

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

Ilya Reshtouski¹ Alkhazur Manakov¹ Ayush Bhandari²
 Ramesh Raskar² Hans-Peter Seidel³ Ivo Ihrke^{1,3,4}
¹Saarland University ²MIT Media Lab ³MPI Informatik ⁴INRIA

Abstract

We investigate the problem of identifying the position of a viewer inside a room of planar mirrors with unknown geometry in conjunction with the room's shape parameters. We consider the observations to consist of angularly resolved depth measurements of a single scene point that is being observed via many multi-bounce interactions with the specular room geometry.

Applications of this problem statement include areas such as calibration, acoustic echo cancelation and time-of-flight imaging. We theoretically analyze the problem and derive sufficient conditions for a combination of convex room geometry, observer, and scene point to be reconstructable. The resulting constructive algorithm is exponential in nature and, therefore, not directly applicable to practical scenarios.

To counter the situation, we propose theoretically devised geometric constraints that enable an efficient pruning of the solution space and develop a heuristic randomized search algorithm that uses these constraints to obtain an effective solution. We demonstrate the effectiveness of our algorithm on extensive simulations as well as in a challenging real-world calibration scenario.

1. Introduction

Mirrors have been used in a number of vision applications in the past. Examples using curved mirror surfaces include catadioptric imaging [21], reflectance [6, 8], and texture [12] measurement systems. Planar mirror systems have been used for multi-view imaging of flat [9] and extended depth samples [7, 19] and for confocal imaging [15, 17]. Most of these systems are designed to work with single bounce reflections and all of them have to be calibrated.

This task relates to the reconstruction of specular surfaces. Great progress has been made in recent years and a variety of methods based on shape from distortion [22, 4], shape from specularities [5], specular flow [1], and polarization analysis [16] have been proposed. The interested reader is referred to the recent survey article [10] for an in-depth discussion. All of the aforementioned techniques assume a single-bounce interaction to occur at the geometry surface, separating the incident light and the camera. An analysis of local multi-bounce specular interaction, both of the reflective and the refractive type has been given by Kutulakos and Steger [13]. Their main result is that a maximum of two specular interactions can be reconstructed from *local* measurements.

In this article, we consider the problem of determining the geometry of multiple planar mirrors and the pose of the camera with respect to this mirror configuration using measurements of apparent depth of a single scene point that is visible via many different multi-bounce interactions with the scene geometry. The pose estimation problem for a calibrated mirror geometry has recently been investigated by Ramalingam et al. [18]. In contrast, we aim at determining the geometry of a room of mirrors *in conjunction* with the camera pose. Our main result is that this is often possible given the depth-resolved imaging of many inter-reflections of a *single* scene point in case of a convex room geometry and we derive sufficient conditions for a configuration to be recoverable. Our analysis is restricted to room geometries of two dimensions. In practice, this lets us deal with $2\frac{1}{2}$ D rooms that have a floor and a ceiling orthogonal to the mirror walls. Simulation experiments show that a large class of room configurations is recoverable, however, the recovery rate decreases with the number of mirror walls. We also demonstrate our method in a practical calibration example where we estimate the mirror geometry and camera pose for a system with many inter-reflections.

By formulating the problem in terms of apparent distance of the observed point via multiple paths, our method applies to time-of-flight measurement systems like RADAR, active SONAR[14], and LIDAR [11], but also to acoustics [3, 2] where room geometry is to be inferred from the travel time of a pulse emitted by a speaker [20, 23]. Recently, the recovery of general geometry from exactly one indirect bounce in a LIDAR context has been described by Velten et al. [24].

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

In particular, our contributions are

- a theoretical model for multi-bounce time-of-flight in convex rooms,
- a general solution method that allows for limited field-of-view, limited observation time, and does not require the identification of reflection order on the measurements, and
- the application of these ideas to the calibration of a planar system of mirrors with previously unachievable generality.

2. Problem Formulation

We consider an abstract angularly resolved distance measurement system in a situation with specular multi-bounce paths, Fig. 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-image is difficult.

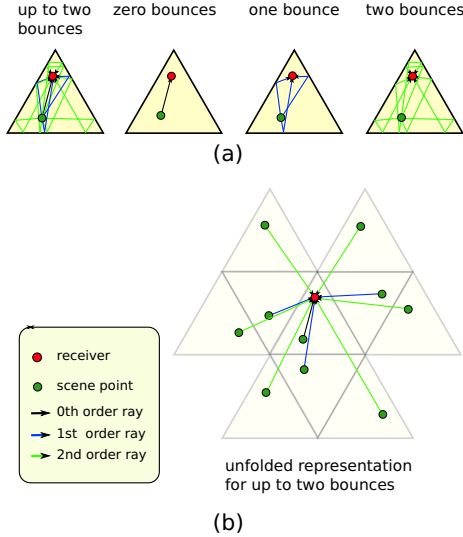


Figure 1. Problem definition: (a) A point is visible to an omni-directional receiver via multiple specular reflections. The receiver is capable of measuring the incidence angle and the distance of the point via the reflected ray paths without being able to detect the specular interaction. (b) To the 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.

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. 1 (b), we show how the receiver could naively interpret the surrounding world as a virtual mirror world consisting of many point objects at different distances. This interpretation is very similar to the unfolding operation introduced by Reshetouski et al. [19] in case the mirror geometry is known. Indeed, it can be taken as the definition of unfolding. In a sense, their paper describes the solution to a dual problem: Given the mirror geometry, compute the object. In the current paper, 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, to $2\frac{1}{2}D$ cases.

Any real mirror system will fail to subdivide the plane perfectly when generating the unfolded representation of the mirror world, the condition for a perfect division being that the mirror operations induced by the walls form a group structure. Therefore, a real system will include “lines of discontinuity” [19] as illustrated through a simulation in Fig. 2. Areas between two such lines indicate a common reflection sequence. The discontinuity lines correspond to mirror corners being hit by a ray bundle after a sequence of reflections. Intuitively, the bundle splits up at these points and traverses different mirror sequences thereafter. This implies that the measurement points are not necessarily




Figure 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 Reshetouski’s “lines of discontinuity” [19] that split areas of different mirror sequences. These lines are present in all but exceptional cases, please see the original paper for details. 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.

visible for every possible combination of mirrors into a sequence.

In the following, let $L_N = (M_1, \dots, M_N)$ be a sequence of N reflections through mirrors M_1 to M_N . As an example, there is exactly one ray among those that traverse the mirror sequence $L_1 = (2)$ that hits the object point. There is no ray amongst those continuing to sequence $L_2 = (2, 3)$ that sees the object point after these two reflections. However, the bundle of rays continuing to sequence $L_2 = (2, 1)$ 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 sequence $L_3 = (2, 3, 1)$ shows. After three reflections, one of the rays with this reflection sequence hits the object point. We refer to the set of all virtual points that are visible through a common reflection sequence as a *chamber*. The polygonal regions in Fig. 2 illustrate the chambers. The *base chamber* is the special chamber enclosing the receiver. In Fig. 2 it is marked in yellow.

In the following, we will discuss a solution for recovering the mirror room geometry including the occlusions introduced by non-group mirror systems as well as the position of the receiver within such system. 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 location 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 our algorithm through simulations and a challenging real-world example.

3. Overview

The paper is organized as follows. 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 configuration (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 Sect. 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, Sect. 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, Sect. 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 Sect. 6 that is based on random graph sampling to improve the reconstruction performance. Finally, we validate our algorithm in Sect. 7 via simulations and a real experiment.

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 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 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.

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 Sect. 6 we then extend our ideas to the case where the system layout is unknown.

Definition: A **doublet** is a pair of points $D = (A, A')$ separated by a single mirror reflection. The sequence of mirror operations leading up to this event is the same for the two points, i.e. $L_N^A = (M_1, \dots, M_N)$ and $L_N^{A'} = (M_1, \dots, M_N, M_{AA'})$, where $M_{AA'}$ is a mirror operation that is completely determined by the two points A and A' .

Properties: A doublet uniquely defines a mirror operation that takes point A into A' and vice versa. If we denote the receiver position by R , then $\overline{RA} < \overline{RA'}$ since an indirect view via mirror $M_{AA'}$ 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) without knowing the remaining room geometry. The doublet from the inset figure (marked) and all other doublets in the system of Fig. 2 are shown in Fig. 3 (left). Doublets belonging to the same physical mirror are marked with the same color. They usually appear in different locations of the virtual mirror world. The properties just mentioned can be verified in the figure.



Definition: A **triplet** is a pair of doublets $T = (D, D')$ that share 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. 3 (middle). It should be noted though that all adjacent doublets form triplets even though we only show a subset of them.

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


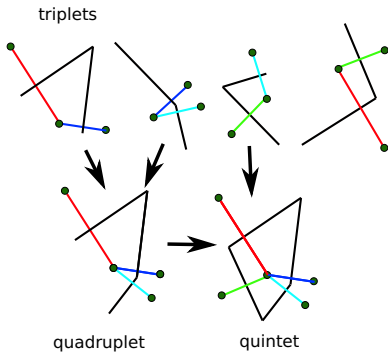


Figure 3. An overview of our approach. *Left*: Doublets are pairs of points that are separated by a single mirror reflection. They are indicative of mirror planes. Equivalent doublets are coded with the same color. As can be seen, doublets usually occur several times, making them stable features of the 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 from which the 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.

building blocks in constructing candidate configurations.

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. 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. 6.

In practice, the identification of doublets is based on their length 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. 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.

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 the inset figure.

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.

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. 3 (right). Assume we have constructed a candidate configuration at the position shown in red. We join the object point C in the reconstructed room with the receiver position R 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 C . If the candidate configuration is correct, this mirror is the last in the reflection sequence leading up to the observation of C . 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 Reshetouski unfolding [19]. Since C 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.

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




Figure 4. *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.

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. 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. 4 (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. 2 and 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.

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. 4 (right).


Our filtering operations exploit the geometric features of a mirror configuration that impose strong constraints on valid distributions of observation 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 AA' is a doublet generated by the mirror $M_{AA'}$, and D is the point of intersection of the mirror $M_{AA'}$ and the side RA' , then the intersection area of the inner part of the cone ARA' and the inner part of the circle with center D and radius AD must not contain any other observed points. To see this, consider the following. Point D is located on the mirror $M_{AA'}$, therefore, D and A belong to the same chamber C .



Using the assumption that our room is convex we can conclude, that AD entirely belongs to chamber C . Therefore, a unique sequence of reflections $L_N^A = (M_1, \dots, M_N)$ leads up to the reflection event at $M_{AA'}$. Therefore, any point A'' inside the cone ARA' is generated by a subsequence or super-sequence of L_N^A . Suppose we would like to generate an additional point A'' by reflecting point A from the mirrors of chamber C . In this case, we need to use more than one reflection of A from the sides of C . Because the direct observation distance is always shortest, the point A'' must be located outside the circle.


Potential Triplets



Triplets cannot be in an arbitrary orientation with respect to the receiver. The left of the figure shows an impossible configuration where A is both an image of B and B' . Conversely, if B and B' were images of A , they would be incorrect

because $\overline{RB} < \overline{RA}$, and similarly, $\overline{RB'} < \overline{RA}$ which is violating the fact that apparent distance increases with every reflection. The two sequences on the right are possible, indicating a sequence $L_{N+2} = (M_1, \dots, M_N, M_{BA}, M_{AB'})$ for the middle case, and a splitting ray bundle with a common subsequence up to A , $L_{N+1}^B = (M_1, \dots, M_N, M_{AB})$ and $L_{N+1}^{B'} = (M_1, \dots, M_N, M_{AB'})$ for the right case, respectively.

Doublet Compatibility




This filter can only determine that two potential 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 the figure 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, $L_{N+1}^{A'} = (M_1, \dots, M_N, M_{AA'})$ is incompatible with $L_{M+1}^{B'} = (M_1, \dots, M_M, M_{BB'})$ within the shaded region. For the middle case, $L_{M+1}^{B'} = (M_1, \dots, M_M, M_{BB'}) \subset L_{N+1}^{A'} = (M_1, \dots, M_N, M_{AA'})$, $M < N$ could be possible, again for the shaded region. In the situation on the right, the two doublets clearly 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. The figure 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 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. 4 (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. 4.2). However, due

to its exponential nature, the search is not very efficient and large problem instances might be practically unfeasible.

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 Sect. 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 Sect. 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 Sect. 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.

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 Sect. 3 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

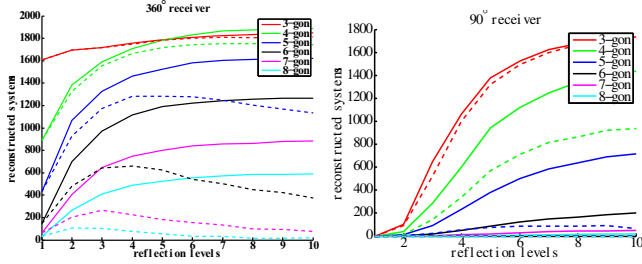


Figure 5. Simulation results for 2000 randomly generated n-gons. The plots show the number reconstructable systems versus the number of reflections considered for the reconstruction task. *Left*: full surround receiver, *Right*: field-of-view limited to 90° . The solid lines indicate the exhaustive search algorithm, the dashed lines the randomized search.




Figure 6. 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.

to a discrete number of positions. For intuition on that, please refer to Fig. 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. Experimental Results

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 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. *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).

For analysis, we performed an exhaustive search, as described in Sect. 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. 5 as solid lines. The plot shows the number of reconstructable room geometries. Then we ran the randomized search algorithm described in Sect. 6 that was working with the filtered graph structure, Fig. 4 (right). We again tested whether the geometry could be recovered. The results of this test are shown in Fig. 5 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. 6 (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. 6 (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.

7.2. 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




Figure 8. *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 fainter.

plan according to which the system was set up manually is shown in Fig. 7. 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 as a super-imposed outline on the sketch. A visual impression of the accuracy of our reconstruction can be gained from Fig. 8 (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 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.

8. Discussion and 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.

In the future we would like to investigate problem instances where the sufficient conditions derived in this paper do not hold. The goal would be to identify the class of reconstructable mirror systems. 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.

Acknowledgements

This work was supported by the German Research Foundation (DFG) through an Emmy-Noether fellowship IH 114/1-1.

References

[1] Y. Adato, Y. Vasilyev, O. Ben-Shahar, and T. Zickler. Towards a Theory of Shape from Specular Flow. In *Proc. ICCV*, pages 1–8, 2007. 1

[2] T. Ajdler, L. Sbaiz, and M. Vetterli. The Plenacoustic Function and its Sampling. *Signal Processing, IEEE Transactions on*, 54(10):3790–3804, 2006. 1

[3] J. Allen and D. Berkley. Image method for efficiently simulating small-room acoustics. *The Journal of the Acoustical Society of America*, 60(S1):S9–S9, 1976. 1

[4] T. Bonfort, P. Sturm, and P. Gargallo. General Specular Surface Triangulation. In *Proceedings of the Asian Conference on Computer Vision*, volume 2, pages 872–881, jan 2006. 1

[5] T. Chen, M. Goesele, and H.-P. Seidel. Mesostructure from specularity. In *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 17–22, 2006. 1

[6] K. Dana. BRDF/BTF Measurement Device. In *Proc. ICCV*, pages 460–466, 2001. 1

[7] K. Forbes, F. Nicolls, G. D. Jager, and A. Voigt. Shape-from-Silhouette with two Mirrors and an Uncalibrated Camera. In *Proc. ECCV*, pages 165–178, 2006. 1

[8] A. Ghosh, W. Heidrich, S. Achutha, and M. O’Toole. A Basis Illumination Approach to BRDF Measurement. 90(2):183–197, 2010. 1

[9] J. Y. Han and K. Perlin. Measuring Bidirectional Texture Reflectance with a Kaleidoscope. In *Proc. SIGGRAPH*, pages 741–748, 2003. 1

[10] I. Ihrke, K. N. Kutulakos, H. P. A. Lensch, M. Magnor, and W. Heidrich. Transparent and specular object reconstruction. *Computer Graphics Forum*, 29(8):2400–2426, 2010. 1

[11] A. Kirmani, T. Hutchison, J. Davis, and R. Raskar. Looking around the Corner using Transient Imaging. In *Proc. ICCV*, pages 159–166, 2009. 1

[12] S. Kuthirummal and S. K. Nayar. Multiview Radial Catadioptric Imaging for Scene Capture. *ACM TOG*, 25(3), 2006. 1

[13] K. N. Kutulakos and E. Steger. A Theory of Refractive and Specular 3D Shape by Light-Path Triangulation. *IJCV*, 76(1):13–29, 2008. 1

[14] B. La Cour. Statistical characterization of active sonar reverberation using extreme value theory. *IEEE Journal of Oceanic Engineering*, 29(2):310–316, april 2004. 1

[15] M. Levoy, B. Chen, V. Vaish, M. Horowitz, I. McDowall, and M. Bolas. Synthetic Aperture Confocal Imaging. *ACM TOG*, 23:825–834, August 2004. 1

[16] D. Miyazaki and K. Ikeuchi. Inverse Polarization Raytracing: Estimating Surface Shapes of Transparent Objects. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 2, pages 910–917, 2005. 1

[17] Y. Mukaigawa, S. Tagawa, J. Kim, R. Raskar, Y. Matsushita, and Y. Yagi. Hemispherical Confocal Imaging using Turtleback Reflector. In *Proc. ACCV*, pages 1–8, 2010. 1

[18] S. Ramalingam, S. Bouaziz, P. Sturm, and P. H. Torr. The Light-Path Less Traveled. In *Proc. CVPR*, pages 3145 – 3152, 2011. 1

[19] I. Reshetouski, A. Manakov, H.-P. Seidel, and I. Ihrke. Three-Dimensional Kaleidoscopic Imaging. In *Proc. CVPR*, pages 353–360, 2011. 1, 2, 5

[20] F. Ribeiro, D. Florencio, D. Ba, and C. Zhang. Geometrically Constrained Room Modeling With Compact Microphone Arrays. *IEEE Transactions on Audio, Speech, and Language Processing*, 20(5):1449–1460, 2012. 1

[21] R. Swaminathan, M. D. Grossberg, and S. K. Nayar. Non-Single Viewpoint Catadioptric Cameras: Geometry and Analysis. *International Journal of Computer Vision (IJCV)*, 66(3):211–229, 2006. 1

[22] M. Tarini, H. P. A. Lensch, M. Goesele, and H.-P. Seidel. 3D Acquisition of Mirroring Objects. *Graphical Models*, 67(4):233–259, 2005. 1

[23] S. Tervo and T. Tossavainen. 3D Room Geometry Estimation from Measured Impulse Responses. In *Proc. ICASSP*, pages 513–516, 2012. 1

[24] A. Velten, T. Willwacher, O. Gupta, A. Veeraraghavan, M. Bawendi, and R. Raskar. Recovering Three-Dimensional Shape around a Corner using Ultrafast Time-of-Flight Imaging. *Nat. Comm.*, 3:745, 2012. 1