IV Structure Reconstruction of Planar Mirror Systems	99
7 Discovering the Structure of a Planar Mirror System from Multiple Observations of a Single Point	101
7.1 Introduction	101
7.2 Problem Formulation	102
7.3 Overview	104
7.3.1 Problem Setting	105
7.3.2 Definitions	107
7.4 Verifying a Candidate Configuration	108
7.4.1 Joining Triplets	108
7.4.2 Conditions for Reconstructability	110
7.4.3 Verification Algorithm	110
7.5 Exhaustive Search Algorithm	111
7.5.1 Geometric Search Space Pruning	112
7.6 Randomized Search Algorithm	115
7.6.1 Extensions	116
7.7 Experimental Results	116
7.7.1 Simulation Results	116
7.7.2 Real World Example	118
7.8 Conclusions	119
Conclusions and Future work	121
Bibliography	125
Appendix	139
A Length of the perpendiculars in Theorem 2	141
B Practical Recommendations for Kaleidoscope Construction	143

PART I

Introduction

Kaleidoscopes in Multi-View and Time-of-Flight Imaging

1.1 Motivation

Mirrors traditionally have an important role in optics. In the last decades, mirror systems have found new applications in areas such as computational photography, where many proposed acquisition systems use mirrors as an instrument for shaping and rearranging camera/projector rays. Among the many kinds of mirror shapes that can be used in computational photography, planar kaleidoscopic mirror configurations play a special role. It was previously shown that a simple kaleidoscope can redistribute the camera/projector resolution into hundreds of sub-images, representing different points of view/illumination. This principle is very promising, since available resolution of today's cameras/projectors is large and still growing. However, to enable the practical use of such systems, several challenging problems need to be solved. In this thesis we address these problems and provide a step towards the ultimate goal of using a single device's available megapixels in the context of multiple-view imaging and acquisition. To advance towards this goal, a combination of three components is necessary.

The first component is the possibility to select a kaleidoscopic mirror configuration that generates a user-defined distribution of views of the given scene on a single camera sensor or/and on a single projector illumination image. Unfortunately, designing kaleidoscopes in order to achieve a desired view distribution is non-trivial. However, our theoretical results provide the

first step towards the possibility of sampling kaleidoscopes.

The second important component is the availability of practical methods to "unpack" and interpret captured data in case of a camera and to be able to illuminate the scene from required directions in case of a projector. In the thesis, this component relates to the solution of the labeling problem.

The third component of our multi-view approach using kaleidoscopes is geometric calibration of the mirrors in conjunction with a camera/projector. A similar problem arises in the area of Time-of-Flight imaging, where a room geometry is to be recovered from time-of-flight measurements. In these scenarios fiducial markers or Time-of-Flight devices allow us to automatically detect 3D positions of different observations of a common point inside the kaleidoscope. The ability to recover the structure of the kaleidoscope using such 3D observations only, would significantly simplify the practical use of the mirror systems and improve the related area in Time-of-Flight imaging. We discuss a solution to this problem in the $2\frac{1}{2}D$ case.

1.2 Thesis Outline

The structure of the thesis is as follows. In the next two chapters, Chapter 2 and Chapter 3 we discuss prior art, introduce the basic concepts underlying this work, and describe experimental prerequisites.

In the theoretical Part II (Chapter 4) we formalize and extend the mirror unfolding concept introduced in Chapter 3 and use this formalization to prove the observer continuity property of planar mirror systems.

The practical aspects of the use of kaleidoscopes for multi-view imaging and acquisition are considered in Part III. Kaleidoscopes allow us to observe the same scene from many different points of view by a single camera. One problem, which arises here is how to distinguish or, in other words, how to label different views, which may even partially overlap. Chapter 5 describes our solution to the labeling problem. Chapter 6 concludes Part III by demonstrating some of the applications, which become possible with our labeling solution.

Another practical aspect for the use of mirror systems is their calibration. This problem is rather complex in kaleidoscopic scenarios. In the last Part IV of the thesis we formulate the calibration problem as the problem of finding the structure of mirrors from multiple observations of a single point. In this formulation the problem is similar to problems in the field of time-of-flight imaging. We describe our solution for the $2\frac{1}{2}D$ case in Chapter 7.

Related Work

2.1 Overview

This chapter is dedicated to review the design and application of mirror systems in computer graphics and computer vision, as well as the related problem of the determination of the geometry of a mirror or mirror system. While less obvious, we point out a connection between mirror calibration or mirror shape estimation and time-of-flight imaging.

Our approach is based on a classification scheme for mirror systems, Fig. 2.1, that builds on the fundamental imaging properties of the employed mirror surfaces. We categorize existing systems into classes based on their mirroring properties and their use in active or passive imaging systems. The main categories for mirror systems are whether the mirrors are planar or curved, whether single or multiple mirrors are used and whether single-bounce or multi-bounce per mirror interaction is employed. However, in this thesis we focus only on planar mirror systems.

In the following sections we discuss passive imaging devices that utilize mirrors, Sect. 2.2. Passive systems have the property that light rays that cover a common scene point do not influence each other. On the other hand, if active illumination is introduced, light can superpose in a scene. We discuss active imaging systems in Sect. 2.3. All systems involving mirrors need to be calibrated, i.e. the geometry and position of the mirrors in the scene has to be determined. For this reason, we review related computer vision methods that aim at determining the shape of specular reflective surfaces or the position of a camera with respect to a known mirror geometry




Figure 2.1: Classification scheme. In our work we are dealing only with two classes: active and passive multiple planar mirrors systems with multiple bounces (indicated in bold in the picture and marked by red squares). Most of the methods developed in the thesis are, however, applicable to any planar mirror system (marked by blue squares).

in Sect. 2.4. Our calibration related approach, Part IV, deals with the recovery of a mirror system's geometry from depth measurements, which is a special case of the calibration problem. However, this problem has its own literature and approaches in the field of time-of-flight imaging and acoustics. We therefore draw connections between the previously discussed techniques and the time-of-flight literature in Sect. 2.5.

2.2 Passive Imaging Systems Utilizing Planar Mirrors

In this section we describe passive imaging devices that utilize planar mirrors in their design. The applications are mostly in stereo, multi-view and




Figure 2.2: Single mirror rectified catadioptric stereo camera [Gluckman00]. (left) Image formation draft. (right) Camera prototype (http://www.cs.columbia.edu/CAVE/projects/cad_stereo/).

panoramic imaging. The advantages of employing mirrors in a system are usually

- a reduction in system cost by utilizing less sensor hardware,
- a simplification of synchronization by compressing several views onto a single sensor, and
- homogeneous radiometric and colorimetric properties of the sensor hardware.

In utilizing these advantages, sensor resolution is usually traded off for an expanded view point coverage of a scene.

2.2.1 Single-Mirror, Single-Bounce

Single planar mirror systems are necessarily single-bounce. They can thus be used to generate two viewpoints in a single image. This feature is often used to produce inexpensive stereo viewers in a dual screen setup [Wu10] and many hobbyists make use of this capability¹.

Similarly, a stereo camera can be built with a single mirror [Gluckman98a] and commercial modifications of standard cameras are being offered². Depending on the mirror orientation with respect to the camera optical axes, the resulting epipolar geometry can be more or less suitable for stereo matching. Gluckman and Nayar [Gluckman00, Gluckman02] describe the conditions for epipolar lines to be parallel and along horizontal scan lines, a case that is particularly easy to handle in matching algorithms, see also Fig. 2.2.

¹http://klub.stereofotograf.eu/dual_monitor.php

²http://hineslab.com/old/Mirror_Stereo.html




Figure 2.3: A mirror array used for light field imaging [Sen05] (left). A fabricated mirror array with an optimized facet distribution [Fuchs12] (right).

In early work, Mitsumoto et al. [Mitsumoto92] describe object triangulation and geometric constraints for 3D reconstruction in case of a single plane mirror symmetry. They time-sequentially move the mirror to different positions and merge the reconstructions to obtain a larger coverage of the object.

Moving planar mirrors are also used to inexpensively generate many different viewpoints, e.g. for light field imaging [Ihrke08] or 3D reconstruction [Murray95, Hu09].

Beamsplitters are often employed to distribute a single view of a scene onto several imaging sensors [McGuire07]. These devices can be considered as a special case of a single mirroring operation for one of the sensors, whereas the beamsplitter appears transparent to the other.

2.2.2 Multi-Mirror, Single-Bounce per Mirror

An increase in complexity and achievable imaging geometry is obtained when introducing several planar mirrors [Gluckman00, Gluckman02]. Restrictions that guarantee a single bounce per mirror are a) that inter-reflections between mirrors are avoided, or b) that all camera rays only encounter mirroring sequences where each of the mirrors participates at most once.

No Inter-Reflections These arrangements are often employed for light field imaging with a single sensor [Levoy04, Sen05, Mukaigawa11, Fuchs12], see also Fig. 2.3. Since light field views differ only slightly from one another, mirror arrangements like the ones shown in the Figure can be suitably employed without too strong requirements on the positioning of the mirrors to avoid inter-reflections. Since views are usually supposed to cover a common viewing area, the carrier surface is chosen in a concave manner. If manufactured on a very small scale, faceted mirrors can be used




Figure 2.4: Design for a four mirror stereo camera or viewing device [Gluckman00] (left). The Sokkia MS27 commercial stereo viewer for aerial imagery (right).

to mimic bidirectional reflection distribution functions (BRDFs) with predefined properties [Weyrich09].

Another way to avoid inter-reflections is to position planar mirrors on a convex surface [Aggarwal01, Tan04] and is realized using pyramidal or truncated pyramid structures. This measure yields out-ward facing views for panoramic imaging [Tan04], or a means of performing aperture splitting of a single image onto several sensors [Aggarwal01], an application that is heavily used in computational photography applications.

In optics, in the area of multi-spectral imaging, especially manufactured mirrors, so called “image slicers” are being used to differently deflect the scan-lines of an image such that vertical sensor space is freed up for sensing spectrally expanded versions of the scan-lines that are obtained by passing them through a diffraction grating [Harvey05, Gao09, Gorman10].

Inter-Reflections with a Single Reflection per Mirror Several mirrors can also be arranged in a sequential sequence which yields a higher flexibility in generating virtual views and purely optical means of image manipulation. The most common commercial applications are probably erecting prisms in SLR viewfinders and other prism-based optical designs that are intended to flip or displace an image without distorting it otherwise [Smith08].

However, several planar mirrors are also used to obtain a higher degree of flexibility in the design of stereo imaging systems [Gluckman98a, Gluckman00, Gluckman02] or in the production of stereo viewing equipment as e.g. produced by Sokkia, see also Fig. 2.4.

In computational photography settings, beamsplitter trees are often employed to deliver a single physical image to different sensor units. The optical path towards each of those sensor units can be modified such that optically differently filtered images are recorded. For an overview of this area the interested reader is referred to [McGuire07, Wetzstein11, Zhou11].




Figure 2.5: The five-view case employing inter-reflections up to second order [Lanman07] – self-occlusion is clearly visible (left). In the case of many inter-reflections (middle) a pixel labeling procedure is necessary, see Chapter 5, that can resolve the view assignment to pixels – up to eight reflection levels have been employed (right).

2.2.3 Multi-Mirror, Multi-Bounce

Multi-bounce planar mirror systems are considerably more difficult to comprehend and to make use of. Early work in mirror-based single-image 3D reconstruction focused on setups consisting of two mirrors arranged such that their normals are in a common plane and that the angle between them is equal to $2\pi/N$. This has been a popular choice for three-dimensional imaging with a single camera with $N = 5$ views [Huang06, Forbes06, Lanman07, Lanman09]. It should be mentioned that this geometry results in a non-Coxeter structure, see Sect. 3.1, and therefore the camera position has to be suitably chosen to hide discontinuous views, see Part II. The multiple view geometry of this setting has been explored in [Ying10].

A common problem with this arrangement, and in fact with any multi-bounce system, is that the object position has to be chosen very carefully. The problem that occurs in the multi-bounce case is that an object might occlude its virtual counter parts, an effect that is easily observed when viewing one-self in a set of opened bathroom mirrors. A solution to this problem is presented later in this thesis, see Chapter 5, and consists in a pixel labeling procedure that determines for every pixel of an image with multiple inter-reflections which virtual view it belongs to, see also Fig. 2.5. This assignment can be computed from a single image and for arbitrary calibrated planar mirror geometries. Because of the kaleidoscopic nature of the resulting images, these systems are referred to as kaleidoscopic imaging systems.

2.3 Active Imaging Systems

Active imaging systems employ a light source in addition to an imaging device. Nowadays, these light sources are typically digital projectors which enable a per-pixel control of the illumination. The use of combinations of

cameras and projectors enables applications such as corrected projection onto curved surfaces, virtual large scale projection displays, 3D structured light scanning, reflectance scanning and more. An overview of the area of camera-projector systems is given in [Bimber08].

The combination of light sources with mirrors introduces additional problems in a measurement setting. Emitted light can superpose in a scene [Fasano03, Ihrke12b], defocusing problems [Baker98, Baker99, Zhang06] are exaggerated since projectors typically employ large apertures for light efficiency. On the other hand, active light helps in coding a scene, as e.g. in structured light scanning, or enables the scanning of surface properties.

Planar mirrors are most often used to multiply the number of physical projectors or to virtually position them in a physically impossible location.

2.3.1 Single-Mirror, Single-Bounce

The most common use of a single planar mirroring device is the use of a beamsplitter to bring a projector and a camera into a coaxial arrangement [Fujii05, Zhang06, Garg06, Ghosh07, Ghosh10, Holroyd10]. This configuration allows for illumination along the same rays that form the camera image and is often part of more complex active imaging systems.

In a different application, the use of a single planar mirror for range scanning inaccessible parts of an object has been reported [Fasano03]. To avoid the super-position of light, the operator has to manually ensure that the real and virtual laser lines are formed in distinct regions and that a distance heuristic can distinguish between the 3D points generated in the real space and in the virtual space, respectively.

2.3.2 Multi-Mirror, Single-Bounce per Mirror

In the active setting, systems of planar mirrors multiply a single projector into a set of virtual projectors, in effect realizing a large aperture projection system. These virtual large apertures have been employed in synthetic aperture confocal imaging techniques [Levoy04, Mukaigawa11] where the superposition of light is a crucial part of the functioning of the device. Confocal imaging systems can slice a volumetric scene via very shallow depth-of-field imaging and illumination. The planar mirrors are arranged tangent to a concave base shape [Levoy04] which is ellipsoidal in the case of [Mukaigawa11]. The mirror array is simultaneously used as a light field imaging unit, Sect. 2.2.2. The geometrical layout and interpretation are as discussed in Section 3.1.

Sequential folding of projection cones is often employed in rear-projection screens to reduce the size of the room that is required behind the screen.

Typically, large-scale front-surface mirrors are employed for this purpose³.

2.3.3 Multi-Mirror, Multi-Bounce

As mentioned in Sect. 2.2.2, the main complication in utilizing multiple ray bounces in a mirror system is that self-occlusion between the object and its virtual counter-parts has to be avoided. The simplest solution to this problem is the imaging of flat objects [Han03, Bangay04]. In [Han03], a kaleidoscopic mirror system was introduced that was capable of scanning the bidirectional texture function (BTF⁴) of a planar sample without moving the acquisition apparatus or the sample. In this case it is possible to observe a surface light field with a single picture and the sample can be illuminated from different directions by using a digital projector that is only highlighting from specific virtual directions. A sampling analysis of this type of system can be found in [Bangay04].

In this thesis we present an extension of planar sample kaleidoscopic reflectance scanning [Han03], which allows to scan extended depth objects, see Section 6.3. The solution is similar to the pixel labeling procedure, Chapter 5, Fig. 2.5 (right), this time applied to the projector coordinate system. If only pixels that have a unique label are illuminated simultaneously, the virtual illumination is guaranteed to come from a single direction without causing illumination overlap in the scene. We combine reflectance scanning with omnidirectional laser-range scanning.

The superposition of light can also be arranged such that a projected pattern perfectly super-positions onto itself. This approach requires orthogonal illumination with a direction that is contained in the plane spanned by the mirror normals. The two-mirror/five-virtual view system mentioned in Sect. 2.2.2 has been used for this purpose [Lanman07, Lanman09].

2.4 Planar Mirror Calibration

In order to successfully use planar mirror systems, they have to be calibrated. Usually this involves the estimation of the mirror position and orientation, potentially its shape (contours), and its radiometric properties [Reshetouski11, Ihrke12b].

Planar mirrors are relatively simple to calibrate since they do not introduce additional distortions into the image. Instead, the image taken by a perspective camera shows different perspective sub-views in parts of the acquired image. It is therefore only necessary to determine the image regions that correspond to a particular view, a task that is often performed manually. Within these viewing regions, standard perspective camera calibration

³<http://www.screen-tech.eu>

⁴See definition in Section 3.2.2

techniques can be employed [Bouguet05]. In the case of single bounce observation, this calibration is usually sufficient.

Single Mirror, Single-Bounce

In case of a moving mirror, it is usually necessary to estimate the mirror pose with respect to the recording camera, since an offline calibration step cannot easily be employed. For this purpose, self-identifying markers that are attached to the mirror can be used [Ihrke08]. Moving platforms are also often employed in the case of robotic applications. The case of a two-planar mirror setup with a moving camera mounted on a robotic platform has been analyzed in [Mariottini09]. The authors derive a calibration procedure for computing the pose of the camera with respect to the mirrors as well as the mirrors' relative position and orientation.

Multiple Mirrors, Multi-Bounce

In the case of multi-bounce observation, the mirror poses as well as the single real camera pose need to be estimated very accurately since the calibration error increases exponentially with the level of reflection. For this reason, special calibration procedures are necessary. In [Ramalingam11] a fixed (and known) mirror geometry is assumed and an algorithm for pose recovery of the real camera that is based on scene point correspondences (without knowing their reflection level) is derived.

The manual identification of reflection levels in a multi-bounce image is tedious and error prone. In Chapter 7, we propose an automatic procedure that can recover the number of mirrors and their parameters without user intervention. Currently, the method is restricted to $2^{\frac{1}{2}}D$ settings.

2.5 Connection to Time-of-Flight Imaging and the Multi-Bounce Problem

The time-of-flight problem, at first hand, appears to be disconnected from the problem settings considered so far and in fact, the literature is largely orthogonal to that of kaleidoscopic systems. In time-of-flight imaging, a pulse is emitted at one spatial location and the time difference until the signal returns is measured by the sensor. The classical time-of-flight technique is RADAR, where radio waves are used as probes. SONAR uses sound waves and LIDAR is using light pulses, usually infrared, to determine the distance of objects. In pulse-based time-of-flight imaging, most commonly, a single reflection of the emitted pulse from the environment is assumed. This situation is equivalent to a single-mirroring operation. In practice, multiple echoes, or multi-bounce signals, can corrupt the detection. Most

often, these echoes are considered to be undesirable noise and filtering procedures are developed to identify first times of arrival, see [Scheuing06] and the references therein.

Multi-bounce analysis in this area is investigating the forward modeling of reverberation and recovery of a room geometry from impulse responses of a room. The forward modeling frequently employs unfolding procedures, Section 3.1, for Coxeter geometries [Allen79], or for arbitrary polyhedral models [Borish84].

The recovery of room geometries from multi-bounce data often considers the special case of a rectangular Coxeter geometry [Ribeiro12] also known as the *shoebox* model, which allows for the interpretation as a perfectly subdivided space. Only recently methods for general convex geometries have started to appear ([Dokmanic11, Antonacci12, Dokmanic13] and the references therein). These methods usually assume the first-bounce, other reflection levels, or the number of walls of the room to be known. For example, in the work proposed by Tervo et. al. [Tervo12], the method assumes, that the position of the real source (zero-bounce source) is known and all the first level reflection sources are visible. Our method, see Part IV, requires only a sufficient amount of virtual sources to be observed and is capable of inferring the room geometry even when the real source and the first level sources are not visible. Moreover, the observation region might be restricted (for example by angle of view of the receiver). However, the method is currently implemented for $2\frac{1}{2}D$ rooms only.

We would also like to point out recent developments that enable the recording of the temporal profile of light for every pixel [Kirmani09]. While the initial *transient imaging* work used a very expensive femto-second laser setup, recently the use of a standard time-of-flight imager for the measurement of transient images has been proposed [Heide13, Kadambi13]. The information acquired with these devices can be used to reconstruct geometry from indirectly observed bounces, i.e. the geometry of hidden objects [Velten12]. An extensive overview of the computer vision literature on time-of-flight is presented in [Grzegorzec13].

2.6 Curved Mirrors

Curved mirrors are not the focus of our work, but we would like to touch on this subject for the sake of completeness.

Curved mirrors are different from planar ones in the sense that they usually do not yield perspective views but rather transform the world according to their surface curvature. One can consider the curved mirror as a surface, that, at each point, has a corresponding planar mirror that is tangent to the surface.

In order to use such mirrors in practice, their geometry and pose with

respect to a recording camera or a projector has to be known very accurately. It is a difficult problem to estimate general mirror shapes precisely, [Ihrke10]. Therefore, in practice, only a limited number of mirror shapes are considered. The classes of mirrors utilized in practical settings so far are restricted to conic sections and to axially symmetric mirrors. These simple types of curved mirrors can be classified into the following groups:

- Circular cone mirrors, including cylinders ([Kuthirummal06]),
- Spherical mirrors ([Nayar88, Unger03, Lensch03, Lanman06a, Ding09b, Taguchi10b, Agrawal13, Lanman06b]),
- Elliptic mirrors ([Mukaigawa07a, Mukaigawa07b]),
- Parabolic mirrors ([Nayar97, Gluckman98b, Gluckman98c, Dana01, Dana04, Ghosh07]),
- Hyperbolic mirrors ([Cabral04, Jang05]), and
- Cylindrical mirrors ([Ding09a]).

We want to highlight circular cone mirrors and mirror surfaces of revolution with piece-wise linear cross-section as those where the light trajectories can be explained with the theoretical apparatus for $2D$ planar mirror systems, developed in Part II, under the condition that imaging rays and the mirror axis are coplanar. This case often occurs in practice.

An in-depth discussion of curved mirrors in computer vision according to our classification, Fig. 2.1, can be found in our recent overview article, [Reshetouski13a].

Basic Concepts, Experimental Setting, Calibration, and Pre-Processing

This chapter introduces the basic theoretical concepts and initial experimental steps which are prerequisites to our work.

Section 3.1 describes the *unfolding principle* and can be considered a gentle introduction to the theoretical Part II. In Section 3.2, we introduce the definition of the visual hull which will be heavily used in Part III of this thesis. We also discuss the necessary terminology for understanding our results in reflectance acquisition in that part. In Section 3.3, we describe the prerequisites (experimental setting, calibration, pre-processing) to our practical solutions which are not the focus of this thesis, but important in order to achieve our goals.

3.1 Unfolding - A Convenient Way for Interpreting Image Formation in Planar Mirror Systems

Our work is based on the *ray unfolding* procedure. According to Coxeter [Coxeter67] one of the first to apply ray unfolding was H. A. Schwarz (1843-1921) for his solution of Fagnano's problem, using a 5-fold unfolding of an acute-angled triangle. Unfolding is also a common operation in




Figure 3.1: Unfolding of a single reflection.

the optical description of prism and mirror systems [Smith08] and in the time-of-flight literature.

In this technique, every mirror interaction is applied to the world instead of the ray. The result is a straight ray that passes through a sequence of virtual copies of the world that is equivalent to the bouncing ray in the real world. This way, complex ray interactions can be visualized in an intuitive manner and a change of coordinate systems can easily be tracked. In the following, we will introduce and apply unfolding to different types of planar mirror systems.

Single-Mirror, Single-Bounce Consider a single planar mirror M and a camera S observing an object point P via a single-bounce reflection, Fig. 3.1.

When a ray of light is hitting the mirror it is mirrored from the plane according to the law of reflection. Instead of mirroring the ray, we can consider that the world is being reflected, creating a virtual world, or as we will call it, a virtual *chamber*. In this case, the ray appears to continue straight into the virtual mirror world. The mirror copy of the scene is an isometric transformation of the real world. The world coordinate system is transformed to the mirrored one by reflecting it in the mirror plane. Left-handed mirror systems transform into right-handed ones and vice versa. The procedure of ray straightening just described is called *unfolding*. Because light paths are reversible, we can consider the ray straightening procedure from the point of view of a scene point or from the point of view of a camera or a projector. Consider a ray from camera S observing a scene point P through a reflection from the planar mirror M . Then from the point of view of the camera, we observe the virtual point P' which is the mirror copy of the real point P . But from the point of view of the point P we are observing the virtual camera S' which is the reflection of the real camera S .

Multi-Mirror, Single-Bounce per Mirror If there are several planar mirrors that are arranged around a camera, as for example in Fig. 3.2 (left),




Figure 3.2: Two planar mirrors: unfolding for two different rays (left) and unfolding for sequential reflection (right).

for rays hitting different mirrors the ray straightening process will introduce a different virtual world (or a different virtual camera if we consider the point of view from the scene). A second possibility is to arrange the planar mirrors such, that there is a sequential ray bouncing from mirror to mirror as shown in Fig. 3.2 (right). In this case the unfolding procedure is applied recursively. Thus, if an even number of reflections is involved, the resulting virtual world (virtual camera) coordinate system will not change its handedness while it changes handedness if the reflection level is odd.

As long as the reflection sequence includes every mirror only once, the recursive unfolding procedure can be applied without ambiguity.

Multi-Mirror, Multi-Bounce However, multiple bounces in systems with several planar mirrors could be such, that the same mirrors are participating in a reflection sequence multiple times. In a theoretical setting, this number could well be infinite.

The simplest such system is an angle constructed from two planar mirrors as in Fig. 3.3. If the angle $\angle ABC$ between the mirrors is $\frac{\pi}{k}$, where $k \in \mathbb{N} \setminus \{1\}$, then the unfolding of all possible rays will introduce a partitioning of the space into continuous regions such that the space is divided into $2k$ different parts. These are the inner part of the original angle (base chamber) and the copies associated with different reflection levels (virtual chambers). The partitioning is, in this case, independent of the origin of the ray, see Fig 3.3.

A simple example involving several mirrors is a bouncing ray inside a rectangular room, see Fig. 3.4. This type of geometry is most often con-




Figure 3.3: Rays are bouncing inside an angle with matching coordinate systems. The magenta and green rays are propagating after the second bounce in the common virtual chamber with coordinate system $O''' X''' Y'''$.




Figure 3.4: Ray bouncing inside the rectangle $ABCD$. Light propagates from point S up to point P .




Figure 3.5: Unfolding of the ray from Fig. 3.4 bouncing inside the rectangle.

sidered in multi-bounce time-of-flight imaging, Sect. 2.5. If we repeatedly unfold the ray while it is propagating in space, we obtain the result seen in Fig. 3.5. In every virtual rectangle (virtual chamber) we have a virtual world that is specific to the sequence of reflections. If we consider all possible ray directions from any possible inner point of the original rectangle, we obtain a partitioning of the space into virtual rectangles. Since the rectangle is a regular structure, unfolding via different reflection sequences yields the same virtual worlds (perfectly overlapping chambers and equal coordinate systems), independent of the sequence of reflections we travel along the ray to reach the virtual rectangle from the real one (see Fig. 3.6).

Unfortunately, only the two-mirror wedges with angles $\frac{\pi}{k}$, a few types of polygons and another single special case (see the note below) produce perfect space partitioning schemes. In these cases, the partitioning is independent of the initial ray position. The polygons (or polyhedra in the 3D case) having this property are known as Coxeter polygons (polyhedra). A polygon is a Coxeter polygon iff all its angles are in the form of $\frac{\pi}{k}$, $k \in \mathbb{N} \setminus \{1\}$. There are only 4 such polygons: rectangles, equilateral triangles, the isosceles right triangles, and right triangles with angles $\frac{\pi}{3}$ and $\frac{\pi}{6}$.

In the 3D case, the dihedral angles $\frac{\pi}{k}$ produce a perfect partitioning. Moreover, we can add one or two orthogonal walls (to both sides) to such types of dihedral angles without losing the perfect partitioning property. Here we can also add trihedral angles consisting of the following dihedral angle triples: $(\frac{\pi}{5}, \frac{\pi}{3}, \frac{\pi}{2})$, $(\frac{\pi}{4}, \frac{\pi}{3}, \frac{\pi}{2})$, $(\frac{\pi}{3}, \frac{\pi}{3}, \frac{\pi}{2})$.

For polyhedra in 3D, the condition to be a Coxeter polyhedron is that




Figure 3.6: Ray bouncing from two different camera locations, S_1, S_2 to the same object point P .




Figure 3.7: The seven Coxeter polyhedra in 3D [Sossinsky12].

all the dihedral angles are of the form $\frac{\pi}{k}$, $k \in \mathbb{N} \setminus \{1\}$. There are only 7 types of Coxeter polyhedra, Fig. 3.7.

Note, that the perfect partition property persists if we remove some of the sides from a Coxeter polygon or polyhedron. For example, a rectangle without one side still has this property.

All other types of planar mirror configurations (including other types of polygons and polyhedra) generate a more complicated space partitioning that depends on the ray origin. We will study this general situation in detail in the next part of the thesis, see Part II.

3.2 Background Concepts

3.2.1 The Visual Hull

The concept of the visual hull in the Euclidean space \mathbb{R}^3 was first introduced by A. Laurentini, see [Laurentini94].

Definition 3.2.1. The visual hull $VH(O, V)$ of an object $O \subset \mathbb{R}^3$ relative to a viewing region $V \subset \mathbb{R}^3$ is the region in \mathbb{R}^3 which consist of those and only those points $\mathbf{p} \in \mathbb{R}^3$ such that for each viewpoint $v \in V$ the half-line starting at v and passing through \mathbf{p} contains at least one point of O .

Another concept, related to the visual hull, is the object's silhouette, which is the projection of the object O onto a camera image. The silhouette information is usually available through a binary segmentation of the camera image into background and foreground. If the object O is situated completely inside the camera's frustrum, then we say that the object is fully observable by the camera. This means that every object point $p \in O$ can be projected onto the camera's image plane and that the projection is inside the camera's image.

In this thesis, we will mostly consider viewing regions V , that consist of a finite set of points - the positions of the projective centers of cameras or virtual cameras. If for all these cameras the object O is fully observable, the visual hull $VH(O, V)$ is equal to the intersection of generalized cones originating at the centers of projection of the cameras and passing through the corresponding object's silhouettes.

This second way of defining the visual hull (image-based visual hull) gives us a direct way of computing it as an intersection of generalized cones.

The visual hull is always larger or equal to the object ($O \subseteq VH(O, V)$) and the main use of visual hulls in computer graphics is the approximation of the object's geometry. As can be seen in Fig. 3.8, the visual hull from the image silhouettes of only 7 views is already close to the original object geometry.

On the other hand, to produce the silhouette of an object at viewpoint $v \in V$ it is sufficient to only know the visual hull $VH(O, V)$.

In this thesis we will use approximations of visual hulls to estimate silhouettes for our labeling solution, Chapter 5.

3.2.2 BRDF

An important property of any material is the way it is redistributing incoming light, or, in other words, its reflectance behavior. For simple homogeneous materials (for example, without subsurface scattering and without visible surface variations), interaction with single wavelength light can be




Figure 3.8: Example of the image-based visual hull of the bunny model obtained from 7 different viewpoints. Ground truth view (left), visual hull image with gray-coded (by 7 different gray levels) indication of influences of different viewpoints (middle), visual hull image with color-coded depth, where pixels closer to the camera are brighter (right).

characterized at the coarse scale by the Bidirectional Reflection Distribution Function (BRDF) $f_r(\omega_i, \omega_o)$ [sr^{-1}]. This function returns the fraction of light that is reflected from the incoming direction ω_i to the outgoing direction ω_o . Knowing the BRDF for a given material allows for the prediction of the material appearance under given (synthetic) lighting conditions, which is highly important in computer graphics.

It is worth to mention, that when it is necessary to characterize the appearance of a complex material at a fine scale, i.e. when surface variations are observable, the appearance can be more adequately described by the Bidirectional Texture Function (BTF).

According to the definition, the BRDF is a four-dimensional function. However, if we want to characterize the reflectance of a 3D object made of multiple materials, we need to assign a BRDF to every surface point of the object. This way, a six-dimensional Spatially Varying Bidirectional Reflectance Distribution Function (SVBRDF) $f_r(\omega_i, \omega_o, u, v)$ [sr^{-1}] can be introduced, where (u, v) are the two-dimensional coordinates on the object's surface.

The classical way to reconstruct reflectance of a given object is to sample the reflectance function and to interpolate it. In this thesis we will use a kaleidoscopic setup to distribute illumination and viewing directions in order to sample the SVBRDF of an object with a static single camera, single projector setup, Section 6.3.

3.3 Experimental Setting, Calibration, and Pre-Processing

In Part III of this thesis, we utilize raytracing in planar mirror systems and capture reflectance data. For this reason, the systems have to be calibrated before our approaches can be applied. The calibration process involves:

the geometric calibration of the mirrors, the camera and the projector (if present); the calibration of the mirrors reflectivity coefficients; and the radiometric calibration of the camera.

Later in this thesis we assume that the necessary calibration steps are already performed and concentrate on the conceptual aspects of the problems being solved.

3.3.1 Geometric Calibration

Geometric calibration is intended to relate camera or/and projector pixels with rays in a common $3D$ space S , as well as to relate mirrors with corresponding sets of $3D$ points in the same space. If the behavior of the camera/projector rays in the virtual setting matches the behavior of the camera/projector pixels in the real setting, we call the system geometrically calibrated. There are different ways to calibrate cameras, projectors or planar mirrors. For our practical applications there is no difference, which calibration method is selected and what kind of camera/projector models (with central projection or not) are used as long as we can predict camera/projector rays with sufficient precision.

In our practical experiments we used a central projection model for the camera and a non-central one for the projector [Manakov11]. Practically, we perform geometric calibration by first estimating the camera intrinsics using a number of checkerboard images using Zhang’s method [Zhang99] as implemented in Bouguet’s calibration toolbox [Bouguet05]. We then place the camera into our setup and observe a checkerboard pattern placed in the kaleidoscope at different heights above the ground plane, Fig. 3.9. This measure helps to keep the calibration consistent in the whole acquisition volume. We remove radial distortion from the images and identify the real image as well as the first-order reflections of the checkerboard. We compute plane equations using the known intrinsics and from those estimate initial guesses for the mirror planes. We perform a bundle adjustment procedure, optimizing the camera extrinsics and the mirror plane parameters while keeping the intrinsics fixed. We then use this initial calibration to predict the position of the second-order reflections. Matching to the extracted corners we optimize again to minimize the reprojection error. We then proceed in a similar manner to add third- and higher-order reflections. The checkerboards are typically well visible up to the sixth reflection order and we use all of this information in a final bundle adjustment of all parameters, camera intrinsics, extrinsics and mirror plane parameters. The reprojection error is typically within 1.2 pixels.

The laser projector was calibrated for its internal parameters using the non-central projection model and procedure of Manakov et al. [Manakov11]. To improve the position and orientation estimate of the laser projector with respect to the mirror system and the camera we place a checkerboard in-



Figure 3.9: Images of our kaleidoscope with the checkerboard at two different heights.

side the mirror system and illuminate a set of projector coordinates, taking an image for every illuminated point. Since the camera/mirror system is calibrated already, the checkerboard allows for the computation of the 3D coordinates of the illuminated point. These are usually in disagreement with the prediction produced by the initial calibration of the laser system. We perform an optimization on the laser intrinsic and extrinsic parameters to improve this prediction. In a final step, we perform a bundle adjustment on the parameters of all system components.

3.3.2 Radiometric Calibration




Figure 3.10: Image of a "Macbeth ColorChecker" chart placed in the kaleidoscope and illuminated by laser "white light".

We linearize the camera response curve by taking exposure sequences of 5 images and applying the technique of Robertson et al. [Robertson03] as implemented in the PFSTools package [Mantiuk06].

For the projector, since we use a laser source that is only roughly white, we compute an affine RGB correction matrix using a "Macbeth ColorChecker" chart positioned inside the mirror system and illuminated by laser "white light", see Fig. 3.10. This color transformation is applied to all captured images.

We estimate the attenuation factor of the mirrors by imaging a uniformly illuminated Spectralon target, an almost perfectly Lambertian reflector, inside the system. Since the mirroring sequence for every virtual camera is known from the system geometry, the individual mirrors' attenuation coefficients can be estimated from observed products of these factors on the Spectralon patch as it is visible in different camera chambers. In our case, the coefficient values were constant for the different color channels.

PART II

Theory of Planar Mirror Systems

A Theory of Projective Imaging within Mirror Systems

The unfolding principle described above in Section 3.1 is very useful. However, if we want to deal with complicated planar mirror configurations, the unfolding approach needs to be extended.

Section 4.1 of this part of the thesis is dedicated to mathematically formulate the physical propagation of idealized light rays inside a planar mirror system and to develop some basic tools for dealing with them. As a result, we propose a natural extension of the unfolding procedure by introducing a so called space partitioning.

Section 4.2 uses the theoretical basics from the previous section to obtain the first non-trivial result: the uniform continuity of a bounded part of the space partitioning under observation point change.

4.1 Mirror Unfolding and Space Partitioning

4.1.1 Basic Definitions

Prerequisites

In the following, we introduce the notation and tools used in the remainder of this thesis. We use the n -dimensional Euclidean space \mathbb{R}^n with its usual two-norm. We make use of the topology of the Euclidean space with its open and closed sets induced by the metric. Many of the mathematical objects

defined below will be sets of points. We denote elements of \mathbb{R}^n as lower-case bold characters (for example \mathbf{x} , \mathbf{y} , \mathbf{v}) and real scalars as italic lower-case characters (for example a , b , c).

Definition 4.1.1. Distances.

The distance $d(\mathbf{p}_1, \mathbf{p}_2)$ between two points \mathbf{p}_1 and \mathbf{p}_2 of \mathbb{R}^n is, as usual, the norm of the difference: $d(\mathbf{p}_1, \mathbf{p}_2) = \|\mathbf{p}_1 - \mathbf{p}_2\|$.

We define the distance $d(S_1, S_2)$ between two sets of points S_1 and S_2 , both from \mathbb{R}^n , as $d(S_1, S_2) = \inf_{\substack{\mathbf{s}_1 \in S_1 \\ \mathbf{s}_2 \in S_2}} d(\mathbf{s}_1, \mathbf{s}_2)$.

In our derivations, we make extensive use of rays.

Definition 4.1.2. Ray.

A ray $\tilde{r} := \tilde{r}(\mathbf{c}, \mathbf{d})$ with origin \mathbf{c} and direction $\mathbf{d} \neq \mathbf{0}$ (see Fig. 4.1) is a set of points $\tilde{r}(\mathbf{c}, \mathbf{d}) = \{\mathbf{c} + \lambda \mathbf{d} \mid \lambda \in \mathbb{R}, \lambda \geq 0\}$.





Figure 4.1: Ray.

A ray is a half-line, the points of which are ordered according to the parameter λ which can be thought of as a time value: For two points \mathbf{p}_1 and \mathbf{p}_2 on the same ray $\tilde{r}(\mathbf{c}, \mathbf{d})$ with $\mathbf{p}_1 = \mathbf{c} + \lambda_1 \mathbf{d}$ and $\mathbf{p}_2 = \mathbf{c} + \lambda_2 \mathbf{d}$ ($0 \leq \lambda_1 < \lambda_2$), we will speak about the ray \tilde{r} passing point \mathbf{p}_1 earlier than point \mathbf{p}_2 . In the following, we will use the ray symbol \tilde{r} in two ways: 1) if it is written as \tilde{r} , i.e. without parameters, it denotes the point set constituting the ray; 2) if parameters are indicated, e.g. $\tilde{r}(\mathbf{c}, \mathbf{d})$, it is acting as a function returning the point set \tilde{r} according to the parameters \mathbf{c} and \mathbf{d} . An expression like $\tilde{r}(\tilde{r} \cap f, \mathbf{d})$ indicates a new starting point for the ray \tilde{r} at the intersection point with the set f (assuming there is only a single one). Typically, f will be a planar set.

Our goal is to describe the propagation of rays in systems of planar mirrors. In the following, we define one-sided mirrors in terms of subsets of hyperplanes that are considered to be the mirroring portion of the hyperplane. The mirroring side is described by the orientation of the hyperplane.

Definition 4.1.3. Oriented hyperplane.

An oriented hyperplane in \mathbb{R}^n , parameterized by the unit normal \mathbf{n} and the offset d , is the set of points $F := F(\mathbf{n}, d) = \{\mathbf{x} \mid \mathbf{n} \cdot \mathbf{x} + d = 0\}$ together with the orientation of its normal.

Figure 4.2: Oriented hyperplane. \mathbf{o} - origin of the space.

A hyperplane induces two half-spaces

$$F_{in} := \{\mathbf{x} \mid \mathbf{n} \cdot \mathbf{x} + d > 0\}, \text{ and}$$


$$F_{out} := \{\mathbf{x} \mid \mathbf{n} \cdot \mathbf{x} + d < 0\},$$

the inside and outside half-spaces, respectively. The unit normal \mathbf{n} points towards the outside half-space, see Fig. 4.2. The hyperplane does not belong to either half-space.

We refer to points relative to a given hyperplane as inside points if they are elements of the inside half-space, as outside points when they are part of the outside half-space, and as points on the plane if they belong to the defining hyperplane.

Definition 4.1.4. Coincident and identical oriented hyperplanes

Two oriented hyperplanes F_1 and F_2 are coincident ($F_1 \sim F_2$) iff they have equal point sets, Fig. 4.3. They are identical ($F_1 = F_2$) iff they are coincident and have the same orientation.

Figure 4.3: Coincident hyperplanes F_1 and F_2 .

Reflections, Mirrors, and Systems thereof

Definition 4.1.5. Mirror. A *mirror* is an ordered pair $M := M(f, F) = (f, F)$ consisting of a *mirror face* f and a supporting oriented hyperplane F , where f is a closed and bounded nonempty subset of F .

The concept is illustrated in Fig. 4.4. Mirrors are one-sided by the orientation of the supporting hyperplane, the mirroring side being in the direction towards the inside halfspace. Some examples for mirrors are a) a planar convex set f , b) unions of such sets sharing the same supporting hyperplane, and more generally, c) arbitrary bounded closed sets with the same property. As a counter-example, the full hyperplane F is un-bounded and is not permissible as a mirror face.

The definition implies that all points $\mathbf{p} \in M$ are also points of the supporting hyperplane. Mirrors are therefore flat even if the outline of their reflective area may be irregular or may contain holes.




Figure 4.4: Mirror.

We classify points in \mathbb{R}^n relative to a given mirror as inside, outside, or points on the plane, if these conditions apply to the mirror's hyperplane. The points are said to be on the mirror if, in addition to being on the hyperplane, they belong to the mirror's face.

Definition 4.1.6. Reflection.

A point reflection operator $m := m_M(\mathbf{x})$ for mirror $M(f, F)$ is a function $m_M : \mathbb{R}^n \rightarrow \mathbb{R}^n$, given by the mirror operation at the corresponding hyperplane F : $\mathbf{x}' = \mathbf{x} - 2(d_M + \mathbf{n}_M \cdot \mathbf{x})\mathbf{n}_M$.

This is the standard definition of the reflection. It can be written in matrix notation as: $\mathbf{x}' = \mathbf{H}_M \mathbf{x} - 2d_M \mathbf{n}_M = (I - 2\mathbf{n}_M \mathbf{n}_M^T) \mathbf{x} - 2d_M \mathbf{n}_M$, where $\mathbf{H}_M = I - 2\mathbf{n}_M \mathbf{n}_M^T$ is a Householder transformation matrix.

Note: We sometimes apply a function that is defined point-wise to a set of points. This notation implies element-wise application, producing a new set: Let $f : \mathbb{R}^n \mapsto \mathbb{R}^n$, $f(A) := \{\mathbf{b} | \mathbf{b} = f(\mathbf{a}), \forall \mathbf{a} \in A\}$. As an example $m_M(\tilde{r})$ is a set of points belonging to a ray \tilde{r} mirrored in mirror M .

The point reflection operator, applied to all of \mathbb{R}^n results in a global isometry, i.e. it is a bijective operator and preserves point-wise distances.

In the following, we consider a mirror system and its surrounding space as sets of points that are in certain relations to each other such as incidence, element of a half space, etc. A global reflection, due to its isometric property, then transforms the whole space together with its objects such that the relations still hold for the mirrored (we also say *virtual*) object versions.

As an example, the inside, outside, and point in hyperplane conditions are maintained with respect to the reflected hyperplane of the mirror.

Our goal is to work with a generic reflection operator that can be applied to all objects of interest (i.e. mirrors, points, rays, ...). For this it is necessary to formally define what is meant by reflecting an object in \mathbb{R}^n . We start by defining the reflection of hyperplanes as building blocks for mirrors and mirror systems.

Lemma 1. *The reflection of an oriented hyperplane $F = F(\mathbf{n}, d) = \{\mathbf{x} \mid \mathbf{n} \cdot \mathbf{x} + d = 0\}$ from a mirror M is another oriented hyperplane $m_M(F) = F' = F(\mathbf{n}', d')$, where $\mathbf{n}' = \mathbf{H}_M \mathbf{n}$ and $d' = d - 2d_M \mathbf{n}_M \cdot \mathbf{n}$.*

Proof. To proof the assertion, we show that $\forall \mathbf{x} \in \mathbb{R}^n$: 1. $\mathbf{x} \in F \Leftrightarrow m_M(\mathbf{x}) \in F'$. 2. $\mathbf{x} \in F_{in} \Leftrightarrow m_M(\mathbf{x}) \in F'_{in}$.

After some algebraic manipulations we find that indeed $\mathbf{n}' \cdot m_M(\mathbf{x}) + d' = \mathbf{n} \cdot \mathbf{x} + d$. Therefore, 1. and 2. are satisfied for all $\mathbf{x} \in \mathbb{R}^n$. \square

We apply the same principle to define the reflection of a mirror M_0 from another mirror M . Because M_0 is a pair consisting of a hyperplane and a point set (F_0, f_0) , we define the reflection $m_M(M_0) = M_0' = (m_M(F_0), m_M(f_0))$ to be the mirror that consists of the oriented hyperplane and face obtained by reflection from mirror M .

Similarly, a ray $\tilde{r} = \tilde{r}(\mathbf{c}, \mathbf{d})$ that is reflected from mirror M is a reflection of all the ray's points in the mirror's hyperplane.

Lemma 2. *A reflected ray is equal to a new ray $\tilde{r}' = m_M(\tilde{r}) = \tilde{r}(\mathbf{c}', \mathbf{d}')$, where $\mathbf{c}' = m_M(\mathbf{c})$ and $\mathbf{d}' = \mathbf{H}_M \mathbf{d}$.*

Proof. $\tilde{r}(\mathbf{c}, \mathbf{d}) = \{\mathbf{c} + \lambda \mathbf{d} \mid \lambda \in \mathbb{R}, \lambda \geq 0\}$. Then

$$\begin{aligned} m_M(\tilde{r}) &= m_M(\{\mathbf{c} + \lambda \mathbf{d} \mid \lambda \in \mathbb{R}, \lambda \geq 0\}) \\ &= \{m_M(\mathbf{c} + \lambda \mathbf{d}) \mid \lambda \in \mathbb{R}, \lambda \geq 0\} \\ &= \{\mathbf{H}_M(\mathbf{c} + \lambda \mathbf{d}) - 2d_M \mathbf{n}_M \mid \lambda \in \mathbb{R}, \lambda \geq 0\} \\ &= \{\mathbf{H}_M \mathbf{c} - 2d_M \mathbf{n}_M + \lambda \mathbf{H}_M \mathbf{d} \mid \lambda \in \mathbb{R}, \lambda \geq 0\} \\ &= \{m_M(\mathbf{c}) + \lambda \mathbf{H}_M \mathbf{d} \mid \lambda \in \mathbb{R}, \lambda \geq 0\} \end{aligned}$$

\square

Having defined the mirroring operation for different structures, i.e. sets of points, oriented hyperplanes, mirrors, and rays, we can now reflect several




Figure 4.5: Physical reflection $p_M(\tilde{r}(\mathbf{c}, \mathbf{d})) = \tilde{r}(\mathbf{c}_1, \mathbf{d}')$ of a ray from mirror $M = M(f, F)$.

objects from one given mirror simultaneously by applying the global mirror operator m_M on the whole of \mathbb{R}^n including its objects.

We will denote the application of the reflection operator for the object types: points, sets of points, mirrors, oriented hyperplanes, and rays, as $m_M(*)$, where $*$ should be substituted by the corresponding object.

Note that if we reflect an object $*$ from a mirror $M(f, F)$, then $F \cap * = F \cap m_M(*)$ because the reflection operation is leaving all points on the mirror's hyperplane F unchanged.

Definition 4.1.7. Physical reflection.

Let $\tilde{r} = \tilde{r}(\mathbf{c}, \mathbf{d})$ be a ray in \mathbb{R}^n and $M(f, F)$ a mirror, also in \mathbb{R}^n . A *physical reflection* (see Fig. 4.5) of the ray \tilde{r} from mirror M is

$$p_M(\tilde{r}) = \begin{cases} m_M(\tilde{r}(\tilde{r} \cap f, \mathbf{d})), & \text{if } \mathbf{c} \in F_{in} \text{ and } \tilde{r} \cap f \neq \emptyset \\ \emptyset, & \text{otherwise.} \end{cases}$$

From the definition of the physical reflection of a ray, we see that rays can be reflected only from the inner side of the mirror and only when the ray intersects the mirror at the mirror's face. In this case, the physical reflection is again a ray.

The physical reflection is undefined if ray \tilde{r} does not intersect any mirrors, if the ray traverses inside a mirror's hyperplane, or if the ray intersects a mirror from its outside half-space. This models the real world behavior of rays of light interacting with mirrors.

To simplify the notation, we introduce a short-hand for the composition of reflections (or physical reflections) from mirrors M_1, M_2, \dots, M_N : Instead of writing $m_{M_N}(m_{M_{N-1}}(\dots(m_{M_1}(*))\dots))$ (or $p_{M_N}(p_{M_{N-1}}(\dots(p_{M_1}(*))\dots))$) we will write $m_{M_N M_{N-1} \dots M_1}(*)$ (or $p_{M_N M_{N-1} \dots M_1}(*)$). The inverse $m_{M_N M_{N-1} \dots M_1}^{-1}(*)$ of a mirror operation is $m_{M_1 M_2 \dots M_N}(*)$.

Having introduced a mirror operator for all our objects in the single mirror case, we now define more complex *mirror systems* that consist of multiple mirrors.

Definition 4.1.8. Mirror system. A *mirror system* (see Fig. 4.6) is a finite set of mirrors $\mathcal{M} = \{M_i = (f_i, F_i) \mid i = 1..N, N \in \mathbb{N}\}$, subject to the condition: if $f_i \cap f_j \neq \emptyset \Rightarrow (F_i \sim F_j \text{ and } F_i \neq F_j)$. N is the number of mirrors in the system.

The mirror system defined such has the following restriction: The only way that a point in the ambient space \mathbb{R}^n can be assigned to more than one mirror is if these mirrors share the same hyperplane but have opposite orientations, i.e. we allow mirrors to be two-sided, but we do not allow mirror faces from different hyperplanes to intersect.

Definition 4.1.9. Set of points of the mirror system. $\mathcal{M}_f := \cup_i f_i$.

Ray Propagation in Planar Mirror Systems

In the following we introduce a formal description of ray propagation in planar mirror systems. Since ray paths in such systems are polygonal, they consist of straight ray segments. We define such segments, giving proper attention to exclude special cases such as rays propagating inside a mirror's hyperplane. Straight segments are again defined as sets of points belonging to them.

Definition 4.1.10. Straight segment.

Let $\tilde{r} = \tilde{r}(\mathbf{c}, \mathbf{d})$ be a ray in a mirror system $\mathcal{M} = \{(f_i, F_i) \mid i = 1..N\}$ such that the ray is not completely inside some mirror's hyperplane where the ray origin is an element of that mirror's face, i.e. $\mathbf{c} \in f_i \Rightarrow \tilde{r} \not\subset F_i$. Then the straight segment $s_{\mathcal{M}}(\tilde{r})$ of the ray \tilde{r} in the mirror system \mathcal{M} is defined as $s_{\mathcal{M}}(\tilde{r}) := \bigcap_i (\tilde{r} \setminus \tilde{r}_i)$, where

$$\tilde{r}_i = \begin{cases} \tilde{r}(\operatorname{arg\,inf}_{\mathbf{a} \in (\tilde{r} \setminus \{\mathbf{c}\}) \cap f_i} d(\mathbf{a}, \mathbf{c}), \mathbf{d}), & \text{if } (\tilde{r} \setminus \{\mathbf{c}\}) \cap f_i \neq \emptyset \\ \emptyset, & \text{otherwise.} \end{cases}$$

Properties: The straight segment can be a finite interval or semi-infinite. The first case (finite interval) occurs if the ray (without its origin) intersects the face of a mirror (under conditions of the definition of the straight segment). In this case, the straight segment is connecting the ray origin to the point on the first mirror, excluding the point on the mirror itself. The second case (semi-infinite interval) occurs if the ray (without its origin) does not hit any mirror face. It continues to propagate infinitely. More precisely:

1. If \tilde{r}_i is non-empty, it is a semi-infinite set starting on mirror M_i , but not from the point \mathbf{c} :
 - If \tilde{r}_i is non-empty, $(\tilde{r} \setminus \{\mathbf{c}\}) \cap f_i \neq \emptyset$.

- Because mirrors are closed sets, $\operatorname{arginf}_{\mathbf{a} \in (\tilde{r} \setminus \{\mathbf{c}\}) \cap f_i} d(\mathbf{a}, \mathbf{c})$ is a point on the mirror M_i if $(\tilde{r} \setminus \{\mathbf{c}\}) \cap f_i \neq \emptyset$.
 - Consider the mirror $M_i \in \mathcal{M}$:
 - (a) If $\mathbf{c} \notin F_i$, which implies $\mathbf{c} \notin f_i$, and $\tilde{r} \cap f_i \neq \emptyset$, then the expression $\operatorname{arginf}_{\mathbf{a} \in (\tilde{r} \setminus \{\mathbf{c}\}) \cap f_i} d(\mathbf{a}, \mathbf{c})$ is equal to $\tilde{r} \cap f_i \neq \mathbf{c}$.
 - (b) If $\mathbf{c} \in F_i$, but $\mathbf{c} \notin f_i$, then, using that $(\tilde{r} \setminus \{\mathbf{c}\}) \cap f_i \neq \emptyset$ and that f_i is a closed set, the expression $\operatorname{arginf}_{\mathbf{a} \in (\tilde{r} \setminus \{\mathbf{c}\}) \cap f_i} d(\mathbf{a}, \mathbf{c})$ is equal to $\operatorname{argmin}_{\mathbf{a} \in (\tilde{r} \setminus \{\mathbf{c}\}) \cap f_i} d(\mathbf{a}, \mathbf{c}) \neq \mathbf{c}$.
 - (c) If $\mathbf{c} \in f_i$, then, using that $(\tilde{r} \setminus \{\mathbf{c}\}) \cap f_i \neq \emptyset$, the ray $\tilde{r} \in F_i$ (two different ray points are on the hyperplane) with the origin $\mathbf{c} \in f_i$ - that is the case excluded by Def. 4.1.10 of a straight segment.
2. *When defined, the straight segment connects \mathbf{c} and the first mirror, excluding the one at the origin. It is equal to the ray if no mirror is intersected:*

- The straight segment of a ray $\tilde{r}(\mathbf{c}, \mathbf{d})$ is the part of the ray from its origin \mathbf{c} to the first intersection with any mirror in the mirror system: According to the definition of \tilde{r}_i we have that $\forall i, j \in \{1, \dots, N\} \Rightarrow (\tilde{r} \setminus \tilde{r}_i) \subseteq (\tilde{r} \setminus \tilde{r}_j)$ or $(\tilde{r} \setminus \tilde{r}_j) \subseteq (\tilde{r} \setminus \tilde{r}_i)$. Therefore, $\exists k \in \{1, \dots, N\}$ such that $s_{\mathcal{M}}(\tilde{r}) = (\tilde{r} \setminus \tilde{r}_k)$: The straight segment is the part of the ray between the point of origin \mathbf{c} and the first intersection with a mirror $M_k \in \mathcal{M}$.
- If there are no mirrors on the ray $\tilde{r} \setminus \mathbf{c}$, then the straight segment of the ray is the ray itself.
- If the ray origin is part of some mirror M_0 , the straight segment is equivalent to the straight segment of the ray in a mirror system $\mathcal{M} \setminus M_0$.

The Double-Sided Mirror

It should be noted that the intersection of the ray with the first mirror may not be unique due to our mirror definition (Def. 4.1.5) that allows for two-sided mirrors.

Ambiguity of the mirror intersection point: We have shown above that for the straight segment, there $\exists k \in \{1, \dots, N\}$ such that $s_{\mathcal{M}}(\tilde{r}) = (\tilde{r} \setminus \tilde{r}_k)$, with $\tilde{r}_k \neq \emptyset$. However, there can be more than one such k , but no more than two:

If $s_{\mathcal{M}}(\tilde{r}) = (\tilde{r} \setminus r_{k_1}) = (\tilde{r} \setminus r_{k_2}) \Rightarrow r_{k_1} = r_{k_2}$. Then, because $r_{k_1} = r_{k_2} \neq \emptyset \Rightarrow$

$$\begin{aligned} \tilde{r}(\operatorname{arginf}_{\mathbf{a} \in (\tilde{r} \setminus \{\mathbf{c}\}) \cap f_{k_1}} d(\mathbf{a}, \mathbf{c}), \mathbf{d}) &= \tilde{r}(\operatorname{arginf}_{\mathbf{a} \in (\tilde{r} \setminus \{\mathbf{c}\}) \cap f_{k_2}} d(\mathbf{a}, \mathbf{c}), \mathbf{d}) \Rightarrow \\ \operatorname{arginf}_{\mathbf{a} \in (\tilde{r} \setminus \{\mathbf{c}\}) \cap f_{k_1}} d(\mathbf{a}, \mathbf{c}) &= \operatorname{arginf}_{\mathbf{a} \in (\tilde{r} \setminus \{\mathbf{c}\}) \cap f_{k_2}} d(\mathbf{a}, \mathbf{c}) \neq \emptyset. \end{aligned}$$

As noted before in our remarks corresponding to property 1), if $\operatorname{arginf}_{\mathbf{a} \in (\tilde{r} \setminus \{\mathbf{c}\}) \cap f_{k_1}} d(\mathbf{a}, \mathbf{c})$ and $\operatorname{arginf}_{\mathbf{a} \in (\tilde{r} \setminus \{\mathbf{c}\}) \cap f_{k_2}} d(\mathbf{a}, \mathbf{c})$ are non-empty, they are points on the mirror M_{k_1} and M_{k_2} , respectively. By definition of the mirror system (Def. 4.1.8), these points can only be equal if the oriented hyperplanes F_{k_1} and F_{k_2} are coincident ($F_{k_1} \sim F_{k_2}$), but not equal ($F_{k_1} \neq F_{k_2}$). Put differently, F_{k_1} and F_{k_2} are differently oriented (i.e. a two-sided mirror). Moreover, in the case of more than two pairwise coincident oriented hyperplanes, at least two of them will be equal, which is impossible by the definition of \mathcal{M} .

Selection of the reflecting mirror side: Let $\tilde{r} = \tilde{r}(\mathbf{c}, \mathbf{d})$ be a ray in a mirror system $\mathcal{M} = \{(f_i, F_i) \mid i = 1..N\}$ and $\mathbf{c} \notin \mathcal{M}_f$. The last condition ensures that the ray \tilde{r} starts from a point that does not belong to any of the mirrors. This implies that the straight segment $s_{\mathcal{M}}(\tilde{r})$ is non-empty. If $s_{\mathcal{M}}(\tilde{r}) \neq \tilde{r}$, i.e. the ray intersects a mirror, then there $\exists k_1 \in \{1, \dots, N\}$ such that $s_{\mathcal{M}}(\tilde{r}) = (\tilde{r} \setminus r_{k_1})$, where $r_{k_1} = \tilde{r}(\operatorname{arginf}_{\mathbf{a} \in (\tilde{r} \setminus \{\mathbf{c}\}) \cap f_{k_1}} d(\mathbf{a}, \mathbf{c}), \mathbf{d}) \neq \emptyset$.

$\operatorname{arginf}_{\mathbf{a} \in (\tilde{r} \setminus \{\mathbf{c}\}) \cap f_{k_1}} d(\mathbf{a}, \mathbf{c})$ is a single point on the mirror $M_{k_1} \in \mathcal{M}$.

If we can select more than one (i.e. two) such k_1 , then we choose the one for which $\mathbf{c} \in F_{in}$ of F_{k_1} . In this case we are encountering a two-sided mirror and we select the side that can reflect the ray \tilde{r} .

Recursive Application of Mirror Reflections

Let us consider the new ray $p_{M_{k_1}}(\tilde{r})$ that is obtained by a physical reflection (Def. 4.1.7) of the initial ray from the mirror M_{k_1} . If the physical reflection is non-empty, we proceed by applying to it the same operations as to the initial ray, realizing the description of a ray tracing scheme.

In particular, the scheme consists in taking the straight segment $s_{\mathcal{M}}(p_{M_{k_1}}(\tilde{r}))$ that ends at mirror $M_{k_2} \in \mathcal{M}$ and continuing the ray reflection there. Since the physical reflection $p_{M_{k_1}}(\tilde{r})$ is not empty, the straight segment $s_{\mathcal{M}}(p_{M_{k_1}}(\tilde{r}))$ is defined and is certainly not empty.

A recursive application of this operation results in a sequence of straight segments: $s_{\mathcal{M}}(\tilde{r})$, $s_{\mathcal{M}}(p_{M_{k_1}}(\tilde{r}))$, $s_{\mathcal{M}}(p_{M_{k_2}M_{k_1}}(\tilde{r}))$, ..., $s_{\mathcal{M}}(p_{M_{k_L} \dots M_{k_2}M_{k_1}}(\tilde{r}))$, ...; and in a corresponding sequence of mirrors that were intersected by the bouncing ray: M_{k_1} , M_{k_2} , ..., M_{k_L} ,

This sequence is finite if at some step l we have: $p_{M_{k_1} \dots M_{k_2} M_{k_1}}(\tilde{r}) = \emptyset$, i.e. the ray hits a mirror from its non-reflecting side or along the mirror's hyperplane, or if $s_{\mathcal{M}}(p_{M_{k_1} \dots M_{k_2} M_{k_1}}(\tilde{r})) = p_{M_{k_1} \dots M_{k_2} M_{k_1}}(\tilde{r})$. The latter condition describes the case where the ray does not intersect any more mirrors.

Definition 4.1.11. Physical ray path. Let $\tilde{r} = \tilde{r}(\mathbf{c}, \mathbf{d})$ be a ray in a mirror system $\mathcal{M} = \{(f_i, F_i) \mid i = 1..N\}$ such that $\mathbf{c} \notin \mathcal{M}_f$. The physical ray path (see Fig. 4.6) in the mirror system \mathcal{M} is the sequence $P_{\mathcal{M}}(\tilde{r}) := (s_{\mathcal{M}}(\tilde{r}), s_{\mathcal{M}}(p_{M_{k_1}}(\tilde{r})), s_{\mathcal{M}}(p_{M_{k_2} M_{k_1}}(\tilde{r})), \dots, s_{\mathcal{M}}(p_{M_{k_L} \dots M_{k_2} M_{k_1}}(\tilde{r})), \dots)$.

Note that $s_{\mathcal{M}}(\tilde{r})$ is defined since $\mathbf{c} \notin \mathcal{M}_f$.

Corollary 1. *The physical ray path $P_{\mathcal{M}}(\tilde{r})$ is C^0 continuous.*

The sequence of straight segments is end-to-beginning connected, i.e. the closure of the end of one segment is the beginning of the next segment in the sequence. The physical ray path describes the bouncing path of the ray inside the mirror system. The first segment is the straight segment between the ray origin \mathbf{c} and mirror M_{k_1} . The following straight segments are connecting mirrors M_{k_i} and $M_{k_{i+1}}$, $i = 1, 2, \dots, L, \dots$. If the sequence $P_{\mathcal{M}}(\tilde{r})$ is finite, it is possible, that the last segment is semi-infinite, i.e. the final ray is not intersecting any more mirrors, or finite, i.e. a mirror back-side is being hit or the final ray intersects a mirror along its hyperplane.




Figure 4.6: Ray propagation in a mirror system: The mirror system \mathcal{M} consists of three mirrors: M_1 , M_2 , and M_3 . The ray $\tilde{r} = \tilde{r}(\mathbf{c}_0, \mathbf{d}_0)$ continues to infinity (dashed lines). Its straight segment $s_{\mathcal{M}}(\tilde{r}(\mathbf{c}_0, \mathbf{d}_0)) = [\mathbf{c}_0; \mathbf{a}_1]$ is determined by M_1 . The physical reflection $p_{M_1}(\tilde{r}(\mathbf{c}_0, \mathbf{d}_0))$ is the original ray reflected from mirror M_1 whereas its straight segment $s_{\mathcal{M}}(p_{M_1}(\tilde{r}(\mathbf{c}_0, \mathbf{d}_0))) = [\mathbf{a}_1; \mathbf{a}_2]$ is only running between M_1 and M_2 . The physical ray path sequence is $P_{\mathcal{M}}(\tilde{r})$: $([\mathbf{c}_0; \mathbf{a}_1], [\mathbf{a}_1; \mathbf{a}_2], [\mathbf{a}_2; \mathbf{a}_3], [\mathbf{a}_3; \mathbf{a}_4], \tilde{r}(\mathbf{a}_4, \mathbf{d}))$.

4.1.2 Ray Unfolding

In the following, we describe ray unfolding, a procedure to turn the physical ray path into a straight line, while keeping track of the required transformations. Often, the other geometric objects in the space are transformed as well. We will use this operation extensively in the remainder of the thesis. As mentioned in Section 3.1, unfolding is usually applied to Coxeter geometries, yielding perfectly overlapping tilings of the surrounding space, or in finite reflection sequences such that overlap is avoided for non-Coxeter geometries.

The goal of the following sections is to introduce an unfolding scheme that is applicable to arbitrary planar mirror systems. We start by defining a per-ray unfolding procedure, later extending the arguments to ray groups with common properties.

Ray Decomposition




Figure 4.7: An example of the decomposition of the ray $\tilde{r} = \tilde{r}(\mathbf{c}, \mathbf{d})$ from its physical ray path. Here, the red semi-interval is the straight segment $s_{\mathcal{M}}(\tilde{r})$ and the green half-line is the ray's physical reflection from mirror M_{K_1} with normal \mathbf{n}_{K_1} and face f_{K_1} .

Lemma 3. *Suppose for a given mirror system $\mathcal{M} = \{M_i = (f_i, F_i) \mid i = 1..N\}$ and a ray $\tilde{r} = \tilde{r}(\mathbf{c}, \mathbf{d}) \mid \mathbf{c} \notin \mathcal{M}_f$, we have the physical ray path sequence: $P_{\mathcal{M}}(\tilde{r}) = (s_{\mathcal{M}}(\tilde{r}), s_{\mathcal{M}}(p_{M_{k_1}}(\tilde{r})), s_{\mathcal{M}}(p_{M_{k_2}M_{k_1}}(\tilde{r})), \dots$,*

$s_{\mathcal{M}}(p_{M_{k_L} \dots M_{k_2} M_{k_1}}(\tilde{r})), \dots)$ with the corresponding mirror sequence: $(M_{k_1}, M_{k_2}, M_{k_3}, \dots, M_{k_L}, \dots)$.

If the physical reflection $p_{M_{k_1}}(\tilde{r}) \neq \emptyset$, the ray \tilde{r} is fully covered by the straight segment $s_{\mathcal{M}}(\tilde{r})$ and the mirrored physical reflection $m_{M_{k_1}}(p_{M_{k_1}}(\tilde{r}))$ of that ray. Moreover, the straight segment and the mirrored physical reflection are disjoint:

$$\tilde{r} = s_{\mathcal{M}}(\tilde{r}) \sqcup m_{M_{k_1}}(p_{M_{k_1}}(\tilde{r})),$$

where \sqcup denotes the union of disjoint sets, see Fig. 4.7.

Proof. According to the conditions of the Lemma, $p_{M_{k_1}}(\tilde{r}) \neq \emptyset$, i.e. a physical reflection exists. This, together with $\mathbf{c} \notin \mathcal{M}_f$, implies that the straight segment $s_{\mathcal{M}}(\tilde{r})$ is non-empty, and, moreover, $s_{\mathcal{M}}(\tilde{r}) = (\tilde{r} \setminus r_{k_1}^{\sim})$, where k_1 is the index of the first intersection with a mirror (M_{k_1}) and $r_{k_1}^{\sim}$ is the semi-infinite interval of \tilde{r} that is starting at the intersection point with the mirror and going to infinity as defined in Def. 4.1.10. In particular, since the intersection point exists, $r_{k_1}^{\sim} = \tilde{r}(\tilde{r} \cap f_{k_1}, \mathbf{d})$.

Moreover, from Def. 4.1.7, the physical reflection is $p_{M_{k_1}}(\tilde{r}) = m_{M_{k_1}}(\tilde{r}(\tilde{r} \cap f_{k_1}, \mathbf{d}))$, i.e., $p_{M_{k_1}}(\tilde{r}) = m_{M_{k_1}}(r_{k_1}^{\sim})$. Undoing the reflection via $m_{M_{k_1}}(p_{M_{k_1}}(\tilde{r}))$, we have $r_{k_1}^{\sim} = m_{M_{k_1}}(p_{M_{k_1}}(\tilde{r}))$ and therefore $s_{\mathcal{M}}(\tilde{r}) = \tilde{r} \setminus m_{M_{k_1}}(p_{M_{k_1}}(\tilde{r}))$. The latter statement is equivalent to the assertion in the Lemma. \square

If we recursively apply Lemma 3 to the second part of the right side of the equality (i.e. to the $p_{M_{k_1}}(\tilde{r})$, $p_{M_{k_2} M_{k_1}}(\tilde{r})$, \dots , $p_{M_{k_L} \dots M_{k_2} M_{k_1}}(\tilde{r})$, \dots), and none of the sets $p_{M_{k_1}}(\tilde{r})$, $p_{M_{k_2} M_{k_1}}(\tilde{r})$, \dots , $p_{M_{k_L} \dots M_{k_2} M_{k_1}}(\tilde{r})$, \dots is equal to \emptyset , we obtain the following ray decomposition \tilde{r}_d' :

$$\begin{aligned} \tilde{r}_d' &= s_{\mathcal{M}}(\tilde{r}) \sqcup m_{M_{k_1}}(p_{M_{k_1}}(\tilde{r})) \\ &= s_{\mathcal{M}}(\tilde{r}) \sqcup m_{M_{k_1}}[s_{\mathcal{M}}(p_{M_{k_1}}(\tilde{r})) \\ &\quad \sqcup m_{M_{k_2}}(p_{M_{k_2} M_{k_1}}(\tilde{r}))] \\ &= s_{\mathcal{M}}(\tilde{r}) \sqcup m_{M_{k_1}}(s_{\mathcal{M}}(p_{M_{k_1}}(\tilde{r}))) \\ &\quad \sqcup m_{M_{k_1} M_{k_2}}(p_{M_{k_2} M_{k_1}}(\tilde{r})) \\ &= \dots \\ &= s_{\mathcal{M}}(\tilde{r}) \sqcup m_{M_{k_1}}(s_{\mathcal{M}}(p_{M_{k_1}}(\tilde{r}))) \\ &\quad \sqcup m_{M_{k_1} M_{k_2}}(s_{\mathcal{M}}(p_{M_{k_2} M_{k_1}}(\tilde{r}))) \\ &\quad \sqcup m_{M_{k_1} M_{k_2} \dots M_{k_L}}(s_{\mathcal{M}}(p_{M_{k_L} \dots M_{k_2} M_{k_1}}(\tilde{r}))) \\ &\quad \sqcup \dots \end{aligned}$$

In case of an infinitely propagating ray (but not necessarily infinitely bouncing), the ray \tilde{r} equals the decomposition \tilde{r}_d' : $\tilde{r} = \tilde{r}_d'$ since all points of the ray are assigned uniquely to one of the unfolded straight segments.

This follows from the fact, that faces of mirrors in the mirror system \mathcal{M} are closed and bounded disjoint sets. Thus, there is a minimal positive distance d between them [Berg08]. The length of any straight segment is not smaller than this distance d . In particular, segments cannot become infinitely small. Therefore, when the ray is infinitely bouncing, the unfolded segments eventually cover the ray and any point of \tilde{r} belongs to a corresponding disjoint element of the ray decomposition.

We refer to the unfolded straight segments, i.e. to the members of the disjoint union, as *unfolded segments*. So far, we are within the limits of validity of Lemma 3.

If, however, at some point the physical ray path hits a mirror in such a way that the physical reflection no longer exists ($p_{M_{k_L} \dots M_{k_2} M_{k_1}}(\tilde{r}) = \emptyset$), then the unfolded ray \tilde{r}'_d stops after a finite distance. In this case, we artificially add a last semi-infinite segment to extend the decomposition to infinity.

$$\tilde{r}_d := \begin{cases} \tilde{r}'_d & \forall L \in \mathbb{N}, p_{M_{k_L} \dots M_{k_2} M_{k_1}}(\tilde{r}) \neq \emptyset \\ \tilde{r}'_d \sqcup (\tilde{r} \setminus \tilde{r}'_d) & \text{otherwise.} \end{cases} \quad (4.1)$$

With these prerequisites, we can state that, for a starting point $\mathbf{c} \notin \mathcal{M}_f$ outside the mirrors, there exists a unique decomposition of the ray $\tilde{r} = \tilde{r}_d$ into classes of points corresponding to a unique mirroring sequence or to a rest class.

Partitioning of the Space Along the Ray

Definition 4.1.12. Chamber and chamber transformation. For a given mirror system \mathcal{M} and a sequence of mirrors $M_1, M_2, \dots, M_N \in \mathcal{M}$, we define a *chamber* $C_{M_1 M_2 \dots M_N}$ to be a transformed version of the whole space \mathbb{R}^n including the contained objects. The associated *chamber transformation* $m_{M_1 M_2 \dots M_N} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is determined by the mirror sequence $M_1, M_2, \dots, M_N \in \mathcal{M}$. Then $C_{M_1 M_2 \dots M_N} := m_{M_1 M_2 \dots M_N}(\mathbb{R}^n, *)$. We refer to N as the chamber level. The associated space without objects is $\mathbb{R}_{M_1 M_2 \dots M_N}^n$. We denote the inverse of a chamber $C_{M_1 M_2 \dots M_N}$ as the chamber with reversed sequence of mirrors: $C_{M_1 M_2 \dots M_N}^{-1} := C_{M_N M_{N-1} \dots M_1}$.

If the sequence of mirrors is empty, the chamber for it is defined to be $C_0 := (\mathbb{R}_0^n, *)$ with level 0 which is \mathbb{R}^n itself including all objects. C_0 is called the *base chamber*, while \mathbb{R}_0^n is the Euclidean space containing the base chamber. Chambers with a higher level than zero are called *virtual chambers*.

\mathbb{R}_0^n is the base space in which the mirror system, its points, rays, hyperplanes, etc. are defined. Together, these entities constitute the base

chamber C_0 . Our goal will be to unfold the physical ray paths that occur within the base chamber into the ambient \mathbb{R}_0^n .

Important note: Two chambers are equal iff the corresponding sequences of mirrors are the same. Two different sequences of reflections may lead to the same transformation of the ambient space \mathbb{R}^n . The corresponding chambers will still be considered as being different. We use this property to associate the corresponding chamber transformation with the chamber with the following notation: $m_{C_{M_1 M_2 \dots M_N}}(*) := m_{M_1 M_2 \dots M_N}(*)$, where the chamber $C_{M_1 M_2 \dots M_N}$ is parameterizing the chamber transformation. We denote its inverse as $m_{C_{M_1 M_2 \dots M_N}}^{-1}(*) := m_{M_N M_{N-1} \dots M_1}(*)$.

Definition 4.1.13. Set of chambers and Extended set of chambers.

For a given mirror system \mathcal{M} , consider the *set of chambers* $C_{\mathcal{M}}$, i.e. all possible chambers $C_{\mathcal{M}} := \{C_0, C_{M_1}, C_{M_2}, \dots, C_{M_{i_1} M_{i_2} \dots M_{i_j}}, \dots\}$ and the special element C_{\emptyset} called the *empty chamber*. The empty chamber has no associated \mathbb{R}^n . The set $\overline{C}_{\mathcal{M}} = C_{\mathcal{M}} \cup \{C_{\emptyset}\}$ is defining the *extended set of chambers*. An element of the extended set of chambers is referred to as an *extended chamber*.

Corollary 2. *If the number of mirrors in mirror system \mathcal{M} is finite, the set $C_{\mathcal{M}}$, and, consequently, the set $\overline{C}_{\mathcal{M}}$ from Def. 4.1.13 is countable.*

Definition 4.1.14. Chamber function. For a given mirror system \mathcal{M} and a point $\mathbf{c} \in (\mathbb{R}_0^n \setminus \mathcal{M}_f)$, we define a *chamber function* $f_{\mathcal{M}, \mathbf{c}} : \mathbb{R}_0^n \mapsto \overline{C}_{\mathcal{M}}$ as follows:

1. $f_{\mathcal{M}, \mathbf{c}}(\mathbf{c}) = C_0$;

Let $\mathbf{p} \in \mathbb{R}_0^n$ and $\mathbf{p} \neq \mathbf{c}$, and consider the ray $\tilde{r} = \tilde{r}(\mathbf{c}, \mathbf{p} - \mathbf{c})$ with the decomposition (4.1). Then:

2. $f_{\mathcal{M}, \mathbf{c}}(\mathbf{p}) = C_0$, if $\mathbf{p} \in s_{\mathcal{M}}(\tilde{r})$;

3. $f_{\mathcal{M}, \mathbf{c}}(\mathbf{p}) = C_{M_{k_1} M_{k_2} M_{k_3} \dots M_{k_L}}$, if $\mathbf{p} \in m_{M_{k_1} M_{k_2} M_{k_3} \dots M_{k_L}}(s_{\mathcal{M}}(p_{M_{k_L} \dots M_{k_3} M_{k_2} M_{k_1}}(\tilde{r})))$;

4. $f_{\mathcal{M}, \mathbf{c}}(\mathbf{p}) = C_{\emptyset}$, if $f_{\mathcal{M}, \mathbf{c}}(\mathbf{p})$ is not assigned by the previous conditions.

As can be seen from the definition, the chamber function $f_{\mathcal{M}, \mathbf{c}} : \mathbb{R}_0^n \mapsto \overline{C}_{\mathcal{M}}$ is defined for any fixed mirror system \mathcal{M} and any fixed point \mathbf{c} that is not part of that mirror system. We refer to \mathbf{c} as an *observation point*.

Lemma 4. *The chamber function $f_{\mathcal{M}, \mathbf{c}} : \mathbb{R}_0^n \mapsto \overline{C}_{\mathcal{M}}$ is a function.*

Proof. We need to prove that $\forall \mathbf{p} \in \mathbb{R}_0^n \Rightarrow (\exists! C \in \overline{C}_{\mathcal{M}} \mid f_{\mathcal{M}, \mathbf{c}}(\mathbf{p}) = C)$. Existence is following directly from the function's definition. We need to prove the uniqueness: If $\mathbf{p} = \mathbf{c} \Rightarrow C = C_0$ and the assignment is unique.

If $\mathbf{p} \neq \mathbf{c} \Rightarrow \exists! \tilde{r} = \tilde{r}(\mathbf{c}, \mathbf{p} - \mathbf{c})$. For that ray we have the uniquely defined ray path sequence $P_{\mathcal{M}}(\tilde{r})$ in the mirror system \mathcal{M} and therefore the uniquely defined decomposition (4.1). Because decomposition (4.1) consists of non-intersecting ray intervals, the point \mathbf{p} belongs to exactly one of them. According to the definition of function $f_{\mathcal{M}, \mathbf{c}}$ this results in the unique assignment of a single value from $\overline{C}_{\mathcal{M}}$. \square

The chamber function $f_{\mathcal{M}, \mathbf{c}}$ of the mirror system \mathcal{M} with the observation point \mathbf{c} is assigning an extended chamber (i.e. an element of $\overline{C}_{\mathcal{M}}$) to each point of \mathbb{R}_0^n . In particular, for a given ray $\tilde{r} = \tilde{r}(\mathbf{c}, \mathbf{d}) \in \mathbb{R}_0^n$, it assigns extended chambers to all of the ray's points according to decomposition (4.1) of the ray \tilde{r} .

Definition 4.1.15. Ray partitioning. The *ray partitioning* $\Pi_{\mathcal{M}}(\tilde{r})$ is an application of the chamber function $f_{\mathcal{M}, \mathbf{c}}$ to all points of ray \tilde{r} . Let \mathcal{M} be a mirror system, $\tilde{r} = \tilde{r}(\mathbf{c}, \mathbf{d})$, $\mathbf{c} \in (\mathbb{R}_0^n \setminus \mathcal{M}_f)$ a ray with origin \mathbf{c} , and let $f_{\mathcal{M}, \mathbf{c}}$ share the origin \mathbf{c} and the mirror system \mathcal{M} with the ray.

Then, the ray partitioning $\Pi_{\mathcal{M}}(\tilde{r}) := \{(\mathbf{p}, f_{\mathcal{M}, \mathbf{c}}(\mathbf{p})) \mid \mathbf{p} \in \tilde{r}\}$ is a set of point/chamber pairs defined for all points of the ray \tilde{r} .

Theorem 1. Ray Unfolding

For any ray $\tilde{r} = \tilde{r}(\mathbf{c}, \mathbf{d})$ where $\mathbf{c} \in (\mathbb{R}_0^n \setminus \mathcal{M}_f)$, without points $\{\mathbf{p} \mid f_{\mathcal{M}, \mathbf{c}}(\mathbf{p}) = C_{\emptyset}\}$,

$$P_{\mathcal{M}}(\tilde{r}) = \{m_{f_{\mathcal{M}, \mathbf{c}}(\mathbf{p})}^{-1}(\mathbf{p}) \mid \mathbf{p} \in \tilde{r}\}.$$

The theorem states that any physical ray path $P_{\mathcal{M}}(\tilde{r})$ can be represented by a straight ray $\tilde{r} \subset \mathbb{R}_0^n$ with a partitioning $\Pi_{\mathcal{M}}(\tilde{r})$ and vice versa. Moreover, the physical ray path is obtained by applying the spatially varying inverse chamber transformation $m_{f_{\mathcal{M}, \mathbf{c}}(\mathbf{p})}^{-1}(\mathbf{p})$ to the straight ray.

The proof is based on the construction of a bijection between the straight segments and the unfolded segments of a ray \tilde{r} .

Proof. According to Def. 4.1.11, the physical ray path consists of straight segments $s_{\mathcal{M}}(p_{M_{k_L} \dots M_{k_2} M_{k_1}}(\tilde{r}))$. Their corresponding unfolded segments in decomposition (4.1) are given by $m_{M_{k_1} M_{k_2} \dots M_{k_L}}(s_{\mathcal{M}}(p_{M_{k_L} \dots M_{k_2} M_{k_1}}(\tilde{r})))$.

Consider one particular such straight segment $a = s_{\mathcal{M}}(p_{M_{a_L} \dots M_{a_2} M_{a_1}}(\tilde{r}))$ and its corresponding unfolded segment $m_{M_{a_1} M_{a_2} \dots M_{a_L}}(a) = m_{M_{a_1} M_{a_2} \dots M_{a_L}}(s_{\mathcal{M}}(p_{M_{a_L} \dots M_{a_2} M_{a_1}}(\tilde{r})))$. According to Def. 4.1.14, a gets assigned a constant chamber $C_{M_{a_1} M_{a_2} \dots M_{a_L}}$. Therefore, the corresponding chamber transformation $m_{C_{M_{a_1} M_{a_2} \dots M_{a_L}}}(\ast) = m_{M_{a_1} M_{a_2} \dots M_{a_L}}(\ast)$ is constant and its inverse can be applied to the whole unfolded segment $m_{M_{a_1} M_{a_2} \dots M_{a_L}}(a)$, yielding

$$\begin{aligned}
m_{C_{M_{a_1} M_{a_2} \dots M_{a_L}}}^{-1} (m_{M_{a_1} M_{a_2} \dots M_{a_L}}(a)) &= \\
m_{M_{a_1} M_{a_2} \dots M_{a_L}}^{-1} (m_{M_{a_1} M_{a_2} \dots M_{a_L}}(a)) &= \\
m_{M_{a_L} M_{a_{L-1}} \dots M_{a_1}} (m_{M_{a_1} M_{a_2} \dots M_{a_L}}(a)) &= \\
m_{M_{a_L} M_{a_{L-1}} \dots M_{a_1} M_{a_1} M_{a_2} \dots M_{a_L}}(a) &= a
\end{aligned}$$

Applying this construction of a bijection between the straight segment a and its corresponding unfolded segment $m_{M_{a_1} M_{a_2} \dots M_{a_L}}(a)$ to all straight segments of the physical path completes the proof. \square

Note that the bijection constructed above is only valid per-straight segment, i.e. that the complete physical ray path can have self-intersections. Thus, in general, there is no bijection between the ray $\tilde{r} \setminus C_\emptyset$ and the full set of points of the physical ray path.

Corollary 3. *The mirror system \mathcal{M} , as an object of \mathbb{R}_0^n is also affected by the chamber transformations according to the partitioning of ray \tilde{r} . In particular, the sequence of mirrors $M_{k_1}, M_{k_2}, \dots, M_{k_L}, \dots$, where reflections of the ray occur for a given ray path sequence $P_{\mathcal{M}}(\tilde{r})$ is transforming into the following sequence of virtual mirrors: $m_{M_{k_1}}(M_{k_1}), m_{M_{k_1} M_{k_2}}(M_{k_2}), \dots, m_{M_{k_1} M_{k_2} \dots M_{k_L}}(M_{k_L})$, the corresponding chambers being $C_{M_{k_1}}, C_{M_{k_1} M_{k_2}}, \dots, C_{M_{k_1} M_{k_2} \dots M_{k_L}}, \dots$*

Corollary 4. *From Corollary 3 we see, that two sequential unfolded segments of the unfolded ray are bordered by a unique virtual mirror, that is defined only by the corresponding sequence of reflections.*

Ray straightening is a very useful operation: 1) In analyzing multi-bounce mirror systems, it allows us to consider unfolded straight rays instead of the chain of straight segments that constitute the physical path, and 2) the associated chamber subdivision of the surrounding space provides us with a proper relative environment (sets of points, mirrors etc.) for any given valid point on the ray.

4.1.3 Space Partitioning

For a given mirror system \mathcal{M} and a point $\mathbf{c} \in (\mathbb{R}_0^n \setminus \mathcal{M}_f)$, if we take any point $\mathbf{p} \in \mathbb{R}_0^n$ with $\mathbf{p} \neq \mathbf{c}$, this point will be on the ray $\tilde{r} = \tilde{r}(\mathbf{c}, \mathbf{p} - \mathbf{c})$, with the origin at \mathbf{c} . For this ray the ray partitioning $\Pi_{\mathcal{M}}(\tilde{r})$ is defined. If the point $\mathbf{p} = \mathbf{c}$, then for any ray with this point as the origin, the ray partitioning will

associate it with chamber C_0 . Thus, all the points in \mathbb{R}_0^n can be classified into extended chambers by the ray partitioning procedure. In other words, the chamber function, in conjunction with the ray partitioning is creating a partitioning of the full space \mathbb{R}_0^n into sets, belonging to chambers from $\overline{C}_{\mathcal{M}}$. In this section we discuss the structure of such partitioning.

Definition 4.1.16. Space partitioning.

A *space partitioning* $\Pi_{\mathcal{M}}^*(\mathbf{c}_0)$ of \mathbb{R}_0^n with respect to a mirror system \mathcal{M} and point $\mathbf{c}_0 \in (\mathbb{R}_0^n \setminus \mathcal{M}_f)$ is the set of ray partitionings for all rays with a common origin \mathbf{c}_0 : $\Pi_{\mathcal{M}}^*(\mathbf{c}_0) = \{\Pi_{\mathcal{M}}(\tilde{r}(\mathbf{c}, \mathbf{d})) \mid \mathbf{c} = \mathbf{c}_0\}$.

Just as the ray partitioning, Def. 4.1.15, the space partitioning $\Pi_{\mathcal{M}}^*(\mathbf{c}_0)$ is a set of point/chamber pairs, this time covering the whole of \mathbb{R}_0^n , i.e. $\forall \mathbf{p} \in \mathbb{R}_0^n \Rightarrow \exists!(\mathbf{p}, C) \in \Pi_{\mathcal{M}}^*(\mathbf{c}_0)$. We will speak about $\Pi_{\mathcal{M}}^*(\mathbf{c}_0)$ assigning a chamber C to the point \mathbf{p} , or, alternatively, a chamber C being associated with point \mathbf{p} by $\Pi_{\mathcal{M}}^*(\mathbf{c}_0)$.

Structure of Space Partitioning $\Pi_{\mathcal{M}}^*(\mathbf{c}_0)$.

Consider a mirror system \mathcal{M} and a point $\mathbf{c}_0 \in (\mathbb{R}_0^n \setminus \mathcal{M}_f)$. If we consider all possible unfolded rays from \mathbf{c}_0 up to the $N \geq 0$ first levels of reflection, including N , we obtain a set of points in \mathbb{R}_0^n . We call this set of points *N-unfolding*. Practically, the decomposition (4.1) can be evaluated up to and including N -unfolded segments for all rays through \mathbf{c}_0 . The following properties for N -unfoldings hold:

Lemma 5.

1. *The chamber, associated with a given point \mathbf{a} in the partitioning $\Pi_{\mathcal{M}}^*(\mathbf{c}_0)$ is uniquely defining the chamber sequence when traveling along the ray from point \mathbf{c}_0 to point \mathbf{a} . In particular, if $f_{\mathcal{M}, \mathbf{c}_0}(\mathbf{a}) = C_{M_{k_1} M_{k_2} \dots M_{k_L}}$, this sequence is $C_0, C_{M_{k_1}}, C_{M_{k_1} M_{k_2}}, \dots, C_{M_{k_1} M_{k_2} \dots M_{k_L}}$. Compare Def. 4.1.14 and decomposition (4.1).*
2. $\forall N \geq 0 \Rightarrow$ *An N -unfolding is a subset of an $(N + 1)$ -unfolding.*
3. *Since all rays \tilde{r} pass through \mathbf{c}_0 and the decomposition (4.1) is connected, an N -unfolding is star-shaped.*

From the lemma above we see, that the space partitioning $\Pi_{\mathcal{M}}^*(\mathbf{c}_0)$ is not just an arbitrary subdivision of the space, but it has a tree-like star-shaped structure, which grows with an increasing number of reflections, see Fig. 4.8.




Figure 4.8: 0-, 1- and 2-unfoldings of the mirror system $\mathcal{M} = \{M_1, M_2, M_3\}$ from the point \mathbf{c}_0 . Green, violet and pink areas depict visible areas (from \mathbf{c}_0) of chambers C_0 , C_{M_i} and $C_{M_i M_j}$, correspondingly.

Visibility in a Space Partition

Definition 4.1.17. Visible point.

Let $\mathbf{c} \in (\mathbb{R}_0^n \setminus \mathcal{M}_f)$, $\mathbf{a} \in \mathbb{R}_0^n$ and \mathcal{M} be a mirror system. The point \mathbf{a} is called *visible* (*visible in direction* \mathbf{d}) from the point \mathbf{c} iff there is a ray $\tilde{r} = \tilde{r}(\mathbf{c}, \mathbf{d})$ such that $\mathbf{a} \in P_{\mathcal{M}}(\tilde{r})$.

In other words, the point is visible from the given point of observation \mathbf{c} iff there is a bouncing path from point \mathbf{c} to it.

If a point \mathbf{a} is visible from \mathbf{c} in \mathcal{M} , then $\exists \tilde{r} = \tilde{r}(\mathbf{c}, \mathbf{d}) \mid \mathbf{a} \in P_{\mathcal{M}}(\tilde{r})$, where $P_{\mathcal{M}}(\tilde{r}) = (s_{\mathcal{M}}(\tilde{r}), s_{\mathcal{M}}(p_{M_{k_1}}(\tilde{r})), s_{\mathcal{M}}(p_{M_{k_2} M_{k_1}}(\tilde{r})), \dots, s_{\mathcal{M}}(p_{M_{k_L} \dots M_{k_2} M_{k_1}}(\tilde{r})), \dots)$. Consequently $\exists K \in \mathbb{N} \cup \{0\} \mid \mathbf{a} \in s_{\mathcal{M}}(p_{M_{k_K} \dots M_{k_2} M_{k_1}}(\tilde{r}))$. Then, according to the space partitioning the straight segment $s_{\mathcal{M}}(p_{M_{k_K} \dots M_{k_2} M_{k_1}}(\tilde{r}))$ has a corresponding unfolded segment on the ray \tilde{r} in the partitioned space. Moreover, it belongs to the chamber $C_{M_{k_1} M_{k_2} \dots M_{k_K}}$. Thus, point $m_{M_{k_1} M_{k_2} \dots M_{k_K}}(\mathbf{a})$ is on the ray \tilde{r} and assigned by the space partitioning to the chamber $C_{M_{k_1} M_{k_2} \dots M_{k_K}}$.

Because the mirror transformation is bijective, we can apply the inverse argument, i.e. if the point \mathbf{a}' is a part of the ray \tilde{r} and belongs to a chamber $C_{M_{k_1} M_{k_2} \dots M_{k_K}}$, then the point $m_{M_{k_K} \dots M_{k_2} M_{k_1}}(\mathbf{a}')$ is visible.

Note, that the ray \tilde{r} and the number K above are not necessarily unique. In other words, a point \mathbf{a} can be reachable from the point of observation \mathbf{c} via different rays and different sequences of reflections.

Definition 4.1.18. Visible point in a chamber. Let $\mathbf{c} \in (\mathbb{R}_0^n \setminus \mathcal{M}_f)$, $\mathbf{a} \in \mathbb{R}_0^n$ and \mathcal{M} be a mirror system. The point \mathbf{a} is called visible in a non-

empty chamber $C_{M_{k_1}M_{k_2}\dots M_{k_K}}$ from point \mathbf{c} iff there is a point \mathbf{a}' in \mathbb{R}_0^n which is associated with chamber $C_{M_{k_1}M_{k_2}\dots M_{k_K}}$ by the space partitioning $\Pi_{\mathcal{M}}^*(\mathbf{c})$, i.e. $\exists \mathbf{a}' | m_{M_{k_K}\dots M_{k_2}M_{k_1}}(\mathbf{a}') = \mathbf{a}$.

Note: We refer to sets of points as being visible in a given chamber iff all the points of the set are visible in that chamber, *ignoring self-occlusion*.

From the definition we have, that if the point is visible in a given chamber, it is reachable by a sequence of bounces from the point \mathbf{c} , see also Fig. 4.8. The sequence of mirror bounces corresponds to the chamber.

Lemma 6. Independence of chambers from the observation point.

In a mirror system \mathcal{M} , all chambers $C \in \bar{C}_{\mathcal{M}} \setminus C_{\emptyset}$ are independent of the position of the point of observation \mathbf{c}_0 and defined only by the mirrors in the system \mathcal{M} . The point of observation \mathbf{c}_0 is only influencing the area that is visible in a given chamber as seen from that point. Changing the position of the observation point can make some chambers visible, or invisible. Some chambers may never be visible independent of the position of point \mathbf{c}_0 .

Proof. The proof is a direct consequence of the definition of a chamber (Def. 4.1.12) and the definition of the space partitioning (Def. 4.1.16). \square

4.2 Continuity of The Space Partitioning

An important property of the space partitioning defined in the previous section is its continuity with respect to changes in observation point. In the following we will formally define and prove this continuity, followed by a discussion of its practical importance.

But let us first introduce some necessary definitions and notations.

Definition 4.2.1. Open ball. Let $r \in \mathbb{R}$, $r > 0$ and $\mathbf{a} \in \mathbb{R}^n$. Then an *open ball* $S_o(r, \mathbf{a})$ of radius r with center \mathbf{a} is $S_o(r, \mathbf{a}) = \{\mathbf{a}' \mid \|\mathbf{a}' - \mathbf{a}\| < r\}$.

Definition 4.2.2. Closed ball. Let $r \in \mathbb{R}$, $r > 0$ and $\mathbf{a} \in \mathbb{R}^n$. Then a *closed ball* $S_c(r, \mathbf{a})$ of radius r with center \mathbf{a} is $S_c(r, \mathbf{a}) = \{\mathbf{a}' \mid \|\mathbf{a}' - \mathbf{a}\| \leq r\}$.

Definition 4.2.3. Distance to the mirror system. Let $\mathcal{M} = \{M_1, M_2, \dots, M_N\}$ be a mirror system in \mathbb{R}^n and S a set of points in \mathbb{R}^n . Then the distance between the mirror system \mathcal{M} and the set S is given by

$$d(\mathcal{M}, S) = \inf_{\mathbf{m} \in \mathcal{M}_f, \mathbf{s} \in S} d(\mathbf{m}, \mathbf{s}).$$

4.2.1 Continuity of the Space Partitioning with Respect to Changes in Observation Point

Definition 4.2.4. Continuity with respect to observation point. A mirror system \mathcal{M} is called *chamber-continuous with respect to a given*




Figure 4.9: Illustrating observer continuity of a mirror system. If \mathbf{a} is visible from \mathbf{c}_0 through a common mirror sequence together with an open ball, then there exists an open ball around \mathbf{c}_0 such that \mathbf{a} is visible from any of its points with the same mirror sequence. The proof is based on an unfolding of the situation, see Fig. 4.10.

chamber $C_{M_{k_1} M_{k_2} \dots M_{k_K}}$ and a point of observation $\mathbf{c}_0 \in (\mathbb{R}^n \setminus \mathcal{M}_f)$ iff $(\forall r > 0, \forall \mathbf{a} \in \mathbb{R}^n \mid \forall \mathbf{a}' \in S_o(r, \mathbf{a}) \Rightarrow \mathbf{a}'$ is visible from \mathbf{c}_0 in the chamber $C_{M_{k_1} M_{k_2} \dots M_{k_K}}) \Rightarrow (\exists r_1 > 0 \mid (\forall \mathbf{c}' \in S_o(r_1, \mathbf{c}_0) \Rightarrow$

a) $\mathbf{c}' \in (\mathbb{R}^n \setminus \mathcal{M}_f)$, and

b) The point \mathbf{a} is visible in the chamber $C_{M_{k_1} M_{k_2} \dots M_{k_K}}$ from the point \mathbf{c}').

The mirror system is called *continuous with respect to change in the observation point* iff it is chamber-continuous for all chambers $\overline{C}_{\mathcal{M}} \setminus C_{\emptyset}$.

The setting is illustrated in Fig. 4.9: Chamber-continuity describes the situation where an open ball around point \mathbf{a} that is visible from \mathbf{c}_0 under a common mirror sequence, implies the existence of an open ball around \mathbf{c}_0 that is visible from point \mathbf{a} through the same mirror sequence as before.

The definition therefore describes the idea that a mirror system is continuous with respect to a point of observation iff the visibility of any point in the system is smooth under changes of the observation point.

In order to proof that any mirror system \mathcal{M} is continuous with respect to the observer position (Theorem 2), we need a number of technical prerequisites.

An Open Ball that is Fully Visible in a Common Chamber does not Contain Mirror Points

The first of these is Lemma 8 which states that an open ball that is completely visible in a common chamber has no mirror intersections. To proof this, we need the help of

Lemma 7. *In a mirror system \mathcal{M} , suppose two points \mathbf{q}_1 and \mathbf{q}_2 are visible from an observation point $\mathbf{c}_0 \in (\mathbb{R}^n \setminus \mathcal{M}_f)$ in a chamber $C = C_{M_{k_1}M_{k_2}\dots M_{k_K}}$, ($K \geq 1$), i.e. a chamber C that is different from the base chamber C_0 . The point \mathbf{q}_1 belongs to the straight segment $s_1 = s_{\mathcal{M}}(p_{M_{k_K}M_{k_{K-1}}\dots M_{k_1}}(\tilde{r}_1))$ of the ray $\tilde{r}_1 = \tilde{r}(\mathbf{c}_0, \mathbf{q}'_1 - \mathbf{c}_0)$ with the origin at point \mathbf{c}_0 and with point $\mathbf{q}'_1 = m_{M_{k_1}M_{k_2}\dots M_{k_K}}(\mathbf{q}_1)$ (see Theorem 1). Let l_1 be a line such that $s_1 \subset l_1$. Then, if*

$$\mathbf{q}_2 \in l_1 \Rightarrow \mathbf{q}_2 \in s_1.$$

Proof. If the point $\mathbf{q}_2 \in l_1$ and $\mathbf{q}_2 \notin s_1$, then, since \mathbf{q}_2 is visible in the same chamber as \mathbf{q}_1 , we have: \mathbf{q}_2 belongs to $s_2 = s_{\mathcal{M}}(p_{M_{k_K}M_{k_{K-1}}\dots M_{k_1}}(\tilde{r}_2))$, where $\tilde{r}_2 = \tilde{r}(\mathbf{c}_0, \mathbf{q}'_2 - \mathbf{c}_0)$, $\mathbf{q}'_2 = m_{M_{k_1}M_{k_2}\dots M_{k_K}}(\mathbf{q}_2)$.

There are two possible cases:

- $s_1 \subset l_1$ and $s_2 \subset l_1$.

In this case \tilde{r}_1 and \tilde{r}_2 belong to the common line $l'_1 = m_{M_{k_1}M_{k_2}\dots M_{k_K}}(l_1)$, while sharing the common origin \mathbf{c}_0 . Then \tilde{r}_1 is equal to \tilde{r}_2 , or \tilde{r}_1 is directed in the opposite direction of \tilde{r}_2 . The first case leads to the contradictory fact, that $s_1 = s_2$. In the second case the ray partitionings $\Pi_{\mathcal{M}}(\tilde{r}_1)$ and $\Pi_{\mathcal{M}}(\tilde{r}_2)$ cannot contain a common chamber, except for the base chamber C_0 (which is excluded by the Lemma's conditions). This is because the same mirror hyperplane cannot be intersected by the line twice. Therefore, the very first mirror in the mirror sequences for the two opposite rays cannot be the same.

- $s_1 \subset l_1$ and $s_2 \not\subset l_1$.

In this case s_2 belongs to another line: $s_2 \subset l_2 \neq l_1$. Then \tilde{r}_1 and \tilde{r}_2 are parts of two different lines: $l'_1 = m_{M_{k_1}M_{k_2}\dots M_{k_K}}(l_1)$ and $l'_2 = m_{M_{k_1}M_{k_2}\dots M_{k_K}}(l_2)$. These lines are intersecting at the point $\mathbf{q}'_2 = m_{M_{k_1}M_{k_2}\dots M_{k_K}}(\mathbf{q}_2)$. Obviously, this is only possible if $\mathbf{q}'_2 = \mathbf{c}_0$. But then, \mathbf{q}_2 must be associated with the base chamber C_0 which is a contrary to the Lemma's conditions.

□

Lemma 8. *If an open ball is completely visible in a common chamber, it does not contain points of the mirror system: Let $\mathbf{a}' \in S_o(r, \mathbf{a})$ and \mathbf{a}' be visible from \mathbf{c}_0 in chamber $C_{M_{k_1}M_{k_2}\dots M_{k_K}}$, $K \geq 0$, then $S_o(r, \mathbf{a}) \cap \mathcal{M}_f = \emptyset$.*

Proof. We proof the Lemma by contradiction. Suppose, an open ball $S_o(r, \mathbf{a})$ is visible from the point of observation $\mathbf{c}_0 \in (\mathbb{R}^n \setminus \mathcal{M}_f)$ in a chamber $C = C_{M_{k_1}M_{k_2}\dots M_{k_K}}$ ($K \geq 0$). If $\exists i \mid M_{k_i} \cap S_o(r, \mathbf{a}) \neq \emptyset$, then $\exists \mathbf{a}_1 \mid (\mathbf{a}_1 \in S_o(r, \mathbf{a}) \text{ and } \mathbf{a}_1 \in M_{k_i})$. There are two possible cases:

1) Chamber $C = C_0$.

Since $\mathbf{a}_1 \in M_{k_i} \Rightarrow$ for the ray $\tilde{r}_1 = \tilde{r}(\mathbf{c}_0, \mathbf{a}_1 - \mathbf{c}_0)$ we have: $s_{\mathcal{M}}(\tilde{r}_1) \subseteq [\mathbf{c}_0; \mathbf{a}_1]$. However, under partitioning $\Pi_{\mathcal{M}}^*(\mathbf{c}_0)$, \mathbf{a}_1 is associated with chamber C_0 which implies that $\mathbf{a}_1 \in s_{\mathcal{M}}(\tilde{r}_1)$ which is a contradiction.

2) Chamber $C = C_{M_{k_1}M_{k_2}\dots M_{k_K}}$, where $K \geq 1$

For point $\mathbf{a}_1 \in M_{k_i}$ the corresponding straight segment passing through it is given by $s_1 = s_{\mathcal{M}}(p_{M_{k_L}M_{k_{L-1}}\dots M_{k_1}}(\tilde{r}_1))$, where $\tilde{r}_1 = \tilde{r}(\mathbf{c}_0, \mathbf{a}'_1 - \mathbf{c}_0)$. Here $\mathbf{a}'_1 = m_{M_{k_1}M_{k_2}\dots M_{k_K}}(\mathbf{a}_1)$. Let l be a line, such that $s_1 \subset l$. Let also $q = l \cap S_o(r, \mathbf{a})$. Obviously, q is an open set. But from Lemma 7, we have $q \subseteq s_1$. This implies that the mirror point \mathbf{a}_1 is contained in the open set q and simultaneously in straight segment s_1 . But, by definition of the straight segment (Def. 4.1.10), its open part, i.e. its interior, cannot contain mirror points which is a contradiction.

□

Corollary 5. *If for a given mirror system \mathcal{M} and a point $\mathbf{c}_0 \in \mathbb{R}^n$ with a positive distance to the mirror system $d = d(\mathcal{M}, \mathbf{c}_0) > 0$ (which implies $\mathbf{c}_0 \in (\mathbb{R}^n \setminus \mathcal{M}_f)$), the open ball $S_o(r, \mathbf{a})$ is visible in a virtual non-empty chamber C , then the distance $d(\mathbf{c}_0, m_C(\mathbf{a})) \geq d + r$.*

Visibility of an Open Ball from One Observation Point Implies Visibility of the Center of that Ball from Nearby Observation Points

The purpose of this subsection is to show that if an open ball of points is visible in a common chamber C from a common observation point \mathbf{c}_0 , then there are points around \mathbf{c}_0 that can see the center of that open ball in the same chamber.

For a given mirror system \mathcal{M} in \mathbb{R}^n , $\mathbf{c}_0 \in \mathbb{R}^n \mid d = d(\mathcal{M}, \mathbf{c}_0) > 0$. Suppose, the open ball $S_o(r, \mathbf{a})$, with the center at \mathbf{a} and radius $r > 0$ is visible from point \mathbf{c}_0 in the chamber $C = C_{M_{k_1}M_{k_2}\dots M_{k_K}}$, $K \geq 0$. Point \mathbf{a} transformed by the chamber C is given by $m_C(\mathbf{a}) = m_{M_{k_1}M_{k_2}\dots M_{k_K}}(\mathbf{a})$ while the open ball $S_o(r, \mathbf{a})$, transformed by the same chamber is the ball $S_o(r, m_C(\mathbf{a}))$.

Statement A: Since all the points inside the open ball $S_o(r, m_C(\mathbf{a}))$ belong to the same chamber C , there is a common sequence of virtual mirrors $Q = (m_{M_{k_1}}(M_{k_1}), m_{M_{k_1}M_{k_2}}(M_{k_2}), \dots, m_{M_{k_1}M_{k_2}\dots M_{k_K}}(M_{k_K}))$ that is intersected by any ray starting from the origin \mathbf{c}_0 and ending at some point inside the open ball $S_o(r, m_C(\mathbf{a}))$, see Fig. 4.10. For the special case $K = 0 \Rightarrow (\mathbf{c}_0 \in C_0, \mathbf{a} \in C_0)$ and there are no intersections with mirrors.

Moreover, the intersection of any virtual mirror from sequence Q with any of the rays is such that it is not along the hyperplane of the virtual mirror since the partitioning assigns a common non-empty chamber to all points of the open ball.

In the following, we argue in a two-dimensional setting that is obtained by embedding a ray through the observation point \mathbf{c}_0 and the center $m_C(\mathbf{a})$ of the open ball $S_o(r, m_C(\mathbf{a}))$ in an arbitrary “cut plane” P . The intersection of this plane P with the world in the full \mathbb{R}^n is used to construct our argument.

Let us refer to the ray $\tilde{r}(\mathbf{c}_0, \mathbf{m}_C(\mathbf{a}) - \mathbf{c}_0)$ as the ‘axial ray’. Consider a $2D$ plane P , i.e, a two-dimensional subspace of \mathbb{R}^n containing the axial ray. The intersection of the open ball $S_o(r, m_C(\mathbf{a}))$ and the plane P is the open disk D with center $m_C(\mathbf{a})$ and radius r , see Fig. 4.10.

The points $m_C(\mathbf{a}')$ and $m_C(\mathbf{a}'')$ are the points of intersection of the closure of the disk D with the tangent lines, passing through the common point \mathbf{c}_0 . The open disk D_0 with the center at \mathbf{c}_0 and with radius d is the intersection of the open ball $S_o(d, \mathbf{c}_0)$ with P .

According to Statement A, virtual mirrors from the original sequence Q intersect the plane P such, that the mirrors appear as open segments with their ends at the different sides of the angle $\alpha = \angle m_C(\mathbf{a}')\mathbf{c}_0 m_C(\mathbf{a}'')$. In the following, let us refer to this sequence of segments as Q , replacing the original meaning with our $2D$ interpretation in plane P . Each segment contained in the sequence is outside the two disks D and D_0 . In the illustration, Fig. 4.10, these open segments are: $(\mathbf{x}_1; \mathbf{y}_1), (\mathbf{x}_2; \mathbf{y}_2), \dots, (\mathbf{x}_K; \mathbf{y}_K)$.

More specifically, all the segments are arranged in such a way, that: a) for each such segment, if we divide the $2D$ plane P into two half-planes along the line containing the segment, one half-plane contains the open disk D_0 and all those segments for which the virtual mirrors appear before the current segment in sequence Q ; while b) the other half-plane contains all the remaining segments and the open disk D .

Lemma 9. *In the above construction, for $K > 0$, consider the ray \tilde{r}_1 from the observation point $\mathbf{c}_1 \in S_o(d, \mathbf{c}_0) \cap P$ and passing through the point $m_C(\mathbf{a})$. If the first mirror reflecting the ray \tilde{r}_1 is mirror M_{k_1} and the reflection point is $\mathbf{z}_1 \in (\mathbf{x}_1; \mathbf{y}_1)$, then the point \mathbf{a} is visible in chamber C as seen from \mathbf{c}_1 .*

Proof. We prove the lemma by mathematical induction. The base for induction is the Lemma’s condition that the ray \tilde{r}_1 is intersecting the mirror M_{k_1} first. Suppose now, the physical ray path for the ray \tilde{r}_1 has u first reflections, $1 \leq u \leq K$, taken in the sequence $M_{k_1}, M_{k_2}, \dots, M_{k_u}$. If $u = K = 1$, then the lemma is true.

Otherwise, let us prove that the physical ray path for the ray \tilde{r}_1 continues to $u + 1$ mirrors in the sequence $M_{k_1}, M_{k_2}, \dots, M_{k_u}, M_{k_{u+1}}$.

The ray \tilde{r}_1 intersects the mirror M_{k_1} , i.e. the mirror coincident with the virtual mirror $m_{M_{k_1}}(M_{k_1})$ at point \mathbf{z}_1 . Moreover, according to the conditions of the Lemma, this is the first intersection along the ray \tilde{r}_1 . Now, the point $\mathbf{z}_1 \in (\mathbf{x}_1; \mathbf{y}_1)$ is inside the angle $\alpha = \angle m_C(\mathbf{a}')\mathbf{c}_0 m_C(\mathbf{a}'')$. Since $m_C(\mathbf{a})$ is contained in the same angle and the angle is a convex set, all the points in the interval $[\mathbf{z}_1; m_C(\mathbf{a})]$ are also inside the angle α .

Moreover, by assumption, the physical path for the ray \tilde{r}_1 encounters the same first u mirrors in the sequence $M_{k_1}, M_{k_2}, \dots, M_{k_u}$ as any other ray connecting \mathbf{c}_0 and $S_o(r, \mathbf{a})$.

Then, according to Corollary 3 and 4, the ray \tilde{r}_1 has the same sequence of virtual mirrors $(m_{M_{k_1}}(M_{k_1}), m_{M_{k_1}M_{k_2}}(M_{k_2}), \dots, m_{M_{k_1}M_{k_2}\dots M_{k_u}}(M_{k_u}))$ and intersects them at the points $\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_u$ correspondingly, where $\mathbf{z}_i \in (\mathbf{x}_i; \mathbf{y}_i)$, $i = 1..u$.

Thus, after intersection with the virtual mirror $m_{M_{k_1}M_{k_2}\dots M_{k_u}}(M_{k_u})$, the ray partitioning $\Pi_{\mathcal{M}}(\tilde{r}_1)$ will have a following unfolded segment $[\mathbf{z}_u; \mathbf{z}^*)$ associated with the chamber $C_{M_{k_1}M_{k_2}\dots M_{k_u}}$, where $\mathbf{z}^* \in \tilde{r}_1$. Let us show, that $\mathbf{z}^* = \mathbf{z}_{u+1}$.

The point \mathbf{z}^* must correspond to a point on a mirror from mirror system \mathcal{M} in the space \mathbb{R}^n , transformed by the chamber transformation associated with chamber $C_{M_{k_1}M_{k_2}\dots M_{k_u}}$.

Then $\mathbf{z}^* \in (\mathbf{z}_u; \mathbf{z}_{u+1})$. Otherwise, the ray \tilde{r}_1 will intersect the point \mathbf{z}_{u+1} first, which is the point on the mirror $M_{k_{u+1}}$ and which is assigned to chamber $C_{M_{k_1}M_{k_2}\dots M_{k_{u+1}}}$ by the ray partitioning $\Pi_{\mathcal{M}}(\tilde{r}_1)$.

But if $\mathbf{z}^* \in (\mathbf{z}_u; \mathbf{z}_{u+1})$, then the ray $\tilde{r}(\mathbf{c}_0, \mathbf{z}^* - \mathbf{c}_0)$ in angle α is intersecting the point \mathbf{z}^* in the chamber $C_{M_{k_1}M_{k_2}\dots M_{k_u}}$ which is a contradiction because the next intersection point belongs to the next chamber $C_{M_{k_1}M_{k_2}\dots M_{k_{u+1}}}$.

Thus, $\mathbf{z}^* = \mathbf{z}_{u+1}$ and the physical ray path for the ray \tilde{r}_1 begins with $u + 1$ reflections in the sequence $M_{k_1}, M_{k_2}, \dots, M_{k_u}, M_{k_{u+1}}$. □


Corollary 6. *The statement of the Lemma remains true for the case $K = 0$, i.e. $C = C_0$, if we change the requirement for the ray \tilde{r}_1 to hit the mirror M_{k_1} first to the requirement that the ray \tilde{r}_1 can be divided into two parts by a point $\xi \in \mathbb{R}^n$ such that:*

1. ξ is inside the angle $\angle m_C(\mathbf{a}')\mathbf{c}_0m_C(\mathbf{a}'')$;
2. Points ξ and \mathbf{c}_0 are either on the same half-plane, relative to the mirror M_{k_1} or ξ is on the mirror M_{k_1} (in case of no mirrors we require the point ξ to be between \mathbf{c}_1 and $m_C(\mathbf{a})$);
3. $[\mathbf{c}_1; \xi) \cap \mathcal{M}_f = \emptyset$.

Mirror Systems are Observer-Continuous

On the basis of the previous Lemma and its Corollary, we now prove the following theorem:

Theorem 2. *In \mathbb{R}^n : $(\forall \mathcal{M}, \text{ Given any chamber } C \text{ and } \mathbf{c}_0 \in \mathbb{R}^n \mid d = d(\mathcal{M}, \mathbf{c}_0) > 0) \Rightarrow (\mathcal{M} \text{ is observation-continuous with respect to point } \mathbf{c}_0 \text{ and chamber } C)$.*

Figure 4.10: Mirror system \mathcal{M} , intersected by the 2D plane P .

Proof. Suppose, chamber $C = C_{M_{k_1} M_{k_2} \dots M_{k_K}}$ with $K \geq 0$. A point $\mathbf{a} \in \mathbb{R}^n$ transformed by chamber C is $m_C(\mathbf{a}) = m_{M_{k_1} M_{k_2} \dots M_{k_K}}(\mathbf{a})$. Suppose also that the point \mathbf{a} is visible in chamber C together with an open ball $S_o(r, \mathbf{a})$. The transformation of the open ball $S_o(r, \mathbf{a})$ by chamber C is given by the open ball $S_o(r, m_C(\mathbf{a}))$.

We define again an 'axial ray', i.e. the ray $\tilde{r}_a = \tilde{r}(\mathbf{c}_0, m_C(\mathbf{a}) - \mathbf{c}_0)$ and a segment connecting the observation point \mathbf{c}_0 and the center of the chamber-transformed ball $S_o(r, m_C(\mathbf{a}))$: $L = d(\mathbf{c}_0, m_C(\mathbf{a}))$, see Fig. 4.11.

To prove the theorem, it suffices to prove that there is a $d_1 > 0$, such that choosing any point $\mathbf{c}_1 \in S_o(d_1, \mathbf{c}_0)$ results in the point \mathbf{a} being visible in chamber C as seen from point \mathbf{c}_1 . An equivalent requirement is that the ray $\tilde{r}_{c_1} = \tilde{r}(\mathbf{c}_1, m_C(\mathbf{a}) - \mathbf{c}_1)$, up to the point $m_C(\mathbf{a})$ intersects the sequence of virtual mirrors $Q = (m_{M_{k_1}}(M_{k_1}), m_{M_{k_1} M_{k_2}}(M_{k_2}), \dots, m_{M_{k_1} M_{k_2} \dots M_{k_K}}(M_{k_K}))$.

Note that if $0 < d_1 < d \Rightarrow S_o(d_1, \mathbf{c}_0) \cap \mathcal{M} = \emptyset$. Thus, for any ray with the origin situated inside the open ball $S_o(d_1, \mathbf{c}_0)$, its partitioning associates all points of the ray inside the ball $S_o(d_1, \mathbf{c}_0)$ with the base chamber C_0 .

If the point $\mathbf{c}_1 \in S_o(d_1, \mathbf{c}_0)$ and \mathbf{c}_1 lies on the line containing the axial ray, the statement is obviously true.

Otherwise, let us consider the $2D$ plane P spanned by the axial ray and the point \mathbf{c}_1 . As in the previous Lemma, the open disk D with center $m_C(\mathbf{a})$ and radius r is the intersection of the open ball $S_o(r, m_C(\mathbf{a}))$ and the plane P . The points $m_C(\mathbf{a}')$ and $m_C(\mathbf{a}'')$ are the points of intersection of the closure of disk D with those tangent lines that are passing through the common point \mathbf{c}_0 . The open disk D_0 with the radius d is the intersection of the open ball $S_o(d, \mathbf{c}_0)$ and the plane P .




Figure 4.11: Sketch of projection of the unfolded mirror system onto the plane P spanned by points \mathbf{c}_0 , \mathbf{c}_1 and $m_C(\mathbf{a})$, when $L^2 > r^2 + d^2$.

There exist two cases: 1) $L^2 > r^2 + d^2$ (Fig. 4.11, the balls around \mathbf{c}_0 and $m_C(\mathbf{a})$ do not intersect), and 2) $L^2 \leq r^2 + d^2$ (Fig. 4.12, the balls intersect).




Figure 4.12: The case, when $[d(\mathbf{c}_0, m_C(\mathbf{a}))]^2 = L^2 \leq r^2 + d^2$.

1. The situation is similar to Fig. 4.10. Consider the two points \mathbf{h}_1 and \mathbf{h}_2 , the intersections of the border of disk D_0 with the segments $[\mathbf{c}_0; m_C(\mathbf{a}')]$ and $[\mathbf{c}_0; m_C(\mathbf{a}'')]$, respectively. Consider now the angle $\angle \mathbf{h}_1 m_C(\mathbf{a}) \mathbf{h}_2$. It can be shown (see Appendix A) that the distance from point \mathbf{c}_0 to the sides of the angle is given by:

$$d(\mathbf{c}_0, \mathbf{b}_1) = d(\mathbf{c}_0, \mathbf{b}_2) = \frac{r \cdot d}{\sqrt{(\sqrt{L^2 - r^2} - d)^2 + r^2}}, \quad (4.2)$$

Let $d_1 = d(\mathbf{c}_0, \mathbf{b}_1) = d(\mathbf{c}_0, \mathbf{b}_2)$. From Eq. 4.2 we see that $d_1 = \frac{r \cdot d}{\sqrt{(\sqrt{L^2 - r^2} - d)^2 + r^2}} > \frac{r \cdot d}{\sqrt{(\sqrt{L^2 - r^2})^2 + r^2}} = \frac{r \cdot d}{L}$. On the other hand, since case 1) covers $L^2 > r^2 + d^2 \Rightarrow d_1 = \frac{r \cdot d}{\sqrt{(\sqrt{L^2 - r^2} - d)^2 + r^2}} < \frac{r \cdot d}{r} = d$.

Consequently:

$$\frac{r \cdot d}{L} < d_1 < d. \quad (4.3)$$

If the point \mathbf{c}_1 is situated inside the open disk with radius d_1 and center \mathbf{c}_0 , then, obviously, the segment $[\mathbf{c}_1; m_C(\mathbf{a})]$ is inside the angle $\angle \mathbf{h}_1 m_C(\mathbf{a}) \mathbf{h}_2$. Simultaneously, point \mathbf{c}_1 is automatically part of the disk D_0 since $d_1 < d$.

Moreover, the point ξ , which is the intersection of the segment $[\mathbf{c}_1; m_C(\mathbf{a})]$ with the border of the disk D_0 (which is a circle) satisfies the conditions of Corollary 6. Thus, the theorem is proven for the case 1).

2. The second case, $L^2 \leq r^2 + d^2$, is sketched in Fig. 4.12. Consider any point $\mathbf{z} \in D_0$. The segment $[\mathbf{z}; m_C(\mathbf{a})]$ is contained completely inside the union of the two open disks D and D_0 . To see this, consider any point on the segment $[\mathbf{z}; m_C(\mathbf{a})]$, for example point \mathbf{f} (see Fig. 4.12), to be located outside of these disks. Then, $d(\mathbf{f}, \mathbf{c}_0) \geq d$ and $d(\mathbf{f}, m_C(\mathbf{a})) \geq r$.

Adding the condition $L^2 \leq r^2 + d^2$ of case 2), we conclude that the angle $\angle \mathbf{c}_0 \mathbf{f} m_C(\mathbf{a}) \leq \frac{\pi}{2}$ and that therefore the angle $\angle \mathbf{z} \mathbf{f} \mathbf{c}_0 = \pi - \angle \mathbf{c}_0 \mathbf{f} m_C(\mathbf{a}) \geq \frac{\pi}{2}$. Thus, $d(\mathbf{c}_0, \mathbf{z}) > d(\mathbf{c}_0, \mathbf{f}) > d$ which is contradicting the assumption that $\mathbf{z} \in D_0$.

On the other hand, by using Lemma 8 we see that there are no real and no virtual mirrors intersecting the disk D for any point $\mathbf{z} \in D_0$. Thus, the union of the open disks $D \cup D_0$ does not contain either real or virtual mirrors.

Thus, the point \mathbf{a} is visible from any point $\mathbf{z} \in D_0$ and, moreover, the chamber $C = C_0$.

□

In plain words, Theorem 2 states that if in a mirror system, the observation point \mathbf{c}_0 is remote from all mirrors of the system, and if, from \mathbf{c}_0 , it is possible to observe some ball with center \mathbf{a} through a common reflection sequence (chamber C), then there is a ball around the observation point \mathbf{c}_0 such that point \mathbf{a} is visible in chamber C from any point within this ball. Put differently, if the point of observation is smoothly changing, then the set of observable points is also smoothly changing.

The next Corollary provides a bound for the radius of the ball around the observation point \mathbf{c}_0 :

Corollary 7. *In \mathbb{R}^n : ($\forall \mathcal{M}$, chamber C and $\mathbf{c}_0 \in \mathbb{R}^n \mid d = d(\mathcal{M}, \mathbf{c}_0) > 0$) \Rightarrow if $S_o(r, \mathbf{a})$ is visible in chamber C , then the point \mathbf{a} is visible from any point of the ball $S_o(d^*, \mathbf{c}_0)$, where*

$$d^* = \min\left(\frac{r \cdot d}{d(\mathbf{c}_0, m_C(\mathbf{a}))}, d\right) > 0. \quad (4.4)$$

4.2.2 Observer-Continuity For Point Sets

Definition 4.2.5. A mirror system \mathcal{M} is called L-uniformly continuous with respect to a given set of points of observation $S \subset \mathbb{R}^n$ iff $\forall r > 0, \exists r_1(r)$ such that $(\forall C \in \overline{C}_{\mathcal{M}} \setminus C_{\emptyset}, \forall \mathbf{a} \in \mathbb{R}^n, \forall \mathbf{c}_0 \in S \mid$

- a) $d(\mathbf{c}_0, m_C(\mathbf{a})) \leq L$;
- b) $\forall \mathbf{a}' \in S_o(r, \mathbf{a}) \Rightarrow \mathbf{a}'$ is visible from \mathbf{c}_0 in chamber C
 $\Rightarrow \forall \mathbf{c}' \in S_o(r_1, \mathbf{c}_0) \Rightarrow \mathbf{a}$ is visible from \mathbf{c}' in chamber C .

The definition states that visibility is uniformly continuous within an L -ball around an arbitrary $\mathbf{c}_0 \in S$, i.e. the radius r_1 of the open ball around \mathbf{c}_0 depends only on the globally chosen radius r of the open ball around $m_C(\mathbf{a})$.

Theorem 3. *If for a given mirror system \mathcal{M} and a set of points $S \in \mathbb{R}^n$, the distance $d(\mathcal{M}, S) = d > 0$, then for any fixed $L > 0$ the mirror system \mathcal{M} is L -uniformly continuous with respect to S .*

Proof. Since $d = d(\mathcal{M}, S) > 0 \Rightarrow (\forall \mathbf{c}_0 \in S \Rightarrow d(\mathcal{M}, \mathbf{c}_0) \geq d)$. Then by Theorem 2 $\forall \mathbf{c}_0 \in S$, chamber $C \in \overline{C}_{\mathcal{M}} \setminus C_{\emptyset} \Rightarrow$ the mirror system \mathcal{M} is observer-continuous with respect to point \mathbf{c}_0 and chamber C . Then $(\forall \mathbf{a} \in \mathbb{R}^n \mid (\forall \mathbf{a}' \in S_o(r, \mathbf{a}) \Rightarrow (\mathbf{a}'$ is visible in chamber C from point $\mathbf{c}_0))) \Rightarrow \mathbf{a}$ is visible from any $\mathbf{c}' \in S_o(r_1, \mathbf{c}_0)$, where $r_1 = \min(\frac{r \cdot d}{d(\mathbf{c}_0, m_C(\mathbf{a}))}, d)$ in accordance with Corollary 7.

Selecting only such pairs \mathbf{a} and C for which $d(\mathbf{c}, m_C(\mathbf{a})) \leq L$ results in $r_1 \geq \min(\frac{r \cdot d}{L}, d)$, fulfilling the theorem's conditions. \square

If we consider a mirror system \mathcal{M} , some finite n -dimensional area of interest $A \in \mathbb{R}^n$, some finite n -dimensional area of observation $S \in \mathbb{R}^n$, $d = d(\mathcal{M}, S) > 0$ and a finite number of reflections N , then the area A in the mirror system \mathcal{M} under the chamber transformation m_C of a level not exceeding N will be at a finite distance to any of the points in S , this distance not exceeding some L .

Theorem 3 implies, that in such case, we may sample all the possible observations of up to N reflections of the set A as observed in the mirror system \mathcal{M} from the set of points S . The sampling can be made accurate up to an arbitrary precision that is given in advance. This is achieved by sampling S with a uniform grid, where the spacing of the grid is determined by Corollary 7.

This situation is of practical interest since in a real scenario, the number of reflections and the areas of interest and observations are limited.

Conclusions

We have presented a theoretical toolbox for interpreting and analysing planar mirror systems. The key concept of this toolbox is the space partitioning based on the well-known ray straightening procedure. With support of the partitioning concept we have shown the continuity of planar mirror systems with respect to observation point change. This result may lead to global optimization schemes for determining a best observer position in a given mirror system. Additionally, the result enables direct triangulation within unfolded space, a capability that may lead to new reconstruction schemes.

In the following part of the thesis we start using our theoretical basics for approaching practical problems involving planar mirror systems in computer vision.

PART III

Three-Dimensional Kaleidoscopic
Imaging

Introduction

In this part of the thesis, we start using the theory, developed in the previous part, for practical applications. The main application class considered here is multi-view imaging and multi-directional illumination.

The classical way to acquire multi-view imagery is the use of multi-camera systems [Wilburn05] or light field camera designs [Levoy96, Gortler96, Ng05]. These options are typically expensive if many cameras are used, or involve time-sequential capture. In addition, available sensor resolution is used inefficiently, mostly recording background information. For in-camera light field imaging, e.g. [Ng05], there is the limited aperture of the main lens and thus the achievable coverage of view-points is limited. Time sequential capture can also be performed with a moving planar mirror [Murray95, Ihrke08, Hu09]. Mirrors have also been used to capture light field information [Levoy04, Sen05] but without considering inter-reflections. Light field imaging has also been performed with a conical mirror [Taguchi10a]. More typically, catadioptric systems are used to achieve a wide field of view as in panoramic imaging, e.g. [Lin06]. 3D reconstruction using a conical mirror device by taking multiple images has also been demonstrated [Kuthirummal06].

Multi-directional illumination is often applied for reflectance field acquisition. The classical ways here are similar to the camera case. They can be classified into those, which capture the scene time sequentially while moving the light source(s) and/or the object [Debevec00, Schwartz13, Holroyd10, Masselus03, Lensch03] and into those, utilizing a large amount of fixed light sources [Schwartz11, Ben-Ezra08].

In contrast to all these techniques, kaleidoscopic systems enable to perform massively multi-view imaging with a single high-resolution camera using only a single image as well as performing multi-directional illumination with a single projector.

The Labeling Problem and its Solution

5.1 Planar Mirror System Containing Objects

In the theoretical Part II of this thesis we introduced unfolding for an arbitrary planar mirror system. However, we assumed an absence of obstacles along the path of rays. When objects are introduced into a mirror system, some rays may be blocked by these objects after a certain number of reflections.

Let \mathcal{M} be a planar mirror system in \mathbb{R}^n , H a set of rays in \mathbb{R}^n ($\tilde{r}(\mathbf{c}, \mathbf{d}) \in H \Rightarrow \mathbf{c} \in (\mathbb{R}^n \setminus \mathcal{M}_f)$), and $O_1 \sqcup O_2 \subset (\mathbb{R}^n \setminus \mathcal{M}_f)$ two distinct sets of points (two objects: foreground O_1 and background O_2) in \mathbb{R}^n , see Fig. 5.1.

Definition 5.1.1. Observable object (in a chamber). In a mirror system \mathcal{M} an object O_1 is observable by the ray $\tilde{r} = \tilde{r}(\mathbf{c}, \mathbf{d}) \in H$ in a non-empty chamber $C_{M_{k_1} M_{k_2} \dots M_{k_L}}$ iff

- $\exists \mathbf{o}_1 \in O_1 \mid$ (point \mathbf{o}_1 is visible, Def. 4.1.17, from \mathbf{c} in direction of \mathbf{d} in a chamber $C_{M_{k_1} M_{k_2} \dots M_{k_L}}$, at a distance $\lambda_1 \geq 0$).
- $\nexists \mathbf{o}_2 \in O_2 \mid$ (point \mathbf{o}_2 is visible from \mathbf{c} in direction of \mathbf{d} , at distance $\lambda_2 < \lambda_1$, $\lambda_2 \geq 0$), i.e. there is no earlier ray intersection with the background object.
- $(\forall \mathbf{o}' \in O_1, C_x \in \{C_0, C_{M_{k_1}}, C_{M_{k_1} M_{k_2}}, \dots, C_{M_{k_1} M_{k_2} \dots M_{k_{L-1}}}\}) \Rightarrow$ (the point \mathbf{o}' is not visible in chamber C_x from the point \mathbf{c} in direction \mathbf{d}), i.e. the object does not block the ray itself at an earlier reflection.




Figure 5.1: Illustrating an object O_1 , its convex bounding set $\overline{O_1}$ and a background O_2 in a mirror system \mathcal{M} containing three mirrors M_1 , M_2 , M_3 .

In our practical experiments we will work in \mathbb{R}^3 . The set H models a camera/projector, where every pixel is represented by one ray, O_1 is an object of interest, while all other physical objects, excluding the mirrors forming the mirror system \mathcal{M} (e.g. background area, camera and/or projector body) will be represented by O_2 .

We also need to introduce $\overline{O_1}$ which is a convex bounding set for the object O_1 , see Fig. 5.1. More precisely, we require $\overline{O_1} \subset (\mathbb{R}^n \setminus \mathcal{M}_f) \mid (\overline{O_1} \cap O_2 = \emptyset, \overline{O_1}$ is a convex set, and $O_1 \subset \overline{O_1}$). The convex bounding set will be necessary for our theoretical argumentation and later, in practical applications, we will use it as the object's bounding box.

Let an object O_1 be observable by ray $\tilde{r} = \tilde{r}(\mathbf{c}, \mathbf{d}) \in H$ in a non-empty chamber $C_{M_{k_1} M_{k_2} \dots M_{k_L}}$ of a mirror system \mathcal{M} . It follows that the virtual¹ object $m_{M_{k_1} M_{k_2} \dots M_{k_L}}(O_1)$ is the first obstacle on the way of the straightened ray \tilde{r} . For our applications it is more convenient to consider different virtual views rather than different virtual objects, Section 3.1. This can be achieved by applying the inverse chamber transformations $m_{M_{k_1} M_{k_2} \dots M_{k_L}}^{-1}$ to those chambers, where the object O_1 is observed. The result is the virtual ray $m_{M_{k_1} M_{k_2} \dots M_{k_L}}^{-1}(\tilde{r})$ which first hits the object O_1 (this is because $m_{M_{k_L} M_{k_{L-1}} \dots M_{k_1} M_{k_1} \dots M_{k_{L-1}} M_{k_L}}(O_1) = O_1$).

Because rays are propagating independently from each other, for a given set of rays H , an object O_1 may be observable in different chambers for

¹From now on, to avoid repetitive statements, the word 'virtual' includes also the real case.

different rays, causing different virtual view points of the object. For example, if all the rays in H have a common origin \mathbf{c}_0 (a projection center), then the object O_1 may be observable from different virtual points of view (according to those chambers where the object O_1 is observable). This is how an object inside a kaleidoscope, observed by a single camera, can be seen from different virtual points of view.

In practice, the overall number of rays is usually much larger than the number of considered chambers. Therefore, it is convenient to perform a ray classification, joining into the same class those and only those rays which observe the object O_1 in a common chamber and those that do not see the object O_1 . More precisely this classification can be summarized in the following

Definition 5.1.2. Labeling function.

Consider a mirror system \mathcal{M} with objects O_1 and O_2 , and a ray $\tilde{r} = \tilde{r}(\mathbf{c}, \mathbf{d})$, $\mathbf{c} \in (\mathbb{R}^n \setminus \mathcal{M}_f)$. The labeling function $Label_{\mathcal{M}, O_1, O_2}(\tilde{r}) =$

$$\begin{cases} C_{M_{k_1} M_{k_2} \dots M_{k_L}}^{-1}, & \text{if the object } O_1 \text{ is observable by } \tilde{r} \text{ in } C_{M_{k_1} M_{k_2} \dots M_{k_L}} \\ \emptyset, & \text{otherwise.} \end{cases}$$

Note that the labeling function is a function as for each ray there is no more than one chamber, where the object is observable.

If H is a set of rays in mirror system \mathcal{M} with objects O_1 and O_2 , then H can obviously be decomposed into $H = H_\emptyset \sqcup H_1 \sqcup H_2 \dots \sqcup H_N \dots$, where $H_\emptyset = Label_{\mathcal{M}, O_1, O_2}^{-1}(\emptyset) \cap H$, $H_1 = Label_{\mathcal{M}, O_1, O_2}^{-1}(C_1^{-1}) \cap H$, $H_2 = Label_{\mathcal{M}, O_1, O_2}^{-1}(C_2^{-1}) \cap H, \dots$, $H_N = Label_{\mathcal{M}, O_1, O_2}^{-1}(C_N^{-1}) \cap H, \dots$

Here $Label_{\mathcal{M}, O_1, O_2}^{-1}(x)$ means the set of all the rays from H for which the labeling function is equal to x . $C_1^{-1}, C_2^{-1}, \dots, C_N^{-1}, \dots$ are the inverted enumerated chambers (chambers are countable, see Part II) of the mirror system \mathcal{M} .

Then the object O_1 is observed from the following sets of virtual rays:

$$V_1 = m_1(H_1), V_2 = m_2(H_2), \dots, m_N(H_N), \dots \quad (5.1)$$

Here m_i , $i \geq 1$, are the chamber transformations corresponding to the chambers C_i . On the other hand, the object O_1 is not observed from the set of rays H_\emptyset .

In real applications, the set of rays H is finite. In this case, no matter if object O_1 is observable in a finite amount of chambers or not, the set in Equation 5.1 contain only a finite amount of non-empty elements.

Definition 5.1.3. Virtual viewset. We refer to an element $m_i(H_i)$ from the set in Equation 5.1 as a virtual viewset V_i .

An obvious property of virtual viewsets is that if all the rays in H have a common origin, any non-empty virtual viewset has a common virtual origin, i.e. a virtual viewpoint.




Figure 5.2: Labeled pixels providing virtual viewpoints. Virtual objects 1 – 4 are fully observable, 5 and 6 are partially observable while 7 is fully occluded. Even if 7 was observable only part of it could be observed due to the space partitioning introduced by the mirrors.

Definition 5.1.4. Virtual viewsets observing the object. Let the set of rays H be finite. Then we define the set of virtual viewsets where the object O_1 is observable as $V(\mathcal{M}, O_1, O_2, H)$, i.e. $V(\mathcal{M}, O_1, O_2, H) := \{V_{i_1}, V_{i_2}, \dots, V_{i_L}\}$, where $V_{i_j} (j = 1..L)$ are those and only those virtual viewsets that are non-empty.

If the labeling function for the whole set of rays H is known, the corresponding set V of virtual viewsets observing the object is well-defined and can be directly used in multi-view applications. Practically, when H is finite, the set of virtual viewsets observing an object may be constructed by classifying each ray in H separately into viewsets. The labeling function $Label_{\mathcal{M}, O_1, O_2}(H)$ (and the set $V(\mathcal{M}, O_1, O_2, H)$) depend on \mathcal{M} , O_1 , O_2 , and H . When all these parameters are known, the labeling (and thus the virtual viewsets) are trivial to obtain. An illustration of this in case of a projective observer C in $2D$ is shown in Fig. 5.2. Differently colored cones encode different viewpoints of the object. Virtual objects 1 – 4 are fully observable, 5 and 6 are partially observable while 7 is fully occluded. Even if 7 was observable, only part of it could be observed due to the space partitioning introduced by the mirrors. As can be seen from the figure, to estimate the labeling function, the object only has to be intersected in all its observable mirrored positions, as determined by the space-partitioning of the system, with occlusion taken into account (see Definition 5.1.2).

In practice, the geometric calibration steps described in Section 3.3.1 provide us with \mathcal{M} and H . We organize the background O_2 in such a way that it is not influencing the observability of the object O_1 (this is possible if, for example, O_2 is outside of the convex hull of the mirror system). However, the geometry and exact pose of O_1 is, usually, unknown. Thus, to recover

the labeling function, it is necessary to infer some geometric structure of the object under consideration.

5.2 Visual Hull as Sufficient Approximation of the Object for the Determination of the Labeling Function

In the previous section we showed, that knowledge of O_1 is required to evaluate the labeling function (and the set of virtual viewsets). In this section we show, that, under special conditions, the labeling function can be recovered if only a certain visual hull (for the definition of a visual hull see Section 3.2.1), of the object O_1 is known. For simplicity, we consider only the case where the observer uses a central projection model (all rays in H have a common origin \mathbf{c}_0), but it is relatively straightforward to generalize the argument to arbitrary sets of rays H (in this case multiple unfoldings with different observation points should be considered).

Theorem 4. *Let a mirror system \mathcal{M} , with foreground and background objects O_1 and O_2 be observed by a bundle of rays H with a common origin \mathbf{c}_0 .*

Let $S = \{C_1, C_2, \dots, C_N, \dots\}$ be the set of chambers observable from point \mathbf{c}_0 , i.e. those and only those chambers, where at least one point in $(\mathbb{R}^n \setminus \mathcal{M}_f)$ is visible from \mathbf{c}_0 .

Let $S^ = \{C_1^*, C_2^*, \dots, C_N^*, \dots\}$ be the set of inverted chambers obtained from S , i.e. $C_i^* = C_i^{-1}$.*

Let $V^ = \{m_1^*(\mathbf{c}_0), m_2^*(\mathbf{c}_0), \dots, m_N^*(\mathbf{c}_0), \dots\}$ be the set of virtual view-points, corresponding to the chambers from S^* , i.e. $m_i^*(\mathbf{c}_0) = m_{C_i^{-1}}(\mathbf{c}_0) = m_{C_i}^{-1}(\mathbf{c}_0)$.*

Suppose $\exists \overline{O_1}$ (a convex bounding set for the object O_1 , $\overline{O_1} \cap \mathcal{M}_f = \emptyset$). Then $\forall \tilde{r} \in H \Rightarrow \text{Label}_{\mathcal{M}, O_1, O_2}(\tilde{r}) = \text{Label}_{\mathcal{M}, VH(O_1, V^) \cap \overline{O_1}, O_2}(\tilde{r})$.*

Proof. We proof the theorem in two steps.

1) If the object O_1 is observable by the ray in the chamber C_i (chamber, inverted to the chamber C_i^*), then the object $VH(O_1, V^*) \cap \overline{O_1}$, containing O_1 , is also observable (in chamber C_i or in one of the preceding chambers).

Let for some ray $\tilde{r} \in H \Rightarrow \text{Label}_{\mathcal{M}, O_1, O_2}(\tilde{r}) \neq \emptyset$. Then $\exists! C_i^* = C_{M_{k_L} M_{k_{L-1}} \dots M_{k_1}} \in S^* \mid \text{Label}_{\mathcal{M}, O_1, O_2}(\tilde{r}) = C_i^*$. But $O_1 \subseteq (VH(O_1, V^*) \cap \overline{O_1}) \Rightarrow \text{Label}_{\mathcal{M}, VH(O_1, V^*) \cap \overline{O_1}, O_2}(\tilde{r}) = C_j^*$, where $C_j^* \in \{C_0, C_{M_{k_1}}, C_{M_{k_2} M_{k_1}}, \dots, C_{M_{k_L} M_{k_{L-1}} \dots M_{k_1}}\}$.

2) A ray from H that has a label under $VH(O_1, V^*) \cap \overline{O_1}$ has the same label under O_1 .

For some ray \tilde{r} , let $\tilde{r} = \tilde{r}(\mathbf{c}_0, \mathbf{d}) \in H \Rightarrow Label_{\mathcal{M}, VH(O_1, V^*) \cap \overline{O_1}, O_2}(\tilde{r}) \neq \emptyset$. Then $\exists! C_i^* = C_{M_{k_L} M_{k_{L-1}} \dots M_{k_1}} \in S^* \mid Label_{\mathcal{M}, VH(O_1, V^*) \cap \overline{O_1}, O_2}(\tilde{r}) = C_i^*$. Because $\overline{O_1}$ is convex and does not contain points from \mathcal{M}_f , the set $\overline{O_1} \cap m_i^*(\tilde{r})$ is connected and does not contain points from \mathcal{M}_f . Thus, the points $m_i(\overline{O_1}) \cap \tilde{r}$ [and, consequently, the points $m_i(VH(O_1, V^*) \cap \overline{O_1}) \cap \tilde{r}$] are assigned to the common chamber C_i (according to the value of the labeling function) in the space partitioning $\Pi_{\mathcal{M}}^*(\mathbf{c}_0)$ and none of these points belongs to $m_i(O_2)$ ($\overline{O_1} \cap O_2 = \emptyset$). On the other hand, $m_i^*(\mathbf{c}_0) \in V^*$ and thus, from the visual hull definition, $m_i^*(\tilde{r}) \cap O_1 \neq \emptyset$ while $m_i^*(\tilde{r}) \cap O_1 \subset m_i^*(\tilde{r}) \cap VH(O_1, V^*)$. Consequently, $\emptyset \neq m_i^*(\tilde{r}) \cap O_1 = m_i^*(\tilde{r}) \cap O_1 \cap \overline{O_1} \subset m_i^*(\tilde{r}) \cap VH(O_1, V^*) \cap \overline{O_1} \Rightarrow$ [if we apply mirror transformation m_i^*] $\emptyset \neq \tilde{r} \cap m_i(O_1) \subset \tilde{r} \cap m_i(VH(O_1, V^*) \cap \overline{O_1})$ which is entirely assigned to the chamber C_i and contains no points from $m_i(O_2)$. Then the ray \tilde{r} is intersecting the virtual object $m_i(O_1)$ in chamber C_i and does not intersect any virtual copies of O_2 or O_1 before this event. Thus, $Label_{\mathcal{M}, VH(O_1, V^*) \cap \overline{O_1}, O_2}(\tilde{r}) = Label_{\mathcal{M}, O_1, O_2}(\tilde{r})$.

Arguments 1) and 2) prove the theorem. \square

Theorem 4 allows us to estimate the labeling function using only the visual hull of an unknown object O_1 , i.e to establish a proper labeling we do not actually need to know the true object geometry, but an object visual hull $VH(O_1, V^*)$ suffices. However, if the number of reflections is unlimited, the set of virtual viewpoints V^* may contain an infinite amount of elements, even if the ray set H is finite.

5.3 Visual Hull Estimation

While the knowledge of the visual hull $VH(O_1, V^*)$ of the foreground object O_1 is sufficient to exactly recover the labeling function (and the set of virtual viewsets), it is more practical to be able to produce a labeling function for arbitrary objects directly from the kaleidoscope images. In this case the visual hull computation is restricted by self-occlusions and partial visibility in different chambers. Our method is to find an approximation of the labeling function using an approximation to the visual hull $VH(O_1, V^*)$. The approximation of the object's visual hull is based on considering rays that do *not intersect* any real or mirrored version of the object.

If, in the kaleidoscopic setup, gaps between an object and its mirror images are observable, these gaps provide the means to perform an approximation of the visual hull $VH(O_1, V^*)$ of the object.

Consider a ray that does not intersect the object nor any of its virtual counterparts. Folding back this ray into the base chamber, we obtain a reflected light path that is guaranteed to be free of intersections. Performing this operation on the set of all rays that do not intersect the object anywhere

in mirror space, we obtain a space carving scheme to determine an approximation of the visual hull $VH(O_1, V^*)$ of the object O_1 : It is computed by successively removing free space from an initial volume (convex bounding set $\overline{O_1}$) that is marked as containing the object O_1 .

Definition 5.3.1. The approximation of the intersection of a visual hull $VH(O_1, V^*)$ with a bounding volume $\overline{O_1}$, obtained by the algorithm above is called **kaleidoscopic visual hull**. We denote it as $KVH := KVH(O_1, \overline{O_1}, V_0^*)$, where V_0^* is the set of virtual viewpoints that is generated by the virtual chambers that are intersected by rays from H .

Using $KVH(O_1, \overline{O_1}, V_0^*) = KVH(O_1, \overline{O_1}, V_0^*) \cap \overline{O_1}$, our labeling function approximation method can be explained by the following three steps:

1. We compute $KVH(O_1, \overline{O_1}, V_0^*)$, where $KVH(O_1, \overline{O_1}, V_0^*) = KVH(O_1, \overline{O_1}, V_0^*) \cap \overline{O_1} \approx VH(O_1, V^*) \cap \overline{O_1}$;
2. Then $Label_{\mathcal{M}, KVH(O_1, \overline{O_1}, V_0^*), O_2}(H) = Label_{\mathcal{M}, KVH(O_1, \overline{O_1}, V_0^*) \cap \overline{O_1}, O_2}(H) \approx Label_{\mathcal{M}, VH(O_1, V^*) \cap \overline{O_1}, O_2}(H)$;
3. Using Theorem 4, $Label_{\mathcal{M}, VH(O_1, V^*) \cap \overline{O_1}, O_2}(H) = Label_{\mathcal{M}, O_1, O_2}(H)$. Then $Label_{\mathcal{M}, KVH(O_1, \overline{O_1}, V_0^*), O_2}(H) \approx Label_{\mathcal{M}, O_1, O_2}(H)$.

The space carving nature of the kaleidoscopic visual hull computation approach provides us with some additional properties:

1. If the calibration of the system is perfect in the sense that all rays from H that do not intersect the object in the mirror space are precise, the kaleidoscopic visual hull $KVH(O_1, \overline{O_1}, V_0^*)$ is always larger or equal to the intersection of the object's visual hull with a bounding set: $VH(O_1, V^*) \cap \overline{O_1}$.
2. The first property remains true if instead of considering complete ray paths, we use only parts of them. For example, we can select only ray paths up to a limited number of reflections, up to a given distance, or even restrict the computation to some set of intervals or points on the ray path. Finally, it is possible to ignore any subset of rays in H (this corresponds to parts of the camera/projector image).
3. If a background ray is selected incorrectly (i.e., it intersects the object in one of the virtual chambers) or if the background ray is badly calibrated (it propagates too far from the estimated trajectory) it may "damage" the visual hull approximation by cutting extra volume, that belongs to the visual hull $VH(O_1, V^*)$, see Fig. 5.4 (e)-(f)-(g).

The properties above are useful in practical situations. They allow us to not propagate rays too far. Typically, no more than 7-9 reflection levels were considered in our experiments. The main reason is the relatively low calibration precision at high levels of reflections. Another reason is to limit the computational costs.




Figure 5.3: Unreliable pixel labeling can occur by a ray intersecting the kaleidoscopic visual hull while the real ray passes the object. Two cases are important: (a) the ray passes at an object boundary and (b) the object has a hole that cannot be recovered from kaleidoscopic silhouette information. (c) The pixels can be grouped into reliable and unreliable pixels according to the number of kaleidoscopic visual hulls they intersect (one and more than one respectively).

5.4 Applicability of the Kaleidoscopic Visual Hull for Image Labeling

The kaleidoscopic visual hull computed as described above can effectively be used as a geometric proxy for the object. By transforming the kaleidoscopic visual hull into the mirror chambers and intersecting the straightened camera rays in unfolded mirror space with the set of mirrored kaleidoscopic visual hulls, we can label the rays with respect to the mirror chamber where the ray first intersects the kaleidoscopic visual hull. This way a virtual view of the object is determined for each pixel. In the following, we discuss the limits of applicability of this scheme.

If we aim to properly label a set of rays H (e.g. an image of a camera), we, ideally, should know the visual hull $VH(O_1, V^*)$, where V^* is defined as above. But the proposed labeling scheme has one important drawback: the kaleidoscopic visual hull is not the object's visual hull $VH(O_1, V^*)$. Especially in cases of overlapping silhouettes it is possible that rays that intersect the kaleidoscopic visual hull do not actually intersect the object's visual hull $VH(O_1, V^*)$, see Fig. 5.3.

In practical situations, the case of Fig. 5.3 (a) is the most important. The kaleidoscopic visual hull as obtained from the mirror system is slightly too large and the ray \tilde{r} is not intersecting object O_1 (and its visual hull $VH(O_1, V^*)$) while intersecting its kaleidoscopic visual hull. The pixel is classified as belonging to base chamber C_0 . In reality, however, the ray intersects the virtual object O'_1 obtained by reflection through mirror 1 and thus should be labeled as belonging to chamber C_1 .

We have been unable to provide a narrow classification rule for these

types of pixels. We can however show that the problem is unsolvable for general objects and that there is thus no simple solution, i.e. one without heuristics. Consider the case of Fig. 5.3 (b). The real object O_1 , which has the same kaleidoscopic visual hull as O_1 of Fig. 5.3 (a), has a very narrow hole (in the limit a Dirac-like opening only permitting a single ray), only visible from \mathbf{c} . The ray passes through the hole but is blocked by virtual object O'_1 . Since the hole direction only permits a single ray, the hole is not visible from any other direction. It is thus impossible to update the kaleidoscopic visual hull to include the hole (in contrast, the visual hull $VH(O_1, V^*)$ will have this hole) and the labeling fails. The proper course of action is to classify all pixels corresponding to rays intersecting more than one real or virtual kaleidoscopic visual hull of the object as *unreliable* which means that they cannot be properly assigned to any particular view, see Fig. 5.3 (c). Note however that there is only a certain number of candidate views, namely the ones corresponding to the intersected virtual kaleidoscopic visual hulls. Future multi-view algorithms could exploit this information.

In practice the hole problem, Fig. 5.3 (b) is insignificant in all cases that we investigated. The case of Fig. 5.3 (a) however, where different virtual views meet side by side, is unavoidable except for simple systems where there is no overlap between mirror images. The better the kaleidoscopic visual hull approximates the visual hull $VH(O_1, V^*)$, the smaller is this error. Since we have a large number of views (typically around 200 usable ones) this error is small and can be mitigated by enforcing a safety region around occlusion boundaries in image space.

5.5 Projector Image Labeling

When a projector is introduced into a planar kaleidoscopic system it is possible to highlight the object from different virtual directions. These can be determined by the same unfolding procedure as in the camera case. This property is especially useful for reflectance measurements, but superposition of light due to possible illumination of the same surface areas from different directions is highly undesirable in such applications.

The solution to this problem is to illuminate the object from a single virtual projector view at a time. This can be easily done if the projector's image is properly labeled.

Unfortunately, in contrast to the camera, it is not obvious how to classify projector rays into those which propagate through the kaleidoscope without intersection with the object and those that highlight the object (we also need to know from which virtual view the object was illuminated). Thus, the visual hull estimation algorithm (see Section 5.3) is not directly applicable in case of a projector.

However, the labeling procedure is rather independent from the way

we obtained the object's visual hull approximation. Thus, to produce an approximate labeling it is possible to use the approximation to the visual hull $VH(O_1, V^*)$ of the object O_1 . In particular, such visual hull approximation may be the kaleidoscopic visual hull $KVH(O_1, \overline{O_1}, V_1^*)$, produced from a different set of rays H_1 (for example from a camera). Then, if the projector is calibrated, the labeling is straightforward.

In Chapter 6 we develop a joint reflectance and geometry scanner, exploiting the possibility of labeling the projector's image by using the kaleidoscopic visual hull obtained from a camera, observing the same scene, but from a different viewpoint.

5.6 Non-central Projection Camera and Projector Image Labeling

In the previous sections we discussed the labeling for central projection cameras and projectors. In such cameras and projectors all the camera/projector rays are originating from a single point, i.e. from the projection center. The projection center is the point of observation in the mirror system and the determining parameter for the space partitioning (or unfolding) into chambers. The space partitioning for the central projection camera or projector is therefore single and common for all image rays.

When the camera or the projector are non-central projection devices, each image ray is originating from its own $3D$ point, i.e. in the non-central projection case each image ray has its own space partitioning. Note, as the mirrors do not change, the virtual chambers remain unchanged for different rays. Only the chamber visibility may vary.

Our labeling procedure performs independent computations for different image rays. Which makes it applicable for non-central projection devices without any modifications. The only difference is that the visualization of the situation (in the sense of unfolding) is becoming more complicated. On the other hand, if the camera or projector imaging properties are close enough to being central-projective and if the image rays are sufficiently bounded, we can still approximate the system with a single space partitioning from the average projection center. The theoretical foundations for such possibility are established in Part II of the thesis, where the mirror systems were shown to be L -uniform continuous under the move of the point of observation. We will use this fact in Section 6.3, where we employ a non-single projection center galvanometric laser scanning system for illumination.

5.7 Segmentation

Our visual hull estimation approach is based on a binary segmentation of the camera image into background (gaps between objects) and foreground (rest of image).

The three properties of our visual hull estimation approach, mentioned in Section 5.3 require our segmentation to be precise in determining the background pixels to avoid damaging the kaleidoscopic visual hull (see Fig. 5.4 (e)-(f)-(g)). On the other hand, we often may leave many background pixels in the foreground class with only minor effect on the KVH quality, see Fig. 5.4 (b). In other words, incorrect silhouettes should never underestimate the object since this leads to erroneous space carving results, while coarse silhouette estimates can in fact be used with such system. Moreover, some of the pure foreground pixels, where the object is visible in some virtual chamber, may be assigned to the background if this chamber and all the chambers succeeding it are not considered by our volume carving algorithm for that particular image pixel.

5.8 Experiments

It should be noted that the tools developed so far are applicable to any system of planar mirrors. In this section we concentrate on a specific kaleidoscopic setup to verify our theoretical developments in practice.

5.8.1 Setup Details

We chose the frustum of a triangular pyramid as our base chamber, the narrow end pointed downward, see Fig. 5.5.

We chose the opening angle of the system to permit up to 10 levels of reflection ($\approx 6.8^\circ$). The setup was recorded by a Canon 5D Mark II equipped with a Canon EF14mm f/2.8L USM lens. The smallest aperture of this lens is $f/22$. We set it to $f/20$ for our experiments to have large enough depth-of-field and to limit diffraction effects. The mirrors are optical front-surface mirrors of a trapezoidal form, see Fig. 5.5 (right), in the low end of the price range ($\approx \text{€}200$ for three of them).

After the calibration of the intrinsic, extrinsic and radiometric parameters of the camera, the geometry, and the reflectivity coefficients of the mirrors, Section 3.3, our system is ready to capture the scene.

5.8.2 Space Carving Implementation

We implemented our space carving scheme as a voxel-based algorithm. The voxel grid is defined in the base chamber. To account for diverging rays we




Figure 5.4: (a) Image of an object (Emily) inside our kaleidoscope. (b) Binary segmentation of the image (black is the background). Note, that the segmentation is far from the perfect one, but the resulting kaleidoscopic visual hull (c) and labeling (d) are good. Moreover, most of the remaining unsegmented background pixels do not contribute to the visual hull estimation anymore. (c)-(d) Kaleidoscopic visual hull and labeled image, based on segmentation (b). (e) Binary segmentation of the image with one small region incorrectly marked as background. (f)-(g) Kaleidoscopic visual hull and labeled image, based on segmentation (e). The small region misclassified as background "damaged" the kaleidoscopic visual hull and the labeling with a hole in one of the wings of Emily. It is not straightforward to fix this hole if the incorrectly segmented area is unknown. This is due to the fact, that there are multiple potential places (see, for example, the number of labels in (g) having a hole in the wing now) in the segmented image, that could have caused the incorrect result.



Figure 5.5: Photograph of our practical setup, left. A view inside showing the mirror space, middle. Draft of one of the three identical mirrors in the kaleidoscope, right.




Figure 5.6: In this figure the estimated mirrors are superimposed as attenuation layers (the original image is shown in Fig.2.5 (middle)). Notice, that the mirror planes are not connected (dark stripes) due to the safety regions.

implemented a sub-sampling scheme for the pixels in the camera view. We also exclude a region of $\epsilon \approx 5\text{mm}$ around the mirrors borders (safety regions) to avoid erroneous ray paths due to potential errors in the estimation of the mirror plane parameters, see Fig. 5.6. To achieve the most efficient use of the available voxel resolution, the size of the bounding box of the voxel grid is chosen to enclose the object tightly. The number N of reflection levels to be used for visual hull reconstruction and pixel labeling is a user parameter to our algorithm.

5.8.3 Results

We recorded our images at a resolution of 3866×2574 pixels. The data sets as well as the computed kaleidoscopic visual hulls, labeling and radiometrically corrected images are shown in Fig. 5.7. Table 5.1 summarizes some statistics. All results have been computed only using the silhouette image and the calibration information. The objects cover different sizes and




Figure 5.7: Results from top to bottom: input image, silhouette image, kaleidoscopic visual hull, corresponding labeling, reliable pixels, radiometrically compensated image.

object	refl.	KVH discretiz.	views	labeled	unrel.	occ. bnd.
Cone	7	256×256×236	128	18.38%	2.31%	0.08%
Cone	8	256×256×236	166	19.60%	4.44%	0.12%
Cone	9	256×256×236	212	19.91%	5.04%	0.20%
Duck	7	256×256×221	128	15.16%	1.44%	0.06%
Duck	8	256×256×221	166	15.94%	3.46%	0.12%
Duck	9	256×256×221	212	16.24%	5.97%	0.19%
Pumpkin	7	300×300×346	128	25.70%	5.89%	0.16%
Pumpkin	8	300×300×346	166	27.20%	8.94%	0.21%
Pumpkin	9	300×300×346	212	27.79%	12.86%	0.26%
Vase	7	256×256×354	128	29.02%	10.64%	0.25%
Vase	8	256×256×354	166	30.41%	13.49%	0.33%
Vase	9	256×256×354	212	31.08%	15.29%	0.39%

Table 5.1: Statistics for the different data sets. From left to right: name of the data set, number of reflection levels used to compute result, discretization of the kaleidoscopic visual hull, number of virtual views that have been used for computation, the number of labeled pixels, unreliable pixels as a percentage of labeled pixels, and number of pixels in a 3-pixel error region to each side of an occlusion boundary. We used 9 sub-samples per pixel for all results shown here.

vary drastically in silhouette quality. The results in the figure were computed using the information from 8 levels of reflection, equaling 166 views. For 9 reflection levels results deteriorated due to insufficient accuracy of the silhouette estimation (chromatic aberration) and residual radial distortion, see Sect. 5.9. The dark streaks in the images are due to the safety region around mirror boundaries. As can be seen in the results, we achieve almost pixel-accurate labeling results: if errors occur, they are conservative, i.e. to the inside of the object. Even if sub-optimal, the labeling is thus still correct. Overlaps between virtual views can be well resolved. The table shows that the unreliable pixel measure is overly conservative and that a heuristic can yield a much higher usable pixel count. Since the objects do not have narrow holes this approach appears reasonable. Further and high-resolution results can be found on our project web page² in the supplemental materials document.

5.9 Limitations

In this section we would like to discuss the physical limitations of kaleidoscopic imaging systems.

²http://ilya.o-x-t.com/kaleidoscopic_imaging/

In our experiments, see Section 5.8, we are operating the camera at the physical limits of the optics. The most important aspect is the limited depth-of-field of real camera systems. Since in our experimental setup the virtual views cover a very wide depth range, both very close to the camera for the imaged real object and very far away, for the distant reflections, a suitable trade-off has to be found. Also, we use a very wide field-of-view to cover the chamber structure imposed by our prototype system. This implies that the higher-order reflections which carry significant silhouette information, since the object is seen from the side in these views, are imaged at the periphery of the image. Thus, any uncorrected radial distortion and in particular other optical aberrations found in these regions are of major concern. Chromatic aberration was a major problem in our experiments. The correct silhouette boundary is not discernible even by a human observer. The spread of the aberration is up to 7 pixels, a much higher value than the geometric calibration error. Use of a different lens (Canon EF 14mm f/2.8L II) reduces this problem.

An inherent property of our data is that they are multi-resolution. Pixels in distant images of higher reflection levels cover a larger surface area of the object. This apparent disadvantage might turn out to be a useful feature in future multi-view reconstruction algorithms. In any case it is a particular property that cannot be achieved easily in different systems.

Finally, scaling a setup might present a problem. However, e.g. Science World in Vancouver is operating a kaleidoscope at human proportions. Foil mirrors can easily be produced up to a size of 4 – 5m. We thus believe that planar kaleidoscopic systems are applicable to real problems of interest.

Application Examples

Overview

As mentioned in the previous Chapter, once the labeling for a camera and/or projector image is available, it can easily be used by classical multi-view algorithms. In this section, we will show three examples of such a use - visual hull estimation, geometry reconstruction using a classical structure from motion technique, and simultaneous geometry and reflectance estimation in a laser scanning setup.

6.1 Visual Hull Approximation

The kaleidoscopic visual hull, which we used in the previous Section for labeling, can be directly used as an approximation to the object's visual hull. This possibility is quite attractive as (besides the necessary calibration) only segmentation is required in order to construct the KVH. However, the restriction that the pixels marked as background have to be absolutely correct and the complex appearance of kaleidoscopic scenes makes the application of state of the art fully automatic segmentation methods extremely challenging. Moreover, if some of the foreground pixels are incorrectly classified by such an automatic algorithm, it is very difficult to locate them afterwards. Thus, the applicability of automatic segmentation algorithms is very restricted even under the final supervision by a human. For humans, the segmentation task is also a non-trivial and tedious procedure.

While it is still possible to produce a reasonable automatic segmentation

for a particular scene or even for a particular movie (see the labeling videos of the moving flower example on our project page¹), it requires considerable tuning and is not applicable as a general solution. Instead, an interactive segmentation tool, where the user can be provided with feedback immediately (in both, the 2D image domain, and in the 3D domain of the visual hull) can be very useful.




Figure 6.1: Showing the application in action: The user draws in the window on the right side. Already drawn patches are rendered with a green checkerboard pattern and are listed in a widget at the right border. The current visual hull and a preview of the current brush are rendered in the left window.

We have developed such a tool [Klehm12]. Our interactive application enables a fast manual foreground/background segmentation of kaleidoscopic images for high-quality results by guiding the user and giving constant feedback. It includes an on-line technique for estimating the visual hull of an object inside a system of arbitrarily positioned planar mirrors. The application provides feedback to the user in the 2D kaleidoscope domain, see Fig. 6.1 (right), as well as in the 3D domain of the visual hull, see Fig. 6.2. In the application we handle the complex computations of drawing, rendering, visual-hull derivation, and labeling by exploiting the GPU. We also explore suitable interaction schemes and visualization techniques that guide the user in rapidly creating an accurate segmentation, Figs. 6.1 and 6.2.

6.2 Multiview Stereo Reconstruction

For our example we used the photometrically compensated image and the corresponding labeling of Fig. 6.3. We applied the software VisualSfm [Wu11a, Wu11b] to this data.

¹http://ilya.o-x-t.com/kaleidoscopic_imaging/



Figure 6.2: Semi-transparent mode (left): Visual-hull voxels are shown in white with a low density. Outside voxels covered by the brush are shown in blue. Finally, visual-hull voxels covered by the brush are shown in red. Image-based shading (middle and right): Pixels are transferred from the kaleidoscope image.




Figure 6.3: Images, used in our experiment: photometrically compensated kaleidoscopic image (left), corresponding labeling (right).

The input set of 98 single-view input images is produced from the compensated kaleidoscopic image, Fig 6.3 (left), by keeping only those image pixels that are corresponding to a single chamber object, while filling the remaining pixels black. Images corresponding to odd levels of reflection are flipped in order to have the same orientation for all images. This flipping is necessary for SIFT feature matching [Lowe04] as SIFT features are not reflection invariant. Using the VLFeat library [Vedaldi08], we then compute SIFT features for the image set generated this way. We exclude those pixels that are outside the object silhouette or on the border of our labeled region. This step allows us to get rid of incorrectly matched feature points that do not fully belong to the object. Finally, because all our virtual cameras have the same intrinsics, we use our calibrated camera intrinsics for all virtual views.

It is worth mentioning, that, since virtual camera positions are estimated by matching SIFT features and bundle adjustment, there is no need to




Figure 6.4: Reconstructed point cloud (5492 points)

provide VisualSFM with the extrinsic parameters of the virtual cameras.

The output of VisualSFM, a point cloud of the object’s surface consisting of 5492 points is depicted in Fig. 6.4.

Note that the example here is not intended to improve on state-of-the-art results in terms of reconstruction quality, but rather to show how easily the kaleidoscopic image can be used in multi-view applications once it is labeled.

6.3 Surround Geometry and Reflectance Acquisition

As an additional application, we describe a system for acquiring reflectance fields of objects without moving parts and without a massively parallel hardware setup. Our system consists of a set of planar mirrors which serve to multiply a single camera and a single projector into a multitude of virtual counterparts. Using this arrangement, we can acquire reflectance fields with an average angular sampling rate of about 120+ view/light pairs per surface point. The mirror system allows for freely programmable illumination with full directional coverage. We employ this setup to realize a 3D acquisition system that employs structured illumination to capture the unknown object geometry, in addition to dense reflectance sampling. We demonstrate for a number of test scenes that the kaleidoscopic approach can acquire complex reflectance properties faithfully. The main limitation is that the multiplexing approach limits the attainable spatial resolution, trading it off for improved directional coverage.

6.3.1 Operational Principle

Our system is built around imaging and projection within systems of planar mirrors. It exploits the multitude of reflections and inter-reflections to compress a full surround field of view into the aperture of a single camera and/or projector. The basis for our design is the analysis of planar mirror systems described in Part II of this thesis. Due to Helmholtz reciprocity, this analysis applies to cameras as well as projectors, see Section 5.5, with one difference:

Simple scene illumination by a projector leads to superposition of light due to different mirror reflections of the light source, which is not desirable.

Point Scanning in Planar Mirror Systems

Consider a two-dimensional mirror system as shown in Fig. 6.5. On the top, we illustrate how a projector is illuminating a single point on the surface of an object. By unfolding the system we can visualize the mirror world with a number of virtual objects. The illuminated point can be considered to be present on all copies of the object. The camera can then observe those illuminated points that are not occluded. The four points observable by the camera are e.g. generated by the reflection sequences (from left to right) $(1, 2, 1)$, (1) , $(\)$, and $(2, 1)$. These numbers indicate the mirrors that are being traversed by the ray before it hits the object.

Alternatively, we can consider the system as consisting of virtual cameras and projectors observing the real object as shown in the middle of Fig. 6.5. In our example, the virtual cameras observe the illuminated point via the reflection sequences (from left to right) $(1, 2, 1)$, $(2, 1)$, (1) , and $(\)$. As expected, these sequences are the same as in the previous case. This alternative interpretation immediately shows that a pair of reflectance values is sampled simultaneously: For one projector illumination direction, we obtain reflectance samples from four viewing directions. Furthermore, since multiple viewing rays are available that observe the same object point, we can triangulate the point and obtain part of the object geometry. Finally, the same object point can be illuminated by different virtual projectors, one at a time. This results in $N_k \times M_k$ reflectance samples per object point \mathbf{x}_k , where N_k is the number of unoccluded views and M_k the number of unoccluded illumination directions. In our example, this yields 20 reflectance samples for the illuminated point.

Of course, point sampling an object is inefficient. It would require the acquisition of $\mathcal{O}(K \times M)$ images for sampling K surface points, and $M = \max_k M_k$. Key to our method is the use of structured light in the present context. The main challenge is that we have to avoid illuminating a surface point from more than one direction in a single image of the scan since this would lead to the superposition of light and thus to summed reflectance




Figure 6.5: Illustration of imaging and projection inside a mirror system.

samples.

Structured Light in Planar Mirror Systems

As shown above, a camera image of a planar mirror system can be decomposed into a number of sets of image coordinates for which a certain virtual camera view is valid, Fig. 6.5 (bottom). These image regions are camera labels. A virtual camera in a virtual multi-view system is then given by this set of image coordinates in conjunction with the calibration parameters that define the projection operation of the virtual camera. Applying the Helmholtz reciprocity principle, we describe virtual projectors, similar to virtual cameras, as a set of coordinates of the original projector (not shown). A virtual projector is given by a projector label, i.e. a set of projector coordinates for which the virtual projector is valid, in conjunction with its calibration parameters that define the projection operation for the virtual projector.

By illuminating only a single projector label at a time, we are able to decompose reflectance acquisition into non-overlapping regions of illumination.

The camera and projector labels enable the decomposition of imagery, obtained from projecting light into a planar mirror system, into different light/direction pairs. Each combination of camera and projector labels gives rise to a unique viewing/lighting direction pair. We are thus able to sample the reflectance field of an object using $N \times M$ (sub-)images. In practice, we achieve an average of about 120 unoccluded view/light directions per surface point. Without exact surface geometry which serves as an interpolation guide, parallax effects, both for view and lighting interpolation become apparent. We therefore opt for structured light projection which enables geometry reconstruction on top of reflectance sampling.

In the following we discuss the hardware design that was used for our experiments.

6.3.2 Hardware design

Our hardware setup consists of four planar mirrors, one digital SLR camera, and one RGB laser projector with a double axis mirror galvanometer for control. Next we discuss our design choices regarding these components.

Mirror System: We employ a kaleidoscopic mirror system consisting of four mirrors: three of them are arranged in the classical configuration of a truncated three-sided pyramid as in [Han03]. As described there, this type of system can only generate hemispherical view distributions which is sufficient to perform exhaustive reflectance sampling if the sample is flat. However, non-flat geometry requires sampling of the full sphere of viewpoints in order to obtain dense coverage of both, surface geometry and reflectance.




Figure 6.6: Images of our acquisition system: Left: the laser/camera system, the red line indicates a laser ray that can be steered inside the mirror system. Middle: an outside view of our system with dimensions. Right: the back mirror is attached to the back entrance of the system once an object is mounted inside.

To enable spherical sampling, we add one additional mirror at the back entrance of the system which generates reflections that also show the object from below, see Fig. 6.7. The mirrors are optical quality first surface mirrors. The mirror on the back side has a hole cut by a water jet. This hole enables an object holder assembly to protrude inside the system. Images of our system are shown in Fig. 6.6.

Projector: For our projector we use a system of coaxial red, green, and blue lasers with approximately 5mW output power per channel. The choice of using a laser system is motivated by the much larger depth-of-field of a laser beam as compared to a standard digital projector: Since our system folds the rays into the mirror system, distances between the apparent closest object (no reflection) and the farthest virtual objects is quite significant (about 4m). Since we need to be able to control the scan in both, x- and y-directions we are using a two-mirror galvanometer scanner.

To cover the large field of view inside the mirror system, we employ a wide angle two-axis mirror galvanometer (Thorlabs GVS012) with an opening angle of 80 degrees in both dimensions. The projector is mounted close to the entrance of the mirror system to make full use of the aperture of the mirror system, i.e. no rays miss the opening at the front side.

Camera: We employ a Canon 5D mark II DSLR equipped with a Canon EF 14mm f/2.8 L II USM lens. This lens has a large field of view and comparatively minor radial distortion and chromatic aberrations. As a compromise between light efficiency and depth-of-field, we use an aperture setting of $f/14$. The camera is mounted such that it is close to the entrance of the mirror system while simultaneously avoiding to image parts of the laser projector.




Figure 6.7: *Top*: Image of an object inside our acquisition system. The inset shows that both, views from the top and views from below are generated by our system. Side views from different directions and at different elevations are visible towards the boundary of the image. *Bottom*: Virtual camera distribution in our setup. The red box is the bounding box of the object, blue cones are virtual cameras in the upper hemisphere, cyan ones are located in the lower hemisphere. The camera marked in green is the real camera location. Units are given in mm.




Figure 6.8: Visualization of the camera (left) and projector labeling (right). Different colors correspond to different labels. The plot simultaneously shows the viewpoint and lighting direction coverage of our system.

6.3.3 Implementation

The labeling computation is based on the segmentation of a flood-lit image of the object positioned inside the mirror system taken with the camera, see Fig. 5.4 (a)-(b) for an example. Using this image, the kaleidoscopic visual hull of the object can be estimated, Fig. 5.4 (c), see also Section 5.3. This geometry is then used to determine the camera and projector labeling. We perform the actual computation of the labels by ray-tracing until the ray intersects the kaleidoscopic visual hull, recording the mirror intersections that occur along the way. Pixels with equal mirror sequences belong to a common chamber and are assigned a common label. This implies that a pixel can only belong to one camera or projector chamber, respectively. The procedure is the same for the virtual cameras and the virtual projectors. The result can be visualized in an image as shown in Fig. 6.8.

After labeling both the camera and the projector, we can prepare a capture session. For this, we scan-convert the projector labels into horizontal scan-lines which we refer to as segments. Due to the geometry of our galvanometric scanning system, this ensures that the scanned laser sheets correspond to three-dimensional plane segments. Restriction of the scan-lines to projector labels guarantees that object points are illuminated only from a single direction and that the plane segment does not split up regardless of the mirroring sequence traversed until the object is hit. Each virtual projector is assigned a list of segments. The total number of segments is typically on the order of 20.000, see Table 6.2. We capture one HDR image for each segment of all virtual projectors.

When capturing is done, we proceed to the final steps: geometry and reflectance estimation. The overview of our pipeline is shown in Fig. 6.9. The implementation details are given in [Ihrke12b].

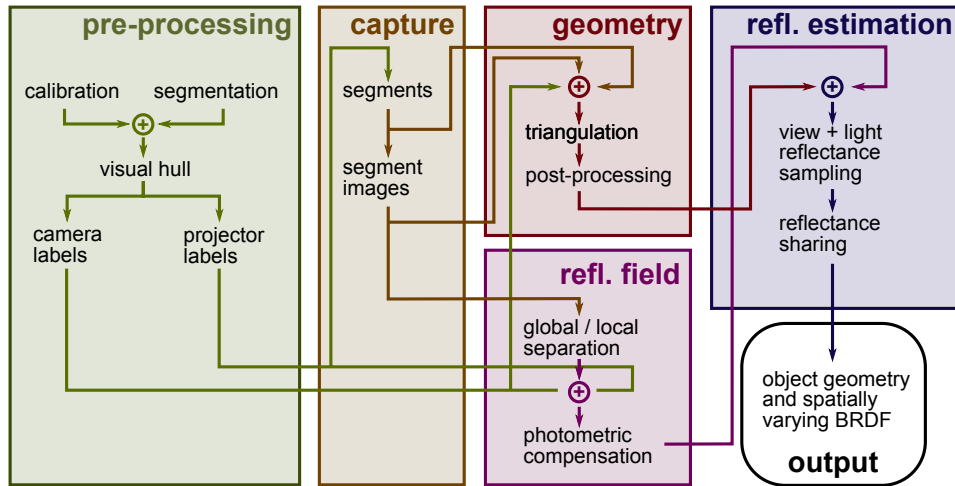


Figure 6.9: Overview of our reconstruction pipeline.

dataset		min	max	$\sqrt{\text{max}}$	avg	$\sqrt{\text{avg}}$
Angel	cam	1	28274	168	7718	88
	proj	32	35947	190	9299	96
Emily	cam	10	20963	145	5744	76
	proj	3	29946	173	6536	81
Max Planck	cam	1	23732	154	7838	89
	proj	19	32283	180	9344	97

Table 6.1: Number of camera and projector pixels for the virtual cameras and projectors. This statistic counts object pixels only, i.e. background is ignored. We give the minimum, the maximum, and the average numbers. The square roots give the side length of an equivalent square of pixels.

6.3.4 Results

We have tested our system on the three objects *Angel*, *Emily*, and *Max Planck* shown in Fig. 6.11. Two of the objects, *Emily* and *Max Planck*, are almost single material, metallic objects, whereas the *Angel* data set shows a larger number of different materials.

For the *Max Planck* data set, a high resolution range scan is available from the Aim@Shape repository. It was acquired by coating the object with white powder and range scanning it with a Minolta vi910 which is at least an order of magnitude more accurate than our system and can be considered as ground truth. In order to evaluate our reconstruction pipeline we registered the geometry recovered by our technique with the ground truth scan using rigid ICP. The error distribution is shown in Fig. 6.10. The accuracy of our system, with an average error of about 5 – 6mm, is low by current standards. However, taking into account the actual resolution achieved by our virtual




Figure 6.10: *First row:* Geometric quality of our recovered geometry as compared to laser scanned geometry (object coated with white powder). Most of the errors are below $6mm$, the gap behind the ear could not be resolved. *Second row:* A comparison of appearance computed with our approximate geometry (left) and with laser scanned ground truth geometry (right). Details appear less blurred in the accurate geometry version.

dataset	#cam	#proj	# segs	avg	max samples	total
Angel	246	144	20177	223	901	31.9 million
Emily	235	139	18863	133	815	20.8 million
Max Planck	239	134	14483	123	606	17.7 million

Table 6.2: System properties for the different data sets. *#cam* and *#proj* are the number of virtual cameras and projectors in the system, *avg* samples is the average number of view and lighting directions per surface point. *Max samples* reports the surface point with the maximum number of reflectance samples. *Total samples* shows the overall number of samples used for reflectance estimation.

multi-view/multi-projection system, the results can be interpreted as quite satisfactory, see Table 6.1: Multiplexing onto a 5 MPixel sensor results in a very low number of pixels that are available for every virtual view and projector. The equivalent square numbers are provided to give an impression of the image size of the object in each virtual view. The fill rate of our sensor with useful information is between 20 and 25%.

In another experiment, we investigated the impact of the low resolution geometry as output by our system. We registered the available ground truth geometry of the Max Planck data set with our reflectance field. As can be seen in the images, more detail is preserved with the more accurate geometry. This effect can be attributed to parallax effects if the geometry is expected in the wrong position. Thus, improving the geometry, e.g. by separately scanning it as in [Schwartz11] might be a viable option for reflectance scanning with a kaleidoscope.




Figure 6.11: Results for our data sets *Angel* (top row), *Emily* (middle row), and *Max Planck* (bottom row). The left two columns show the distribution of the number of unoccluded reflectance samples (view/light pairs) on the object surface. The center right column shows the objects illuminated from a single direction (*Angel* and *Emily*) and from four directions for the *Max Planck* data set. In the right column, we show the objects illuminated by environment lighting.

Conclusions

In this part of the thesis we proved that the decomposition of a kaleidoscopic image into different viewpoints is possible and accurate when the specific visual hull $VH(O_1, V^*)$ of an object of observation is known. The requirement to have the visual hull of an object of interest is, in many cases, impractical. For this reason, we proposed to use the kaleidoscopic visual hull, which can be easily computed if the kaleidoscopic image is segmented into background and foreground, instead of the visual hull. Additionally, we developed a kaleidoscopic segmentation software which makes the production of the kaleidoscopic visual hull very comfortable. Finally, we demonstrated the feasibility of our techniques in a number of different multi-view applications.

In the next part of this thesis we turn to a different problem: the recovery of the structure of planar mirror systems.

PART IV

Structure Reconstruction of Planar
Mirror Systems

Discovering the Structure of a Planar Mirror System from Multiple Observations of a Single Point

7.1 Introduction

In the previous part of the thesis we described some novel applications of kaleidoscopic systems in the multi-view area of computational photography. However, all these applications strongly rely on the geometric calibration of mirrors. Thus, if we want to apply the techniques from Part III in practice, we need sophisticated methods for the geometric reconstruction of complex mirror designs.


This task is complicated by the fact that we typically do not have access to the chamber information at this point. Tedious human intervention is therefore necessary to perform the calibration, see Section 3.3.1.

In this part we aim at developing a fully automatic technique for determining the geometry of multiple planar mirrors and the pose of the camera with respect to this mirror configuration. Our work is showing the feasibility of solving the problem in a $2\frac{1}{2}D$ convex room setting. Because we aim at demonstrating the principal feasibility, we consider a minimal case. Using only the measurements of apparent depth of a single scene point that is visible via many different multi-bounce interactions with the mirrors, we develop an algorithm to identify the geometry of the room, and the position and orientation of the camera. Our setting is considerably more general

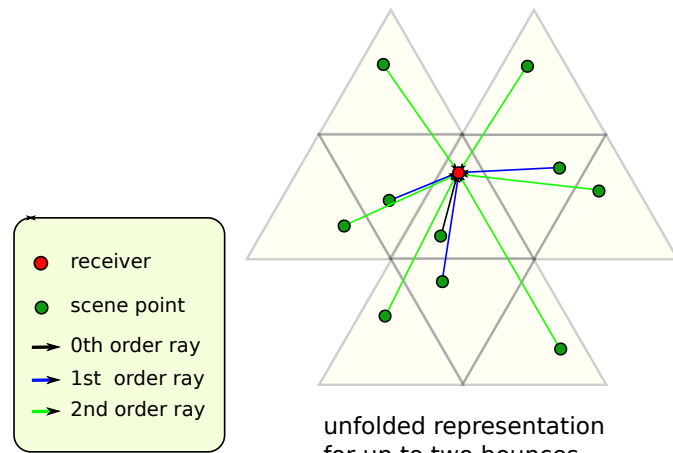
than state-of-the-art solutions [Ribeiro12, Antonacci12, Tervo12] that require zero- and first-bounce reflections to be observable and that can deal only with a very limited class of mirror shapes in practice.

7.2 Problem Formulation

We consider an abstract angularly resolved distance measurement system in a situation with specular multi-bounce paths, Fig. 7.1 (a). The scene consists of a room with specular walls and one scene point (dark green), as well as a recording device (red). The interpretation of the left-most sub-



(a)



(b)

Figure 7.1: Problem definition: (a) A point is visible to an omni-directional receiver via multiple specular reflections. The receiver is capable of measuring the incidence angle and the distance of the point via the reflected ray paths without being able to detect the specular interaction. (b) To the receiver, the situation appears as if there are multiple points at different distances from its own position. The mirroring room geometry (faint yellow) is unknown and has to be recovered from the point measurements.

image is difficult. It contains all ray paths up to second order that hit

the receiver. For this reason the following sketches show the ray paths of different reflection orders separately. We assume that the receiver is capable of measuring the angle of incidence of the rays as well as the apparent distance of the point along each reflected ray path. The task is to reconstruct the positions of the unknown mirror planes and to locate the receiver with respect to the mirror system.

In Fig. 7.1 (b), we show how the receiver could naïvely interpret the surrounding world as a virtual mirror world consisting of many point objects at different distances. This interpretation corresponds to the unfolding operation, see Part II, if the mirror geometry is known. Instead, we ask for the mirrors geometry given the unfolded representation. In a sense, Part III of this thesis describes the solution to the dual problem: Given the mirror geometry, compute the object. In the current part, we consider the object geometry to be given (a single point) at different virtual locations but the mirror room geometry has to be recovered. In the following we will present all analyses in two dimensions. This restricts our practical examples, Sect. 7.7, to $2\frac{1}{2}D$ cases.

Here and later in the text we consider only rays starting from the observation point \mathbf{c} (red dot in Fig. 7.2). Thus, to interpret the rays behaviour, the unfolded representation can be considered, see a simulation in Fig. 7.2. In this picture, polygonal regions illustrate areas of the corresponding chambers that are visible from the observation point. The the base chamber is marked in yellow.

As discussed above, in Section 3.1 and Part II, any real mirror system will fail to subdivide the plane perfectly when generating the unfolded representation of the mirror world. Therefore, a real system includes lines of discontinuity as illustrated in Fig. 7.2. Areas between two discontinuity lines indicate a common reflection sequence. The discontinuity lines correspond to mirror corners being hit by a ray bundle after a sequence of reflections. The bundle splits up at these points and traverses different mirror sequences thereafter. This implies that the measurement points are not necessarily visible for every possible combination of mirrors into a sequence. As an example, there is exactly one ray among those that traverse the mirror sequence (2) that hits the object point. There is no ray amongst those continuing to mirror sequence (2, 3) that sees the object point after these two reflections (i.e. the object is not in the visible part of the chamber $C_{2,3}$). However, the bundle of rays continuing to mirror sequence (2, 1) contains such ray. Moreover, the invisibility of the object point after a certain subsequence of reflections does not imply that the point will be invisible in the future as the continued reflections sequence (2, 3, 1) shows. After exactly three reflections, one of the rays with this reflection sequence hits the object point.

In the following, we will discuss a solution for recovering the mirror room geometry as well as the position of the receiver within such system.




Figure 7.2: Simulation of a general four-sided room geometry. We show 4 levels of reflection. The geometry is greyed out to emphasize that we only consider the point distribution to be available. The numbers indicate mirrors in the base chamber (faint yellow) and their respective reflections. The pink lines mark lines of discontinuity that split areas of different mirror sequences. The room geometry and these lines have to be predicted in conjunction in order to use the data points as measurements. Again, red is the receiver and the virtual point objects are marked green. The spatial locations of these points are the only input to our method.

Our method allows for the recovery of convex room geometries with an arbitrary unknown number of mirror walls. The virtual point positions and the receiver coordinates are the only input to our algorithm. In particular, the reflection levels through which the virtual points are seen are assumed to be unknown. We also do not require any particular reflection (for example the direct view) to be available. We demonstrate the performance of our algorithm through simulations and a challenging real-world example.

7.3 Overview

In order to discuss our recovery algorithm it is necessary to introduce a few definitions as well as to explore some properties of mirror systems. Our main tool for recovery is a validation procedure: Given a candidate config-

uration (consisting of mirror geometry, observer position and scene point position), determine if this configuration is compatible with the observations. We describe this part of the algorithm in Section 7.4. We derive sufficient conditions for a configuration to be reconstructable. A naïve algorithm then consists in checking all possible configurations which we term the exhaustive search algorithm, Section 7.5. This algorithm operates on a graph structure, that, in the absence of all optimizations, is a fully connected graph. The computational costs of performing this search are exponential and we develop theoretically sound graph pruning strategies that ensure that no false negatives occur, Section 7.5.1. The exhaustive search algorithm can be executed on the pruned graph, ensuring that a solution is found if the configuration is reconstructable in our sense, however, at a significant cost. For this reason, we introduce a heuristic search algorithm in Section 7.6 that is based on random graph sampling to improve the reconstruction performance. Finally, we validate our algorithm in Section 7.7 via simulations and a real experiment.

7.3.1 Problem Setting

Our algorithm is based on a number of assumptions about the scene that directly inform the constraints we may apply and the algorithmic strategies we employ. In particular, we consider

- the room geometry to be convex,
- the room and scene to be essentially planar, and
- the scene to consist of a single object point.

This list allows for rooms with gaps in the mirrors, e.g. to place a camera. The planarity constraint permits rooms with walls that are orthogonal to a common ground plane and a common ceiling (which may be mirrors) while having a convex layout in the ground plane. The single object point constraint avoids a matching procedure to identify images of a common world point.

For the discussion of the basic algorithm idea we introduce some further restrictions. These are didactic in nature and will be relaxed later, Sect 7.6.1:

- the room geometry is closed,
- the room geometry is irregular,
- the object point is in a general position with respect to the mirror planes, and
- the perspective center of the receiver is inside the convex hull of the room.

By a closed convex room we mean a room that is equal to its convex hull. Irregularity relates to asymmetric configurations. These are easier to deal




Figure 7.3: An overview of our approach. *Left*: Doublets are pairs of points that are separated by a single mirror reflection. They are indicative of mirror planes. Equivalent doublets are coded with the same color. As can be seen, doublets usually occur several times, making them stable features of the system. Moreover, combinations of doublets are repeated throughout the system. The sketch shows 4 reflection levels. *Middle*: Doublets can be joined into triplets that are indicative of room corners. By identifying common doublets, these triplets can be joined together in an iterative process. This way, a candidate room geometry can be recovered in some virtual location. *Right*: In order to verify the correctness of a particular configuration, we reflect the reconstructed geometry along a line of sight (black) until it contains the receiver. The sequence of reflections is shown in a color-coded fashion. In this position, the candidate geometry serves as a base chamber of the mirror system from which a representation of the mirror world can be computed by an unfolding procedure. This last step enables the comparison of predicted and recorded point positions and therefore the validation of the candidate geometry.

with initially because mirror walls cannot be exchanged by an invariant transformation that keeps the mirror system apparently unchanged. Further, we require that mirrors produce unique actions on the object point. This condition is satisfied if the point is at different distances to each of the mirror walls. The condition that the perspective center of the receiver is within the mirroring room is needed to uniquely predict the pose of the receiver with respect to the reconstructed mirror geometry.

We emphasize that we do not require the number of reflections for a particular ray path to be known and that we do not need to observe special identifiable reflections like the direct image (0 bounce) or the first order reflections. Our algorithm is designed to be agnostic to this information. While this type of information would simplify the task considerably, we aim for a general method that can work with limited data, both in the field-of-view of the receiver and in the depth range that can be reliably measured. We also consider the number of mirror walls of the room geometry to be

unknown.

7.3.2 Definitions

The current discussion is based on an ideal setting where the system layout is known. Under these conditions, we first derive constraints and conditions that allow for the reconstruction of the mirror geometry from the virtual point distribution. Later, in Section 7.6, we then extend our ideas to the case where the system layout is unknown.

Definition 7.3.1. Let \mathbf{a} be a point in the base chamber C_0 , visible in chambers $C_{M_{K_1}M_{K_2}\dots M_{K_N}}$ and $C_{M_{K_1}M_{K_2}\dots M_{K_N}M_{K_{N+1}}}$, $N \geq 0$. Then a **doublet** is a pair of points $D = (\mathbf{a}', \mathbf{a}'')$, where $\mathbf{a}' = m_{M_{K_1}M_{K_2}\dots M_{K_N}}(\mathbf{a})$, $\mathbf{a}'' = m_{M_{K_1}M_{K_2}\dots M_{K_N}M_{K_{N+1}}}(\mathbf{a})$.




Figure 7.4: A doublet $(\mathbf{a}', \mathbf{a}'')$ and corresponding receiver position \mathbf{c} . This sketch is an enlarged version of the marked triangle in Fig. 7.3 (left).

Properties:

- A doublet uniquely defines a virtual mirror $M_{\mathbf{a}'\mathbf{a}''}$ that takes point \mathbf{a}' into \mathbf{a}'' ($\mathbf{a}'' = m_{M_{\mathbf{a}'\mathbf{a}''}}(\mathbf{a}')$). On the other hand, according to the doublet definition, the virtual mirror $M_{\mathbf{a}'\mathbf{a}''} = m_{M_{K_1}M_{K_2}\dots M_{K_N}}(M_{K_{N+1}})$, i.e. mirror $M_{K_{N+1}}$ in the corresponding chamber $C_{M_{K_1}M_{K_2}\dots M_{K_N}}$.
- If we denote the receiver position by \mathbf{c} , then $d(\mathbf{c}, \mathbf{a}') < d(\mathbf{c}, \mathbf{a}'')$ since an indirect view via virtual mirror $M_{\mathbf{a}'\mathbf{a}''}$ has a longer path length than the direct view.
- Doublets never intersect with discontinuity lines since they are indicative of a real mirror.
- Even though doublets define a mirror transformation, it is in general impossible to transfer them to the base chamber (in effect reconstructing a single wall of the room ($M_{K_{N+1}}$ in our example)) without knowing the remaining room geometry.

The doublet in Fig. 7.4 and all other doublets in the system of Fig. 7.2 are shown in Fig. 7.3 (left). Doublets belonging to the same physical mirror are marked with the same color. They usually re-appear in different locations of the virtual mirror world. The properties just mentioned can be verified in the figure.

Definition 7.3.2. A **triplet** is a pair of doublets $T = (D, D')$ that shares a common point. It consists of three observed points.

Properties:

- A triplet defines the angle between two mirror planes. This often is a corner of the mirror room reflected to some position in the virtual mirror world. It may happen though that two doublets that do not correspond to directly adjacent mirrors form a triplet. Nevertheless, the angle between these two mirror planes is fixed by the triplet. We will refer to the doublets that constitute the triplet as its *legs*.
- As in the case of doublets, a triplet can typically not be transformed to the base chamber without knowing the remaining geometry of the mirror world since an unspecified sequence of reflections lies between the observed position and the canonical position of the legs in the base chamber.

Some examples of triplets in conjunction with the mirror corners defined by them are shown in Fig. 7.3 (middle). It should be noted though that all adjacent doublets form triplets even though we only show a subset of them.

7.4 Verifying a Candidate Configuration

The core of our algorithm is based on being able to verify a given configuration against the observed data. The basis for the verification step are the triplets just defined. They serve as building blocks in constructing candidate configurations.

7.4.1 Joining Triplets

Given that a triplet defines the relative position and orientation of two walls of the mirror room it is natural to attempt to join them into quadruplets, quintets, sextets, and so forth, until the relative position and orientation of all mirror planes with respect to each other has been fixed. This would constitute a reconstruction of the room geometry.

For this scheme to work we need to be able to relate different triplets that are observed in different locations of the virtual mirror world. Two triplets fit together if they have one common leg, i.e. they share a doublet.




Figure 7.5: Joining procedure

In general, the different observations are related by an odd or an even number of reflections through the base chamber which is unknown. However, an even sequence of reflections is equivalent to a rotation, and an odd sequence is equivalent to a rotation and a flip. Given two triplets with a common leg, we therefore have two options of joining them resulting in two candidate quadruplets that fix two potentially consistent relative positions and orientations of three mirror planes.

The joining process can be continued by joining a quadruplet with another triplet, yielding a quintet.

In the current discussion we assume that doublets and triplets that are being observed in different locations can be identified and that no erroneous doublets or triplets exist. We will discuss the extension to the case with erroneous information in Sect. 7.6. In practice, the identification of doublets is based on their length (distance between points) which is twice the distance to the corresponding mirror plane. Since we assumed that the object point is in general position with respect to all mirror planes this identification can easily be performed, yielding equivalence classes of doublets. Equivalence classes of triplets are formed by considering the two constituting doublets. In Fig. 7.3 (left,middle), the equivalence classes are color-coded with doublets of the same color, and triplets of the same color pair, belonging to the same class, respectively. In the following, we will drop the explicit mention of the equivalence classes, simply referring to them as doublets or triplets, it should be understood, however, that individual doublets or triplets are only representatives of their class.

7.4.2 Conditions for Reconstructability

A necessary condition for the proposed algorithm to work is that all doublets are being observed by the system, i.e. all mirrors in the room must be observed by their action on a set of two points. In addition to that, a sufficient number of triplets must be observed in order to recover the complete geometry. Consider the case of N doublets (which in the perfect case considered here corresponds to exactly N mirror walls in the room geometry), then the *minimum* amount of triplets that could yield a solution is $N - 1$. This is the case if the triplets can be joined in a sequential manner as indicated in Fig. 7.5.

The amount of triplets that *certainly* yields a solution is $(N - 1)(N - 2)/2 + 1$. Consider a graph structure that we refer to as the *doublet graph* where the nodes are doublets and there are edges if a triplet with the two doublets in question exists. The condition for the room to be recoverable is that a connected component covering all nodes exists. The meaning of this is that all mirrors can be related to one another via pairwise relative position and orientation. In the worst case we have a fully connected component of $N - 1$ nodes with a single unconnected node. This structure has $(N - 1)(N - 2)/2$ edges. If one additional edge is known the complete graph is fully connected.

In practice, the number of triplets is somewhere between these two extremes. The condition for recovery is that the full graph is covered by at least one connected component. In the perfect case, all such connected components are equivalent and yield the same solution.

Summarizing the previous discussion, sufficient conditions for a mirror/source/observer configuration to be reconstructable are:

1. All doublets are observed,
2. the available triplets contain all doublets, and
3. the doublet graph is connected.

7.4.3 Verification Algorithm

The algorithm for joining triplets into a candidate room configuration is then the discovery of connected components. The discovery of one such component suffices in theory. However, each triplet might occur in two flavors, corresponding to whether it resulted from an odd or even sequence of reflections from the unknown base chamber. Denoting by M the number of edges (triplets) in the connected component, there are 2^M possible room configurations in the worst case. If cycles are present, they reduce this complexity. Intuitively, cycles correspond to consistent subsets of triplets that can only have two possible orientations.

Graph theoretical arguments cannot differentiate between these possible solutions. We therefore need a way to verify them by comparison to the actual data.

The idea for verification is depicted in Fig. 7.3 (right). Assume we have constructed a candidate configuration at the position shown in red. We join the virtual object point \mathbf{a}^* in the reconstructed room with the receiver position \mathbf{c} by a line of sight shown in black. The line of sight intersects exactly one mirror plane which must be the mirror that produced the virtual object point \mathbf{a}^* . If the candidate configuration is correct, this mirror is the last in the reflection sequence leading up to the observation of \mathbf{a}^* . We can therefore undo this operation yielding the geometry in green. Since we have assumed that the receiver is inside the convex room, this process can be repeated until the condition holds. In the illustration, this is the case for the yellow position of the geometry.

This position corresponds to the unknown base chamber assuming the candidate configuration is correct. Therefore, we can determine visibility inside the mirror system by unfolding. Since \mathbf{a}^* is transformed to the base chamber in conjunction with the geometry, we can now simulate where the point would appear in the system even if we have not observed the direct view. Comparing the predicted point distribution with the observed one we find whether the candidate configuration is a correct reconstruction.

The verification operation consists of a linear time (in the number of doublet classes, i.e. mirrors) spanning tree computation on the doublet graph and an exponential (in the number of triplets, i.e. mirror corners) determination of the correct flipping configuration.

7.5 Exhaustive Search Algorithm

In the previous section, we assumed a known candidate configuration that could be verified. Unfortunately, we do not have access to which pairs of the input points correspond to doublets. This implies that the triplets are unknown as well. A naïve algorithm for discovering an unknown configuration would have to consider all possible pairs of points, i.e. all potential doublets. Clearly, the number of nodes in the doublet graph is quadratic in the number of observations. Further, the number of all possible triplets is quadratic in the number of nodes in the doublet graph. In the naïve algorithm, the exponential verification procedure, Sect. 7.4.3, has to be performed starting at every possible triplet.

An illustration of all potential doublets in an example configuration is depicted in Fig. 7.6 (left). In this figure, the real doublets are shown in red whereas pairs of points that are no doublets are shown in light blue. The mirror system shown is the same as in Figs. 7.2 and 7.3. As can be seen, the number of false doublets is far larger than the number of real ones, making

an exhaustive search strategy on the full graph structure extremely costly.




Figure 7.6: *Left:* In the case of unknown doublets and triplets, naïvely, all possible pairs of points have to be considered. *Right:* Our geometric constraints are able to remove most of the false connections. The weakly visible unfolding geometry is shown only for illustration. It is still unknown at this point of our algorithm.

7.5.1 Geometric Search Space Pruning

We therefore derive a number of filtering operations that are intended to reduce the number of potential doublets and the potential number of triplets that are built from them, effectively pruning the search space. A result of our filtering operations is shown in Fig. 7.6 (right).

Our filtering operations exploit the geometric features of a mirror configuration that impose strong constraints on valid distributions of observed points. In addition, they are conservative, i.e., no false negatives are generated. The filter operations come in four flavors, these being filters on

- individual point pairs, i.e. potential doublets, $f_1(D)$,
- pairs of potential doublets, i.e. potential triplets, $f_2(T)$,
- compatibility between two potential doublets, $f_3(D, D')$, and
- compatibility between two potential triplets, $f_4(T, T')$.

Each of these filter types can determine impossible configurations, which, however, does not imply correctness of the filter argument. We next describe the individual filters.

Potential Doublets If $(\mathbf{a}', \mathbf{a}'')$ is a doublet generated by the mirror $M_{\mathbf{a}'\mathbf{a}''}$, and \mathbf{b} is the point of intersection of the virtual mirror $M_{\mathbf{a}'\mathbf{a}''}$ and the side \mathbf{ca}'' , then the intersection area of the inner part of the angle $\angle \mathbf{a}'\mathbf{ca}''$ and




Figure 7.7: For the receiver position c , if (a', a'') is a doublet, then the shaded area must be empty.

the inner part of the circle with center b and radius $a'b$ must not contain any other observed points, see Fig. 7.7. To see this, consider the following: Point b is located on the mirror $M_{a'a''}$, therefore, using the assumption that our room is convex, we can conclude, that the half-closed interval $[a'; b)$ entirely belongs to chamber $C_{M_{K_1}M_{K_2}\dots M_{K_N}}$. Therefore, any point a^* inside the angle $\angle a'ca''$ is associated with one of the following chambers: C_0 , $C_{M_{K_1}M_{K_2}\dots M_{K_i}}$ ($i \geq 1$), or with C_\emptyset . Note, that if $i > N$, then we have a super-sequence of the mirror sequence $(M_{K_1}, M_{K_2}, \dots, M_{K_N})$. Suppose we would like to generate an additional point a^* by reflecting point a' from the virtual mirrors of chamber $C_{M_{K_1}M_{K_2}\dots M_{K_N}}$. In this case, we need to use more than one reflection of a' from the virtual mirrors in $C_{M_{K_1}M_{K_2}\dots M_{K_N}}$. Because the direct observation distance is always shortest, the point a^* must be located outside the inner part of the circle.

Potential Triplets Triplets cannot be in an arbitrary orientation with re-




Figure 7.8: Examples of the invalid (left) and valid (middle, right) triplet orientations.

spect to the receiver, see Fig. 7.8. The left of the figure shows an impossible configuration where a'' is both an image of a' and a^* . Conversely, if a' and

\mathbf{a}^* were images of \mathbf{a}'' , they would be incorrect because $d(\mathbf{c}, \mathbf{a}') < d(\mathbf{c}, \mathbf{a}'')$, and similarly, $d(\mathbf{c}, \mathbf{a}^*) < d(\mathbf{c}, \mathbf{a}'')$ which is violating the fact that apparent distance increases with every reflection. The two sequences on the right are possible, indicating a sequential bouncing (i.e. $\mathbf{a}' = m_{M_{K_1} M_{K_2} \dots M_{K_N}}(\mathbf{a})$, $\mathbf{a}'' = m_{M_{K_1} M_{K_2} \dots M_{K_N} M_{K_{N+1}}}(\mathbf{a})$, $\mathbf{a}^* = m_{M_{K_1} M_{K_2} \dots M_{K_N} M_{K_{N+1}} M_{K_{N+2}}}(\mathbf{a})$) for the middle case, and a splitting ray bundle with a common subsequence up to \mathbf{a}' (i.e. $\mathbf{a}' = m_{M_{K_1} M_{K_2} \dots M_{K_N}}(\mathbf{a})$, $\mathbf{a}'' = m_{M_{K_1} M_{K_2} \dots M_{K_N} M_{K_{N+1}}}(\mathbf{a})$, $\mathbf{a}^* = m_{M_{K_1} M_{K_2} \dots M_{K_N} M_{K_{N+1}}}(\mathbf{a})$) for the right case, respectively.

Doublet Compatibility This filter can only determine that two potential




Figure 7.9: Examples of the incompatible doublets (left) and compatible ones (middle, right).

doublets are incompatible with each other. It cannot determine which of the potential doublets is the violating one. We exploit this condition by setting up a compatibility matrix for all pairs of potential doublets. When building potential triplets, we verify that the constituent doublets are compatible with each other. Otherwise the triplet is incorrect. The left part of Fig. 7.9 shows an impossible case since proper doublets have a unique reflection sequence in the complete triangle spanned by the receiver and the doublet. In this case \mathbf{a}' must be in a preceding chamber to the chamber of \mathbf{b}' and vice versa, which is impossible. For the middle and right cases, the two doublets do not conflict.

Triplet Compatibility is based on the doublet compatibility constraint. By definition, a potential triplet consists of compatible doublets. Unfortunately, the doublet compatibility relation is not transitive since it can only report incorrect pairs of potential doublets. Therefore, two potential triplets even though consistent themselves can contain a combination of inconsistent potential doublets. Fig. 7.10 shows two triplets built from potential doublets D_1 and D_2 , and D_1 and D_3 , respectively. For the resulting quadruplet to be consistent, D_2 and D_3 have to be compatible via doublet compatibility as well. The same considerations hold for larger assemblies of triplets. When




Figure 7.10: For two triplets $T = (D_1, D_2)$ and $T' = (D_1, D_3)$ with a common leg D_1 to be compatible it is necessary, that doublets D_2 and D_3 are compatible themselves.

a new triplet is joined in, all doublets have to be checked for compatibility, again due to the non-transitivity of the relation.

As can be seen from Fig. 7.6 (right), our filtering strategies are very effective in pruning false point pairs. However, we cannot remove all incorrect pairs. Therefore, the doublet graph contains false doublets. An exhaustive search strategy can be employed on the pruned doublet graph and will find a solution if it exists (in the sense of Sect. 7.4.2). However, due to its exponential nature, the search is not very efficient and large problem instances may be practically unfeasible.

7.6 Randomized Search Algorithm

We therefore adapt our search strategy for connected components that is used to build candidate configurations. In the ideal case considered in Section 7.4.1, the doublet graph contained only valid doublets. In the case of the doublet graph containing invalid doublets, we maintain the search for a connected component. However, we perform this search in a randomized manner by employing a forward search strategy that is using importance sampling to decide on likely transitions for the triplet joining procedure. In every step, we perform a validation of the current configuration via the method outlined in Section 7.4.3. The importance scores are based on doublet and triplet statistics (correct doublets and triplets occur more often) and compatibility checks as outlined in Section 7.5.1. Each connected component discovered such is checked in all its possible flipped configurations. If unsuccessful, the search ends after an upper number of triplet additions (provided by the user) has been reached or if no triplet can be added in a consistent manner. In this case, the procedure is restarted until an upper number of trials has been reached.

7.6.1 Extensions

So far, we have concentrated on an idealized simulation setting. In reality, a number of issues might occur. The most important aspect is measurement noise. It not only influences the measured positions of the point data that we use for our reconstruction algorithm, it also complicates the process of establishing equivalence classes for doublets. In practice, we use a user supplied ϵ on the length of potential doublets to account for the expected variation. This in turn puts constraints on the point's position with respect to the mirror planes. Our current algorithm is not designed to handle the issue of incorrect class assignments or mixed equivalence classes explicitly. However, even in this case there is a chance that the correct mirror structure is recovered due to the randomized nature of the algorithm.

Another constraint that we introduced in Section 7.3.1 is that the receiver has to be positioned in the base chamber. We were using this constraint to terminate the repeated backward mirror operation necessary for the validation of candidate configurations. In practice, the camera can be positioned outside the base chamber. Our algorithm will still recover the geometry if sufficient data is available. However, the position of the camera with respect to the mirror geometry can only be recovered up to a discrete number of positions. For an intuitive understanding, please refer to Fig. 7.3 and consider that the backward mirroring operation is stopped early, yielding one possible position and orientation with respect to the reconstructed mirror geometry for each possible stopping position. Depending on the location of the reconstruction in the virtual world, it might be necessary to perform forward mirroring along the line-of-sight as well.

Finally, gaps in the mirrors only lead to missing data, i.e. some of the data points that would be observed otherwise are missing. Our algorithm therefore can be applied unchanged if, e.g. mirrors are not meeting at a corner or if the field of view of the receiver is restricted.

7.7 Experimental Results

7.7.1 Simulation Results

We performed extensive simulations with our algorithm to investigate the stability of the results and to study recoverability of the geometry with respect to the number of sides in a polygon and the number of observed inter-reflections. We randomly generated 2000 different convex mirror systems with random object point positions for each n-gon, where $n \in [3..8]$. We simulated both a full surround receiver and a field-of-view that was restricted to 90° . In the case of a surround receiver, the direct observation will always be observed. However, this is not the case in the limited field-of-view




Figure 7.11: Simulation results for 2000 randomly generated n -gons. The plots show the number of reconstructable systems versus the number of reflections considered for the reconstruction task. *Left*: full surround receiver, *Right*: field-of-view limited to 90° . The solid lines indicate the exhaustive search algorithm, the dashed lines the randomized search.

example. We did not pay attention to include or exclude any particular reflection level such as the direct observation in our simulated mirror systems. To avoid a bias in the statistics due to extreme configurations, we limited the systems to a ratio of 3 : 1 between the largest mirror and the smallest one.




Figure 7.12: Simulation results for 2000 randomly generated n -gons. The plots show reasons for reconstruction failure. *Left*: failure due to missing doublets, *Right*: failure due to missing triplets. The solid lines indicate the 360° receiver, the dashed lines the 90° receiver.

For analysis, we performed an exhaustive search, as described in Section 7.4.1 on the clean doublet graph structure, i.e. only true doublets and triplets participate. This way, we established an upper bound on performance for any triplet-based reconstruction algorithm. The results are shown

in Fig. 7.11 as solid lines. The plot shows the number of reconstructable room geometries. Then we ran the randomized search algorithm described in Section 7.6 that utilized the filtered graph structure, Fig. 7.6 (right). We again tested whether the geometry could be recovered. The results of this test are shown in Fig. 7.11 as dashed lines.

First, analyzing the exhaustive search results to have a comparison baseline, we find that the case of a surround receiver is favorable, especially in the case of a low number of usable reflections. Another interesting aspect is that an increasing number of reflections leads to a convergence in the recoverable geometries. This suggests an intrinsic bound for the information in our data: higher levels of reflection are so fractured that no additional useful information can be observed. If the number of mirror walls in the geometry increases, our chances of success decrease rapidly. In the vast majority of cases, the reason for a failure to reconstruct the geometry is that doublets are missing from the observation, see Fig. 7.12 (left). This indicates that some mirror planes are never observed via a direct reflection in many cases. Once sufficiently many doublets are observed, the failure to identify sufficiently many triplets is not a serious problem, Fig. 7.12 (right).

The results for our randomized search strategy, which constitutes our practical reconstruction algorithm, show that for a low number n of mirror walls we can perform a reasonable job. Again, the results deteriorate quickly with larger n . We ran the algorithm with a fixed user threshold of 50 connected component recovery trials. The decrease of the performance curves with larger number of reflections shows that the complexity of the graph structure increases and that the randomized algorithm has less success in discovering one of the correct configurations. We would like to mention that if a solution is found, it is exact since we work in a noise-less setting.

7.7.2 Real World Example

To test our algorithm in a real setting, we performed a calibration experiment. We set up a system of six planar mirrors containing a checkerboard and took a photograph with an intrinsically calibrated camera. The setup together with a sketch of the ground plan according to which the system was set up manually is shown in Fig. 7.13. It contains four walls that are approximately orthogonal to both the ground plane and the ceiling which also consist of mirrors, realizing a $2\frac{1}{2}$ D setup. We triangulated the midpoint of the checkerboard in various apparent locations. This data matches the requirements of our algorithm. We then reconstructed the polygonal outline of the four walls. The result is shown in Fig. 7.13 as a super-imposed outline on the sketch. A visual impression of the accuracy of our reconstruction can be gained from Fig. 7.14 (right) where we rendered the multiply reflected mirror planes as semi-transparent polygons. The results show that we can reconstruct a mirror geometry even from real-world samples. The




Figure 7.13: *Left*: Photograph of our system with mirrors indicated. The top mirror is removed to show the inside. *Right*: A sketch of the ground plan according to which we built our system super-imposed with the reconstruction result (green).

only change to our algorithm is that we use 3D uncertainties available from the triangulation procedure to evaluate our potential configurations. The remaining mismatches can be attributed to the manual setup and alignment of the mirrors as well as to imperfect orthogonality between the ground and ceiling planes and the mirror walls.




Figure 7.14: *Left*: View inside the mirror system that was used for reconstruction. *Right*: The reconstructed mirrors are super-imposed on the image as attenuating layers. Further reflection levels appear brighter.

7.8 Conclusions

We have shown that it is possible to reconstruct the geometry of a convex room of mirrors from the measurement of a single scene point and we

have identified sufficient conditions for doing so. For this it is necessary to measure its distance to the receiver via many different inter-reflection light paths. Our technique relates to time-of-flight measurements and could possibly be used beneficially in areas such as active SONAR, RADAR and LIDAR where “ghosts” are a frequent problem. Our work shows that these ghosts carry valuable information about the scene.

Conclusions and Future work

Conclusions

In this thesis we have introduced a general framework for dealing with systems of planar mirrors. Its goal is to simplify the interpretation of light ray trajectories in planar mirror systems. The framework's main tool is the space partitioning concept. Using this concept, we have shown that planar mirror systems are continuous with respect to change of the observation point.

On the practical side, we have illustrated the value of the proposed framework in solving important but challenging problems arising in computer vision and time-of-flight imaging that involve multiple interactions with planar mirrors.

For example, we have shown that generalized kaleidoscopic imaging systems can be used to obtain dense spherical multi-view data for extended (in contrast to planar) objects that are calibrated both geometrically and photometrically. We have also shown that the output of these techniques is directly usable in standard multi-view reconstruction algorithms as well as for reflectance acquisition. All these properties allow for inexpensive solutions when compared to traditional multi-view acquisition systems and allows for an improved use of the available sensor area and/or projector resolution. Additional advantages include perfect synchronization and common radiometric and colorimetric properties for all views and/or illumination directions.

The theoretical framework has also enabled progress in the solution of the problem of reconstructing a planar mirror system from multiple observations of a single point. In particular, we have enlarged the family of reconstructable planar mirror configurations in $2\frac{1}{2}D$. In contrast to previous results in this area, under certain constraints, our algorithm is capable to recover the mirror structure even when the point is not observable directly and/or through a single reflection.

Future Work

The tools and methods developed in this thesis open new possibilities for the use of virtual mirror worlds generated by planar kaleidoscopes in computer vision. However, to achieve a true freedom in the use of planar mirror systems, we still have to find solutions to several questions.

Sampling of kaleidoscopes. The principal problems associated with depth-of-field, field-of-view and multi-resolution are related to the specific design of the kaleidoscope. Other, better kaleidoscope configurations may minimize these drawbacks. In the future, better configurations of planar mirror systems could be established by sampling the space of planar mirror systems. We have already obtained partial result towards this goal in Part II: the observer-continuity of a fixed system allows for a sampling of the best observation point. For a full optimization, however, continuity has to be proven for a change in the mirror configuration as well.

Kaleidoscopic visual hull and labeling. Future work includes development of rigorous bounds for the pixel labeling error at occlusion boundaries. Since this is impossible in the general case, suitable object classes that cover all scenes of interest while still enabling an estimation of a tight error bound have to be found. Future algorithmic developments include the incorporation of inherently multi-resolution data in multi-view reconstruction algorithms as well as investigating techniques to differentiate between the limited number of views that are potentially responsible for unreliable pixels. Another direction for future work is the improvement of the space carving implementation. We used a simple but relatively slow and imprecise voxel-based algorithm. In the future, the polyhedral-based algorithm [Franco03] could be adapted to improve the current solution.

Acquisition of dynamic scenes. The acquisition of dynamic scenes, i.e. both geometry and reflectance, is one of the promising areas of application of kaleidoscopes. This is due to the perfect synchronization between different views. However, the problem of automatically labeling a dynamic scene has to be solved.

Structure reconstruction of planar mirror systems. In the future we would like to investigate problem instances where the sufficient mirror reconstruction conditions derived in Chapter 7 do not hold. The goal would be to identify the class of reconstructable mirror systems, i.e., the necessary and sufficient conditions for reconstructability. In particular, our definitions of doublets and triplets rely on a single bounce separation of two observed points. There may be problem instances where mirrors can only be observed through second or higher-order bounces. A different way to overcome the insufficient amount of connections via doublets and triplets, as well as the burden of reconstruction complexity, is to use multiple distinct observation points in the system. Another task for the future is the generalization of

the described $2\frac{1}{2}D$ approach to truly $3D$ cases. In this situation the concept of doublets and triplets remains the same, while assembling them together into quadruplets, quintets etc. will not be performed in the $2D$ plane, but in $3D$ space.

Curved mirrors and refractive systems. Curved mirror systems or systems with refractive elements are even less understood and harder to predict than planar mirror systems. This thesis is dedicated to planar mirror systems, but the theory developed here may be partially used in the curved mirror case or in the case when there is a refraction of rays (or even combinations thereof). In particular, the space partitioning concept may be adopted to the curved or refractive cases in future work. Note, that in these cases, chambers, except maybe the base chamber, will have a non-Euclidean structure.

Bibliography

- [Aggarwal01] M. AGGARWAL AND N. AHUJA. Split Aperture Imaging for High Dynamic Range. In *Proceedings of IEEE International Conference on Computer Vision*, volume 2, pages 10–17 vol.2, 2001. (Cited on page 9)
- [Agrawal13] AMIT AGRAWAL AND SRIKUMAR RAMALINGAM. Single Image Calibration of Multi-Axial Imaging Systems. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 1399–1406, 2013. (Cited on page 15)
- [Ajdler06] T. AJDLER, L. SBAIZ, AND M. VETTERLI. The Plenacoustic Function and its Sampling. *IEEE Trans. on Signal Processing*, 54(10):3790–3804, 2006. (Not cited.)
- [Allen79] J ALLEN AND D BERKLEY. Image Method for Efficiently Simulating Small-Room Acoustics. *The Journal of the Acoustical Society of America*, 65(4):943–950, 1979. (Cited on page 14)
- [Antonacci12] F. ANTONACCI, J. FILOS, M. R. P. THOMAS, E.A.P. HABETS, A. SARTI, P. A. NAYLOR, AND S. TUBARO. Inference of Room Geometry From Acoustic Impulse Responses. *IEEE Trans. on Audio, Speech and Language Processing*, 20(10):2683–2695, 2012. (Cited on pages 14 and 102)
- [Baker98] SIMON BAKER AND SHREE K. NAYAR. A Theory of Catadioptric Image Formation. In *Proceedings of IEEE International Conference on Computer Vision*, pages 35–42, 1998. (Cited on page 11)
- [Baker99] SIMON BAKER AND SHREE K. NAYAR. A Theory of Single-Viewpoint Catadioptric Image Formation. *In-*

- ternational Journal of Computer Vision*, 35(2):175–196, 1999. (Cited on page 11)
- [Bangay04] SHAUN BANGAY AND JUDITH D. RADLOFF. Kaleidoscope Configurations for Reflectance Measurement. In *Proceedings of AFRIGRAPH*, pages 161–170, 2004. (Cited on page 12)
- [Ben-Ezra08] M. BEN-EZRA, JIAPING WANG, B. WILBURN, XIAOYANG LI, AND LE MA. An LED-Only BRDF Measurement Device. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, pages 1–8, June 2008. (Cited on page 65)
- [Berg08] C. BERG AND KØBENHAVNS UNIVERSITET. MATEMATISK AFDELING. *Complex Analysis: Christian Berg*. Matematisk Afdeling, Københavns Universitet, 2008. (Cited on page 43)
- [Bimber08] OLIVER BIMBER, DAISUKE IWAI, GORDON WETZSTEIN, AND ANSELM GRUNDHÖFER. The Visual Computing of Projector-Camera Systems. *Computer Graphics Forum*, 27(8):2219–2245, 2008. (Cited on page 11)
- [Borish84] JEFFREY BORISH. Extension of the Image Model to Arbitrary Polyhedra. *Journal of the Acoustic Society of America*, 75(6):1827–1836, 1984. (Cited on page 14)
- [Bouguet05] JEAN-YVES BOUGUET. Camera Calibration Toolbox for Matlab, 2005. (Cited on pages 13 and 25)
- [Cabral04] E.L.L. CABRAL, J.C. DE SOUZA, AND M.C. HUNOLD. Omnidirectional stereo vision with a hyperbolic double lobed mirror. In *Pattern Recognition, 2004. ICPR 2004. Proceedings of the 17th International Conference on*, volume 1, pages 1–9 Vol.1, Aug 2004. (Cited on page 15)
- [Coxeter67] HAROLD S. M. COXETER AND SAMUEL L. GREITZER. *Geometry Revisited*. The Mathematical Association of America, Washington, 1967. (Cited on page 17)
- [Dana01] KRISTIN DANA. BRDF/BTF Measurement Device. In *Proceedings of IEEE International Conference on Computer Vision*, pages 460–466, 2001. (Cited on page 15)
- [Dana04] KRISTIN J. DANA AND JING WANG. Device for Convenient Measurement of Spatially Varying Bidirectional

- Reflectance. *Journal of the Optical Society of America A*, 21:1–12, January 2004. (Cited on page 15)
- [Debevec00] PAUL DEBEVEC, TIM HAWKINS, CHRIS TCHOU, HAARM-PIETER DUIKER, WESTLEY SAROKIN, AND MARK SAGAR. Acquiring the Reflectance Field of a Human Face. In *Proceedings of SIGGRAPH*, pages 145–156, 2000. (Cited on page 65)
- [Ding09a] YUANYUAN DING, JING XIAO, KAR-HAN TAN, AND JINGYI YU. Catadioptric Projectors. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 2528–2535, 2009. (Cited on page 15)
- [Ding09b] YUANYUAN DING, JINGYI YU, AND PETER STURM. Multiperspective Stereo Matching and Volumetric Reconstruction. In *Proceedings of IEEE International Conference on Computer Vision*, pages 1827–1834, 2009. (Cited on page 15)
- [Dokmanic11] IVAN DOKMANIC, YUE M. LU, AND MARTIN VETTERLI. Can one hear the shape of a room: The 2-D polygonal case. In *International Conference on Acoustics, Speech, and Signal Processing*, pages 321–324, 2011. (Cited on page 14)
- [Dokmanic13] I. DOKMANIC, R. PARHIZKAR, A. WALTHER, Y. M. LU, AND M. VETTERLI. Acoustic Echoes Reveal Room Shape. *Proceedings of the National Academy of Sciences*, (30):12186–12191, 2013. (Cited on page 14)
- [Fasano03] ANDREA FASANO, MARCO CALLIERI, PAOLO CIGNONI, AND ROBERTO SCOPIGNO. Exploiting Mirrors for Laser Stripe 3D Scanning. In *Proc. of 4th International Conference on 3D Digital Imaging and Modeling (3DIM 2003)*, Banff (Canada), 2003. (Cited on page 11)
- [Forbes06] KEITH FORBES, FRED NICOLLS, GERHARD DE JAGER, AND ANTHON VOIGT. Shape-from-Silhouette with two Mirrors and an Uncalibrated Camera. In *Proceedings of European Conference on Computer Vision*, pages 165–178, 2006. (Cited on page 10)
- [Franco03] JEAN-SÉBASTIEN FRANCO AND EDMOND BOYER. Exact Polyhedral Visual Hulls. In *Proceedings of British Machine Vision Conference*, volume 1, pages 329–338, 2003. (Cited on page 122)

- [Fuchs12] MARTIN FUCHS, MARKUS KÄCHELE, AND SZYMON RUSINKIEWICZ. Design and Fabrication of Faceted Mirror Arrays for Light Field Capture. In *Proceedings of Vision, Modeling, and Visualization (VMV)*, pages 1–8, 2012. (Cited on page 8)
- [Fujii05] K. FUJII, M.D. GROSSBERG, AND S.K. NAYAR. A Projector-Camera System with Real-Time Photometric Adaptation for Dynamic Environments. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, volume 2, pages 1180–1187 vol. 2, 2005. (Cited on page 11)
- [Gao09] LIANG GAO, ROBERT T. KESTER, AND TOMASZ S. TKACZYK. Compact Image Slicing Spectrometer (ISS) for Hyperspectral Fluorescence Microscopy. *Optics Express*, 17(15):12293–12308, 2009. (Cited on page 9)
- [Garg06] GAURAV GARG GARG, EINO-VILLE TALVALA, MARC LEVOY, AND HENDRIK P. A. LENSCH. Symmetric Photography: Exploiting Data-Sparseness in Reflectance Fields. In *Proceedings of Eurographics Symposium on Rendering*, pages 251–262, June 2006. (Cited on page 11)
- [Ghosh07] A. GHOSH, S. ACHUTHA, W. HEIDRICH, AND M. O’TOOLE. BRDF Acquisition with Basis Illumination. In *Proceedings of IEEE International Conference on Computer Vision*, pages 183–197, 2007. (Cited on pages 11 and 15)
- [Ghosh10] A. GHOSH, W. HEIDRICH, S. ACHUTHA, AND M. O’TOOLE. A Basis Illumination Approach to BRDF Measurement. *International Journal of Computer Vision*, 90(2):183–197, 2010. (Cited on page 11)
- [Gluckman98a] JOSHUA GLUCKMAN AND SHREE K. NAYAR. A Real-Time Catadioptric Stereo System Using Planar Mirrors. In *Proc. of Image Understanding Workshop*, 1998. (Cited on pages 7 and 9)
- [Gluckman98b] JOSHUA GLUCKMAN AND SHREE K. NAYAR. Ego-Motion and Omnidirectional Cameras . In *Proceedings of IEEE International Conference on Computer Vision*, pages 999–1005, 1998. (Cited on page 15)

- [Gluckman98c] JOSHUA GLUCKMAN AND SHREE K. NAYAR. Real-Time Omnidirectional and Panoramic Stereo. In *Proc. of Image Understanding Workshop*, 1998. (Cited on page 15)
- [Gluckman00] JOSHUA GLUCKMAN AND SHREE K. NAYAR. Rectified Catadioptric Stereo Sensors. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 380–387, 2000. (Cited on pages 7, 8, and 9)
- [Gluckman02] J. GLUCKMAN AND S.K. NAYAR. Rectified Catadioptric Stereo Sensors. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(2):224–236, 2002. (Cited on pages 7, 8, and 9)
- [Gorman10] A. GORMAN, D. W. FLETCHER-HOLMES, AND A. R. HARVEY. Generalization of the Lyot Filter and its Application to Snapshot Spectral Imaging. *Optics Express*, 18(6):5602–5608, 2010. (Cited on page 9)
- [Gortler96] S. GORTLER, R. GRZESZCZUK, R. SZELINSKI, AND M. COHEN. The Lumigraph. In *Proceedings of SIGGRAPH*, pages 43–54, 1996. (Cited on page 65)
- [Grzegorzec13] Marcin Grzegorzec, Christian Theobalt, Reinhard Koch, and Andreas Kolb, editors. *Time-of-Flight and Depth Imaging. Sensors, Algorithms, and Applications*, volume 8200 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, 2013. (Cited on page 14)
- [Han03] JEFFERSON Y. HAN AND KEN PERLIN. Measuring Bidirectional Texture Reflectance with a Kaleidoscope. In *Proceedings of SIGGRAPH*, pages 741–748, 2003. (Cited on pages 12 and 89)
- [Harvey05] ANDREW R. HARVEY, DAVID W. FLETCHER-HOLMES, AND ALISTAIR GORMAN. Spectral Imaging in a Snapshot. In *Proc. SPIE 5694*, pages 1–10, 2005. (Cited on page 9)
- [Heide13] FELIX HEIDE, MATTHIAS HULLIN, JAMES GREGSON, AND WOLFGANG HEIDRICH. Low-Budget Transient Imaging using Photonic Mixer Devices. *ACM Transactions on Graphics*, 32(4):45:1–45:10, 2013. (Cited on page 14)
- [Holroyd10] MICHAEL HOLROYD, JASON LAWRENCE, AND TODD ZICKLER. A Coaxial Optical Scanner for Synchronous Acquisition of 3D Geometry and Surface Reflectance. *ACM*

- Transactions on Graphics*, 29(4):article no. 99, 2010. (Cited on pages 11 and 65)
- [Hu09] BO HU. It's All Done with Mirrors: Calibration-and-Correspondence-Free 3D Reconstruction. In *Proceedings of the Conference on Computer and Robot Vision*, pages 148–154, 2009. (Cited on pages 8 and 65)
- [Huang06] PO-HAO HUANG AND SHANG-HON LAI. Contour-Based Structure from Reflection. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 165–178, 2006. (Cited on page 10)
- [Ihrke08] IVO IHRKE, TIMO STICH, HEIKO GOTTSCHLICH, MARCUS MAGNOR, AND HANS-PETER SEIDEL. Fast Incident Light Field Acquisition and Rendering. *Journal of WSCG*, 16(1-3):25–32, 2008. (Cited on pages 8, 13, and 65)
- [Ihrke10] IVO IHRKE, KIRIAKOS N. KUTULAKOS, HENDRIK P. A. LENSCH, MARCUS MAGNOR, AND WOLFGANG HEIDRICH. Transparent and Specular Object Reconstruction. *Computer Graphics Forum*, 29(8):2400–2426, 2010. (Cited on page 15)
- [Ihrke12a] IVO IHRKE, ILYA RESHETOUSKI, ALKHAZUR MANAKOV, AND HANS-PETER SEIDEL. Three-Dimensional Kaleidoscopic Imaging. In *Computational Optical Sensing and Imaging*, page CTu4B.8. Optical Society of America, 2012. (Cited on pages viii and xi)
- [Ihrke12b] IVO IHRKE, ILYA RESHETOUSKI, ALKHAZUR MANAKOV, ART TEVS, MICHAEL WAND, AND HANS-PETER SEIDEL. A Kaleidoscopic Approach to Surround Geometry and Reflectance Acquisition. In *Proceedings of IEEE International Workshop on Computational Cameras and Displays*, pages 29–36, 2012. (Cited on pages viii, xi, 11, 12, and 92)
- [Jang05] GIJEONG JANG, SUNGHO KIM, AND INSO KWEON. Single Camera Catadioptric Stereo System. In *Workshop on Omnidirectional Vision, Camera Networks and Nonclassical Cameras (OMNIVIS2005)*, pages 1–7, 2005. (Cited on page 15)
- [Kadambi13] ACHUTA KADAMBI, REFAEL WHYTE, AYUSH BHANDARI, LEE STREETER, CHRISTOPHER BARSII, ADRIAN

- DORRINGTON, AND RAMESH RASKAR. Coded Time of Flight Cameras: Sparse Deconvolution to Address Multipath Interference and Recover Time Profiles. *ACM Trans. Graph.*, 32(6):167:1–167:10, November 2013. (Cited on page 14)
- [Kirmani09] A. KIRMANI, T. HUTCHISON, J. DAVIS, AND R. RASKAR. Looking around the Corner using Transient Imaging. In *Proceedings of IEEE International Conference on Computer Vision*, pages 159–166, 2009. (Cited on page 14)
- [Klehm12] OLIVER KLEHM, ILYA RESHETOUSKI, ELMAR EISEMANN, IVO IHRKE, AND HANS-PETER SEIDEL. Interactive Geometry-Aware Segmentation for the Decomposition of Kaleidoscopic Images. In *17th International Workshop on Vision, Modeling and Visualization*, pages 9–14, Magdeburg, Germany, November 2012. Eurographics Association. (Cited on pages viii, xi, and 84)
- [Kuthirummal06] SUJIT KUTHIRUMMAL AND SHREE K. NAYAR. Multiview Radial Catadioptric Imaging for Scene Capture. *ACM Transactions on Graphics*, 25(3), 2006. (Cited on pages 15 and 65)
- [Lanman06a] DOUGLAS LANMAN, DANIEL CRISPELL, MEGAN WACHS, AND GABRIEL TAUBIN. Spherical Catadioptric Arrays: Construction, Multi-View Geometry, and Calibration. In *Proceedings of the International Symposium on 3D Data Processing, Visualization and Transmission (3DPVT)*, pages 81–88, 2006. (Cited on page 15)
- [Lanman06b] DOUGLAS LANMAN, MEGAN WACHS, GABRIEL TAUBIN, AND FERNANDO CUKIERMAN. Reconstructing a 3D Line from a Single Catadioptric Image. In *Proceedings of the International Symposium on 3D Data Processing, Visualization and Transmission (3DPVT)*, pages 1–8, 2006. (Cited on page 15)
- [Lanman07] DOUGLAS LANMAN, DANIEL CRISPELL, AND GABRIEL TAUBIN. Surround Structured Lighting for Full Object Scanning. In *International Conference on 3D Digital Imaging and Modeling (3DIM)*, pages 107–116, 2007. (Cited on pages 10 and 12)
- [Lanman09] DOUGLAS LANMAN, DANIEL CRISPELL, AND GABRIEL TAUBIN. Surround Structured Lighting: 3-D Scanning

- with Orthographic Illumination. *Computer Vision and Image Understanding*, 113(11):1107–1117, 2009. (Cited on pages 10 and 12)
- [Laurentini94] ALDO LAURENTINI. The Visual Hull Concept for Silhouette-Based Image Understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 150–162, 1994. (Cited on page 23)
- [Lensch03] HENDRIK P. A. LENSCH, JAN KAUTZ, MICHAEL GOESELE, WOLFGANG HEIDRICH, AND HANS-PETER SEIDEL. Image-based Reconstruction of Spatial Appearance and Geometric Detail. *ACM Transactions on Graphics*, 22(2):234–257, April 2003. (Cited on pages 15 and 65)
- [Levoy96] M. LEVOY AND P. HANRAHAN. Light Field Rendering. In *Proceedings of SIGGRAPH*, pages 31–42, 1996. (Cited on page 65)
- [Levoy04] MARC LEVOY, BILLY CHEN, VAIBHAV VAISH, MARK HOROWITZ, IAN MCDOWALL, AND MARK BOLAS. Synthetic Aperture Confocal Imaging. *ACM Transactions on Graphics*, 23:825–834, August 2004. (Cited on pages 8, 11, and 65)
- [Lin06] SHIH-SCHÖN LIN AND RUZENA BAJCSY. Single-viewpoint, Catadioptric Cone Mirror Omnidirectional Imaging Theory and Analysis. *Journal of the Optical Society of America A*, 23:2997–3015, May 2006. (Cited on page 65)
- [Lowe04] DAVID G. LOWE. Distinctive Image Features from Scale-Invariant Keypoints. *Int. J. Comput. Vision*, 60(2):91–110, November 2004. (Cited on page 85)
- [Manakov11] ALKHAZUR MANAKOV, HANS-PETER SEIDEL, AND IVO IHRKE. A Mathematical Model and Calibration Procedure for Galvanometric Laser Scanning Systems. In *Proceedings of Vision, Modeling, and Visualization (VMV)*, pages 207–214, 2011. (Cited on page 25)
- [Mantiuk06] RAFAŁ MANTIUK. PFSCalibration, 2006. (Cited on page 27)
- [Mariottini09] GIAN LUCA MARIOTTINI, STEFANO SCHEGGI, FABIO MORBIDI, AND DOMENICO PRATTICHIZZO. Planar Catadioptric Stereo: Single and Multi-View Geometry for Cal-

- ibration and Localization. In *Proceedings of IEEE International Conference on Robotics and Automation*, pages 2711–2716, 2009. (Cited on page 13)
- [Masselus03] VINCENT MASSELUS, PIETER PEERS, PHILIP DUTRÉ, AND YVES D. WILLEMS. Relighting with 4D incident light fields. *ACM Transactions on Graphics*, 22(3):613–620, July 2003. (Cited on page 65)
- [McGuire07] MORGAN MCGUIRE, WOJCIECH MATUSIK, HANSPETER PFISTER, BILLY CHEN, JOHN F. HUGHES, AND SHREE K. NAYAR. Optical Splitting Trees for High-Precision Monocular Imaging. *IEEE Comput. Graph. & Appl.*, 27(2):32–42, 2007. (Cited on pages 8 and 9)
- [Mitsumoto92] HIROSHI MITSUMOTO, SHINICHI TAMURA, KOZO OKAZAKI, NAOKI KAJIMI, AND YUTAKA FUKUI. 3-D Reconstruction Using Mirror Images Based on a Plane Symmetry Recovering Method. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14:941–946, September 1992. (Cited on page 8)
- [Mukaigawa07a] YASUHIRO MUKAIGAWA, KOHEI SUMINO, AND YASUSHI YAGI. High-Speed Measurement of BDRF Using an Ellipsoidal Mirror and a Projector. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, 2007. (Cited on page 15)
- [Mukaigawa07b] YASUHIRO MUKAIGAWA, KOHEI SUMINO, AND YASUSHI YAGI. Multiplexed Illumination for Measuring BDRF Using an Ellipsoidal Mirror and a Projector. In *Proceedings of Asian Conference on Computer Vision*, pages 246–257, 2007. (Cited on page 15)
- [Mukaigawa11] YASUHIRO MUKAIGAWA, SEIICHI TAGAWA, JAEWON KIM, RAMESH RASKAR, YASUYUKI MATSUSHITA, AND YASUSHI YAGI. Hemispherical Confocal Imaging Using Turtleback Reflector. In *Proceedings of the 10th Asian conference on Computer vision - Volume Part I, ACCV'10*, pages 336–349, Berlin, Heidelberg, 2011. Springer-Verlag. (Cited on pages 8 and 11)
- [Murray95] D. W. MURRAY. Recovering Range using Virtual Multi-Camera Stereo. *Computer Vision and Image Understanding*, 61(2):285–291, March 1995. (Cited on pages 8 and 65)

- [Nayar88] SHREE K. NAYAR. Sphero: Determining Depth using Two Specular Spheres and a Single Camera. In *SPIE Vol. 1005 Optics, Illumination, and Image Sensing for Machine Vision III*, pages 245–254, 1988. (Cited on page 15)
- [Nayar97] S.K. NAYAR. Catadioptric Omnidirectional Camera. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 482–488, June 1997. (Cited on page 15)
- [Ng05] REN NG, MARC LEVOY, MATHIEU BRÉDIF, GENE DUVAL, MARK HOROWITZ, AND PAT HANRAHAN. Light Field Photography with a Hand-Held Plenoptic Camera. Technical Report Computer Science CSTR 2005-02, Stanford University, 2005. (Cited on page 65)
- [Ramalingam11] SRIKUMAR RAMALINGAM, SOFIEN BOUAZIZ, PETER STURM, AND PHILIP H.S. TORR. The Light-Path Less Traveled. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 3145 – 3152, 2011. (Cited on page 13)
- [Reshetouski11] ILYA RESHETOUSKI, ALKHAZUR MANAKOV, HANS-PETER SEIDEL, AND IVO IHRKE. Three-Dimensional Kaleidoscopic Imaging. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 353–360, 2011. (Cited on pages viii, xi, and 12)
- [Reshetouski13a] ILYA RESHETOUSKI AND IVO IHRKE. Mirrors in Computer Graphics, Computer Vision and Time-of-Flight Imaging. In Marcin Grzegorzec, Christian Theobalt, Reinhard Koch, and Andreas Kolb, editors, *Time-of-Flight and Depth Imaging. Sensors, Algorithms, and Applications*, volume 8200 of *Lecture Notes in Computer Science*, pages 77–104. Springer Berlin Heidelberg, 2013. (Cited on pages viii, xi, and 15)
- [Reshetouski13b] ILYA RESHETOUSKI, ALKHAZUR MANAKOV, AYUSH BANDHARI, RAMESH RASKAR, HANS-PETER SEIDEL, AND IVO IHRKE. Discovering the Structure of a Planar Mirror System from Multiple Observations of a Single Point. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 89–96, 2013. (Cited on pages viii and xi)

- [Ribeiro12] F. RIBEIRO, D. FLORENCIO, D. BA, AND CHA ZHANG. Geometrically Constrained Room Modeling With Compact Microphone Arrays. *IEEE Transactions on Audio, Speech, and Language Processing*, 20(5):1449–1460, 2012. (Cited on pages 14 and 102)
- [Robertson03] MARK A. ROBERTSON, SEAN BORMAN, AND ROBERT L. STEVENSON. Estimation-Theoretic Approach to Dynamic Range Enhancement using Multiple Exposures. *Journal of Electronic Imaging*, 12(2):219–228, April 2003. (Cited on page 27)
- [Scheuing06] J. SCHEUING AND B. YANG. Disambiguation of TDOA Estimates in Multi-Path Multi-Source Environments (DATEMM). *IEEE ICASSP*, 4:837–840, May 2006. (Cited on page 14)
- [Schwartz11] CHRISTOPHER SCHWARTZ, MICHAEL WEINMANN, ROLAND RUITERS, AND REINHARD KLEIN. Integrated High-Quality Acquisition of Geometry and Appearance for Cultural Heritage. In *Proceedings of International Symposium on Virtual Reality, Archeology and Cultural Heritage (VAST)*, pages 25–32, 2011. (Cited on pages 65 and 95)
- [Schwartz13] CHRISTOPHER SCHWARTZ, RALF SARLETTE, MICHAEL WEINMANN, AND REINHARD KLEIN. DOME II: A Parallelized BTF Acquisition System. In *Eurographics Workshop on Material Appearance Modeling: Issues and Acquisition*, pages 25–31. Eurographics Association, June 2013. (Cited on page 65)
- [Sen05] PRADEEP SEN, BILLY CHEN, GAURAV GARG, STEPHEN R. MARSCHNER, MARK HOROWITZ, MARC LEVOY, AND HENDRIK P. A. LENSCH. Dual photography. *ACM Transactions on Graphics*, 24:745–755, July 2005. (Cited on pages 8 and 65)
- [Smith08] WARREN J. SMITH. *Modern Optical Engineering, 4th edition*. McGraw Hill Professional, 2008. (Cited on pages 9 and 18)
- [Sossinsky12] ALEKSEI BRONISLAVOVICH SOSSINSKY. *Geometries*, volume 64 of *Student mathematical library*. American Mathematical Society, 2012. (Cited on page 22)

- [Taguchi10a] YUICHI TAGUCHI, AMIT AGRAWAL, SRIKUMAR RAMALINGAM, AND ASHOK VEERARAGHAVAN. Axial Light Field for Curved Mirrors: Reflect Your Perspective, Widen Your View. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, 2010. (Cited on page 65)
- [Taguchi10b] YUICHI TAGUCHI, AMIT AGRAWAL, ASHOK VEERARAGHAVAN, SRIKUMAR RAMALINGAM, AND RAMESH RASKAR. Axial-Cones: Modeling Spherical Catadioptric Cameras for Wide-Angle Light Field Rendering. *ACM Trans. Graph.*, 29:172:1–172:8, 2010. (Cited on page 15)
- [Tan04] KAR-HAN TAN, HONG HUA, AND NARENDRA AHUJA. Multiview Panoramic Cameras Using Mirror Pyramids. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(7):941–946, 2004. (Cited on page 9)
- [Tervo12] S. TERVO AND T. TOSSAVAINEN. 3D Room Geometry Estimation from Measured Impulse Responses. In *Proc. ICASSP*, pages 513–516, 2012. (Cited on pages 14 and 102)
- [Unger03] JONAS UNGER, A. WENGER, TIM HAWKINS, A. GARDNER, AND PAUL DEBEVEC. Capturing and Rendering with Incident Light Fields. In *Proceedings of Eurographics Workshop on Rendering*, pages 141–149, 2003. (Cited on page 15)
- [Vedaldi08] A. VEDALDI AND B. FULKERSON. VLFeat: An Open and Portable Library of Computer Vision Algorithms. <http://www.vlfeat.org/>, 2008. (Cited on page 85)
- [Velten12] A. VELTEN, T. WILLWACHER, O. GUPTA, A. VEERARAGHAVAN, M.G. BAWENDI, AND R. RASKAR. Recovering Three-Dimensional Shape around a Corner using Ultrafast Time-of-Flight Imaging. *Nature Communications*, 3:745, 2012. (Cited on page 14)
- [Wetzstein11] GORDON WETZSTEIN, IVO IHRKE, DOUGLAS LANMAN, AND WOLFGANG HEIDRICH. Computational Plenoptic Imaging. *Computer Graphics Forum*, 30(8):2397–2426, 2011. (Cited on page 9)
- [Weyrich09] TIM WEYRICH, PIETER PEERS, WOJCIECH MATUSIK, AND SZYMON RUSINKIEWICZ. Fabricating Microgeometry

- for Custom Surface Reflectance. *ACM Transactions on Graphics*, 28(3):32:1–32:6, 2009. (Cited on page 9)
- [Wilburn05] BENNETT WILBURN, NEEL JOSHI, VAIBHAV VAISH, EINO-VILLE TALVALA, EMILIO ANTUNEZ, ADAM BARTH, ANDREW ADAMS, MARK HOROWITZ, AND MARC LEVOY. High Performance Imaging using Large Camera Arrays. *ACM Transactions on Graphics*, 24(3):765–776, 2005. (Cited on page 65)
- [Wu10] HSIEN-HUANG P. WU AND SHIH-HSIN CHANG. Design of Stereoscopic Viewing System Based on a Compact Mirror and Dual Monitor. *Optical Engineering*, 49(2):027401–1–027401–6, 2010. (Cited on page 7)
- [Wu11a] CHANGCHANG WU. VisualSFM: A Visual Structure from Motion System. <http://ccwu.me/vsfm/>, 2011. (Cited on page 84)
- [Wu11b] CHANGCHANG WU, SAMEER AGARWAL, BRIAN CURLESS, AND STEVEN M. SEITZ. Multicore bundle adjustment. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 3057–3064. IEEE, 2011. (Cited on page 84)
- [Ying10] XIANGHUA YING, KUN PENG, REN REN, AND HONGBIN ZHA. Geometric Properties of Multiple Reflections in Catadioptric Camera with Two Planar Mirrors. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 1 – 8, 2010. (Cited on page 10)
- [Zhang99] ZHENGYOU ZHANG. Flexible Camera Calibration By Viewing a Plane From Unknown Orientations. In *Proceedings of IEEE International Conference on Computer Vision*, pages 666–673, 1999. (Cited on page 25)
- [Zhang06] LI ZHANG AND SHREE K. NAYAR. Projection Defocus Analysis for Scene Capture and Image Display. *ACM Transactions on Graphics*, 25:907–915, July 2006. (Cited on page 11)
- [Zhou11] CHANGYIN ZHOU AND SHREE NAYAR. Computational Cameras: Convergence of Optics and Processing. *IEEE Transactions on Image Processing*, 20(12):3322–3340, December 2011. (Cited on page 9)

Appendix

Length of the perpendiculars in Theorem 2

We refer to the sketch in Figure 4.11. Remember that $L = d(\mathbf{c}_0, m_C(\mathbf{a}))$, $r = d(m_C(\mathbf{a}), m_C(\mathbf{a}'))$, and $d = d(\mathcal{M}, \mathbf{c}_0)$.

We aim at expressing $d(\mathbf{c}_0, \mathbf{b}_1)$ in terms of L , r and d .

Since $(\angle \mathbf{h}_1 \mathbf{b}_1 \mathbf{c}_0 = \angle \mathbf{h}_1 m_C(\mathbf{a}') m_C(\mathbf{a}) = \frac{\pi}{2}$ and $\angle \mathbf{c}_0 \mathbf{h}_1 \mathbf{b}_1 = \angle m_C(\mathbf{a}') \mathbf{h}_1 m_C(\mathbf{a}) \Rightarrow \triangle \mathbf{c}_0 \mathbf{b}_1 \mathbf{h}_1 \sim \triangle m_C(\mathbf{a}') m_C(\mathbf{a}) \mathbf{h}_1$.

Then $\frac{d(\mathbf{c}_0, \mathbf{b}_1)}{r} = \frac{d(\mathbf{c}_0, \mathbf{h}_1)}{d(m_C(\mathbf{a}), \mathbf{h}_1)} \Rightarrow d(\mathbf{c}_0, \mathbf{b}_1) = \frac{r \cdot d(\mathbf{c}_0, \mathbf{h}_1)}{d(m_C(\mathbf{a}), \mathbf{h}_1)} = \frac{r \cdot d}{d(m_C(\mathbf{a}), \mathbf{h}_1)}$.

Applying the Pythagorean theorem to the triangle $\triangle \mathbf{c}_0 m_C(\mathbf{a}') m_C(\mathbf{a})$, we have: $d(\mathbf{c}_0, m_C(\mathbf{a}')) = \sqrt{(d(\mathbf{c}_0, m_C(\mathbf{a})))^2 - (d(m_C(\mathbf{a}), m_C(\mathbf{a}')))^2} = \sqrt{L^2 - r^2}$. Then $d(\mathbf{h}_1, m_C(\mathbf{a}')) = \sqrt{L^2 - r^2} - d$.

Applying the Pythagorean theorem to the triangle $\triangle \mathbf{h}_1 m_C(\mathbf{a}') m_C(\mathbf{a})$ yields: $d(m_C(\mathbf{a}), \mathbf{h}_1) = \sqrt{(d(\mathbf{h}_1, m_C(\mathbf{a}')))^2 + (d(m_C(\mathbf{a}), m_C(\mathbf{a}')))^2} = \sqrt{(\sqrt{L^2 - r^2} - d)^2 + r^2}$. Then:

$$d(\mathbf{c}_0, \mathbf{b}_1) = \frac{r \cdot d}{\sqrt{(\sqrt{L^2 - r^2} - d)^2 + r^2}}.$$

$d(\mathbf{c}, \mathbf{b}_2)$ is computed similarly to $d(\mathbf{c}, \mathbf{b}_1)$. Since $d(m_C(\mathbf{a}), m_C(\mathbf{a}'')) = d(m_C(\mathbf{a}), m_C(\mathbf{a}')) = r$ and $d(\mathbf{h}_2, \mathbf{c}) = d(\mathbf{h}_1, \mathbf{c}) = d \Rightarrow d(\mathbf{c}, \mathbf{b}_2) = d(\mathbf{c}, \mathbf{b}_1)$.

Practical Recommendations for Kaleidoscope Construction

Before starting the design and the construction of a real kaleidoscopic imaging system it is very important to plan the experimental details in advance. We suggest to select a mirror configuration having the following questions in mind:

- Which objects will be imaged in the kaleidoscope? What are their physical dimensions?
- Which cameras/projectors will be used (including lens selection)? What are their intrinsic parameters, including field-of-view? What are their depth-of-field parameters? What are their physical dimensions?
- How many virtual views are needed? How should the virtual views be distributed?
- What is the maximum size of the mirrors?
- How well does the kaleidoscopic visual hull approximate the visual hull?

Our approach is to perform computer simulations before the actual kaleidoscope construction. For this purpose we render views from inside the kaleidoscopes in PBRT¹. Renderings should be performed with representative objects to evaluate the approximation of the visual hull by the corresponding kaleidoscopic visual hull.

¹<http://www.pbrt.org>

The knowledge of the real camera/projector intrinsics allow us to properly set up the camera parameters of the rendering part, while the knowledge of the dimensions of the object allows us to estimate the dimensions of the mirrors.

Once the mirror size is determined and the camera/projector, the object, and the illumination are placed into the virtual scene, we can render the camera/projector image. This image may be processed in the same fashion as the image from the real kaleidoscope (note, all the mirror positions are exactly known in the simulation, i.e. the virtual scene is perfectly calibrated). In other words, having the rendered image, we can produce the kaleidoscopic visual hull, the labeling and the list of virtual cameras/projectors. The result now can be used to estimate if our configuration satisfies additional important requirements:

- Are the virtual objects in the depth-of-field range?
- Is the number and the distribution of the virtual views satisfactory?
- Is the kaleidoscopic visual hull a satisfactory approximation of the visual hull?

Although the acceptability of the kaleidoscopic visual hull for one particular object does not guarantee, in general, the acceptability of the kaleidoscopic visual hull for other objects of similar size, it is a good indicator for major problems.

Once a mirror design satisfying the above-mentioned requirements is found, it needs to be checked further (in order to accept, modify, or reject it) with the following questions:

- How can objects be inserted into / removed from the kaleidoscope (without touching the mirrors)?
- How will the scene be illuminated? Is there sufficient space in/around the kaleidoscope? Is there sufficient illumination power?
- How to physically build the proposed kaleidoscope (frame structure and attachment)?
- How will the kaleidoscope be calibrated?