DOI: 10.1111/cgf.12201

COMPUTER GRAPHICS *forum*
Volume 32 (2013), number 8 pp. 246–257

Design and Fabrication of Faceted Mirror Arrays for Light Field Capture

Martin Fuchs[1], Markus Kächele[2] and Szymon Rusinkiewicz[3]

[1]University of Stuttgart, Germany
martin.fuchs@visus.uni-stuttgart.de
[2]University of Ulm, Germany
markus.kaechele@uni-ulm.de
[3]Princeton University, Princeton, NJ, USA
smr@princeton.edu

Abstract

The high resolution of digital cameras has made single-shot, single-sensor acquisition of light fields feasible, though considerable design effort is still necessary in order to construct the necessary collection of optical elements for particular acquisition scenarios. This paper explores a pipeline for designing, fabricating and utilizing faceted mirror arrays which simplifies this task. The foundation of the pipeline is an interactive tool that automatically optimizes for mirror designs while exposing to the user a set of intuitive parameters for light field quality and manufacturing constraints. We investigate two manufacturing processes for automatic fabrication of the resulting designs: one is based on CNC milling, polishing, and plating of one solid work piece, while the other involves assembly of CNC-cut mirror facets. We demonstrate results for refocusing in a macro photography scenario. In addition, we observe that traditional photographic parameters take novel roles in the faceted mirror array setup and discuss their influence.

Keywords: catadioptric, fabrication, light field recording, mirror design

ACM CCS: 2012 Computing Methodologies; Artificial Intelligence; Computer Vision; Image and Video Acquisition; Computational Photography

ACM CCS 1998: I.4.1 Image Processing and Computer Vision; Digitization and Image Capture; Imaging geometry

1. Introduction

There is an inherent tension in the design of digital cameras: on the one hand, they must enable capturing the moment spontaneously ('point-and-shoot'); on the other, they are to deliver beautiful pictures that result from a painstaking choice of photographic parameters (such as viewpoint, focus and aperture). Manufacturers have tried to give non-experts the ability to obtain high quality by providing automatic assistants (most prominently auto-focus lenses and auto-exposure metering). These, however, have evolved into a plethora of disparate motif programs—recent consumer cameras often provide dozens of scene modes, including presets for fireworks, pets, food, children, etc. and software options for automatically focusing on human faces and detecting expressions such as smiles or blinks. The user is once again faced with having to make so many choices *before* a picture is taken that spontaneity and great photography seem ultimately incompatible.

Going beyond quality and usability, there is one design parameter in which camera manufacturers have enjoyed unquestioned success: resolution. There is, however, a profound and increasing discrepancy between typical display and capture resolution: modern digital SLR cameras record more than ten times as many pixels as an HDTV display can show. This suggests that, before it is displayed, most recorded footage is either cropped or rescaled, making no use of the available resolution.

The excess sensor resolution may be better used to facilitate *post hoc-photography*: shoot now, choose parameters later. Doing this requires collecting many rays from the scene—a light field [LH96]—and previous systems have used modified cameras with lenslet arrays

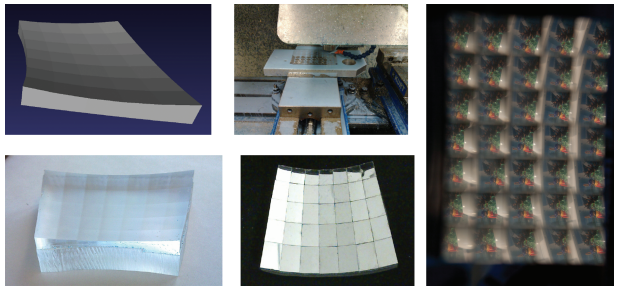

Figure 1: *Design and fabrication pipeline. After computing an optimal shape, we use CNC machining to form a piece of acrylic into a support structure for a set of precision-cut polished stainless steel segments. This mirror may be used for single-shot capture of a light field with 35 views.*

to capture sufficient data for refocusing and small viewpoint change [NLB*05, GZC*06].

We investigate an alternate approach based on combining high-resolution cameras with segmented mirrors (Figure 1). These systems can be used in conjunction with conventional digital cameras, do not require modifying camera internals (such as by breaking apart the camera body or cutting open a lens), are not heavier or more sensitive to transport than commodity DSLR lenses, and can be purpose-built for roughly the same cost as quality lenses (under USD 1000).

This paper presents a pipeline for designing, manufacturing and utilizing such mirror arrays for catadioptric light field capture. Although we restrict the design space for our faceted mirrors, there remains a bewildering set of design choices. Moreover, most of the design parameters are not directly related to the user's goals, which are phrased in terms of how practical the mirror will be to manufacture and the resolution, coverage and quality of the resulting light field. Therefore, the heart of our pipeline is an interactive tool (Section 3): the user is free to specify constraints and explore trade-offs phrased in terms of intuitive quantities, while the program automatically optimizes for mirror parameters and provides feedback on manufacturing cost. We explore several case studies of mirrors designed using our program, with very different angular coverage, working volume, and intended imager (Section 4).

Once a mirror has been designed, it must be fabricated, and we have investigated two different manufacturing technologies (Section 6). The first approach is to CNC-mill a complete multi-facet mirror surface out of a single block of brass, then polish and plate it. The second approach is to cut discrete mirror facets and assemble them onto a support structure CNC-milled from plastic. We will discuss the challenges these methods pose and their respective advantages.

In extension to a previous publication at VMV 2012 [FKR12], this invited version of the paper also discusses the novel roles of traditional photographic parameters—aperture settings, zoom and small changes in position (Section 5)—and provides additional visualizations in the form of ray diagrams and epipolar plots.

2. Basics and Related Work

Light fields (introduced to graphics by Levoy and Hanrahan [LH96], and, under the name 'lumigraph', by Gortler *et al*. [GGSC96]) exhaustively model the radiance emitted by a static scene, and, outside participating media, allow computing images for any describable camera. Thus, they enable photographic applications such as refocusing, viewpoint change, and synthetic aperture modification. As a data structure, a light field maps a ray in space identified by four spatial coordinates to the radiance transported along that ray.

The task of light field acquisition is to sample the 4-D space of rays, usually along 2-D slices. Initial approaches assumed a static scene, and therefore could record it with a moving camera. This could either be hand-held [GGSC96] or moved using an automated gantry [LH96] or motorized translation stage [UWH*03, IMG00]. Later techniques sped up the acquisition, often enabling capture of dynamic scenes:

Camera arrays combine the simultaneous exposures of dozens of cameras to acquire the light field [YEBM02, WJV*05]. This enables video capability, and can be used for feeding a 3D display in real time [MP04], but requires a large and costly setup.

Lens arrays rely on optical components to slice the light field. Yang [Yan00] arranged 88 lenses in front of a flatbed scanner in an inexpensive setup for static scenes. Other designs require only a single, static 2-D sensor and allow near instantaneous capture: Ng *et al*. [NLB*05, Ng06] added a micro-lens array on the sensor of a digital SLR camera, thus splitting the rays going through the camera aperture into subpixel sets according to their direction, enabling a compact, transportable, stable setup, but it requires camera internals to be modified and limits the application: no larger aperture can be simulated, and the viewpoint remains largely unchanged, as all captured rays have come through the camera lens aperture. Georgiev *et al*. [GZC*06] explored a lens array embedded into the main lens, improving on the usable sensor area. Levin *et al*. [LHG*09] presented a 4-D frequency analysis of a multi-lens-based setup optimized for refocusing. Cossairt *et al*. [CNR08] used a lens array at a considerable distance from the main camera, outside its body for light field transfer. This enables a wide sampling range; however, since the lens array has reflective surfaces facing the camera, either the light from this side must be blocked (enlarging the camera body), or the lens must be employed in a carefully chosen (indoor) setup that eliminates stray reflections.

Masks in the optical path close to the camera sensor permit reconstruction of full-resolution pictures for certain scene types, as demonstrated by Veeraraghavan *et al*. [VRA*07, VAR*08]. An additional mask close to the camera primary lens adds trade-offs on time versus spatial and angular resolution [AVR10]. Wetzstein *et al*. [WIH10] showed a high-dynamic-range application for a modulating camera, while Ihrke *et al*. [IWH10] investigated the interaction between spatial and Fourier multiplexed patterns.

Mirrors use reflections off surfaces to slice the light field. Thus, in contrast to designs based around refractive elements, they do not need to avoid these reflections on the camera-visible surfaces, do not require the optical path to be encased in large, light-tight camera bodies, and place no restrictions on scene illumination. Analogously to multiple, spatially arranged lenses, designs with multiple mirrors

enable instantaneous capture, as demonstrated with an array of planar mirrors by Levoy *et al*. [LCV*04]. Curved mirrors increase the field of view that can be observed; they can also be arranged in a plane and still record a dense light field, which makes the construction easier. Unger *et al*. [UWH*03] used a grid of 12 × 12 mirror balls for light field capture, while Lanman *et al*. [LCWT06] used sections of 31 spherical mirrors. Especially for full spheres, a wide field of view can be achieved, as has been shown by Taguchi *et al*. with an array of mirrors [TAV*10] and with a single mirror and a moving camera [TARV10]. Outside light field acquisition, curved mirrors have been used to great effect to observe a single surface point (or small patch) under many directions at once, making use of an ellipsoid/spherical [War92, MSY07] or paraboloid [Dan01] geometry. Ghosh *et al*. [GHAO10] used the same mirror setup also for illumination, while Peers *et al*. [PHD06] used a rough mirror for relighting. Mukaigawa *et al*. [MTK*11] designed a polyhedral mirror based on the geometric properties of ellipsoids for a design that allows the placement of objects at one focal point and a camera at the other to perform hemispherical confocal imaging.

This work strives to construct a setup that functions as an external augmentation to existing digital cameras, providing single-exposure recoding capability. It should be compact and inexpensive, and, once manufactured, remain stable without additional fine-tuning. We wish to allow designs optimized for a variety of rendering applications, ranging from refocusing to large synthetic apertures and modest changes in viewpoint. This leads us to choose a mirror-based construction built with rapid manufacturing tools.

With regards to geometry, we choose a piece-wise planar design, as it is does not introduce distortions while allowing effective recording of small working volumes. In contrast to [LCV*04], we aim for a continuous surface, as this avoids gaps and occlusions. In contrast to Mukaigawa *et al*. [MTK*11], we provide a user guided optimization pipeline for the derivation of the geometric configuration.

3. Designing Optimal Mirror Geometry

3.1. Choosing a family of mirrors

Let us consider two possible extremes for the shape of a mirror surface. A *planar* mirror yields a single virtual viewpoint with field of view identical to the camera's. In contrast, a mirror in the shape of an *ellipsoid of revolution*, as used by Ward *et al*. [War92] for BRDF measurement, enables observing a fixed point from a large, continuous range of directions. Its shape is uniquely defined by its focal points—the camera position in A, the scene centre in B—and the sum of distances from its surface to A and B (Figure 2).

We can combine the above ideas into a *segmented ellipsoidal mirror* composed of planar facets: it observes a scene from many viewpoints, each of which contains a conventional virtual camera. For any ellipsoidal basis shape, we can define the facets by intersecting one (central) ray per segment from the camera against the ellipsoid, and intersecting the planes that contain the intersection points and are tangent to the ellipsoid against their neighbours. By distributing the intersection points of the viewing rays equally on the image sensor, we can make sure that roughly the same number of pixels per segment contributes to the light field.

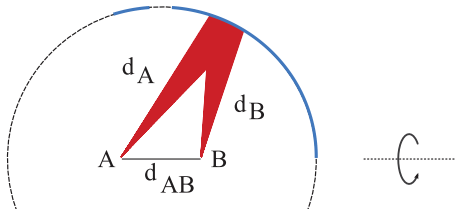

Figure 2: *Consider a camera in point A and a scene point B. The mirror surface that reflects rays from each camera direction into B is the one that keeps the length of each ray pair $d_A + d_B$ constant. It is a subsection of an ellipsoid of revolution uniquely defined by these variables. Not considering global translation and rotation, the shape depends only on the distance d_{AB} between A and B and the sum $d_A + d_B$.*

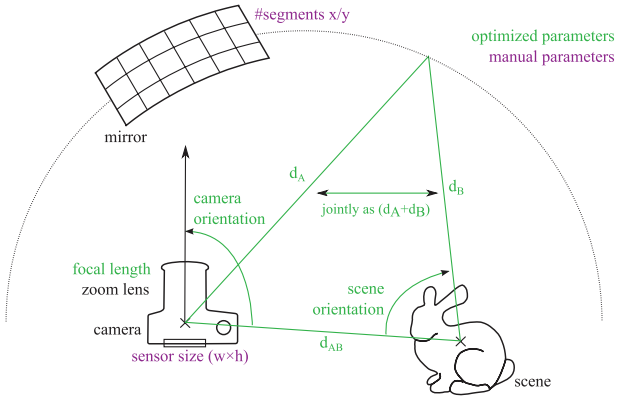

Figure 3: *Optimized and manually chosen parameters for a faceted mirror light field recording setup.*

In determining the parameters of the ellipsoid basis shape, it is sufficient to model its behaviour in a 2-D cross-section that contains both points A and B, as rotational ellipsoids are symmetric with respect to the axis $\overline{AB}$. We align $\overline{AB}$ with the user's choice of either of the image dimensions.

This leaves the following parameters to be chosen (see Figure 3): for the camera, the *focal length* and the *orientation* (i.e. which part of the ellipse it observes, expressed as rotation around the 'up' axis); for the setup: the *distance between scene and camera* d_{AB} and, if the scene is not spherical, the *orientation of the scene*; and, finally, for the mirror, the *sum of distances* $d_A + d_B$ and the *size of segments* (both in-plane along the ellipse and around the axis of revolution). The latter can be computed if the number of segments (corresponding to the number of achievable views) and the sensor dimensions of the camera are known. We treat these as user input.

3.2. Mirror design goals

These parameters interact in complex ways, and one of our key observations is that there is not an immediate connection between them and a designer's high-level goals for light field capture. We propose an automatic optimization of these parameters

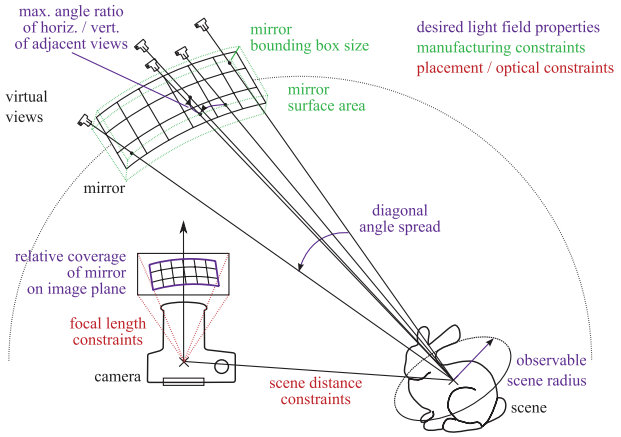

Figure 4: *Desiderata of the optimization, including mirror manufacturing/setup constraints, and light field quality/efficiency measures.*

so as to minimize an objective function (Figure 5), which encodes a trade-off between different desiderata for the light field quality (Figure 4):

- **Resolution** of the acquired light field: the number of views (that is, the manually chosen number of segments), and the number of sensor pixels per view (this follows from the fixed resolution of the sensor and the relative area of the sensor used, expressed below).
- **Coverage** of the light field: the angular coverage/spread of observed viewpoints, and the maximum scene size that can be observed with a given mirror/camera combination.
- **Efficiency** of the acquisition: the *directional uniformity* of views, and the *relative area* of the sensor used. (The sensor is rectangular, but the mirror's projection into the camera view is not: the better the mirror fills the available sensor space, the more rays can be observed. We can also trade off efficiency for completeness of coverage by permitting the outer segments to partially exceed the sensor.)

In addition, the mirror shape cannot be chosen arbitrarily, but is subject to two types of limitations:

- **Manufacturing constraints:** the cost of mirror and support material increases with volume; as does the cost of manufacturing (a CNC mill, for instance, needs to run longer if it is to ablate more material) and may further increase with surface area (when electroplating, the surface needs polishing after cutting with a dense tool path).
- **Placement/optical constraints:** the physical size of equipment dictates a minimum and maximum distance between the scene and the camera, while the availability of objective lenses places limits on focal length.

In order to design a mirror array satisfying the user's goals and constraints, we turn to optimization. We express the desirables and limitations in a single (non-linear) objective function (Figure 5), and use the `levmar` [Lou04] optimizer to find the best-compromise mirror parameters.

$$\begin{aligned}\mathcal{E} = {} & T^{\max}_{\text{surface area}} + T^{\max}_{\text{bounding box width}} \\ & + T^{\max}_{\text{bound. box height}} + T^{\max}_{\text{bound. box depth}} + T_{\text{rel. sensor coverage}} \\ & + T_{\text{scene radius}} + T_{\text{diagonal angle spread}} \\ & + T_{\text{angle uniformity}} + T^{\max}_{\text{focal length}} + T^{\min}_{\text{focal length}} \end{aligned} \quad (1)$$

where

$$T_x := \lambda_x^2 \cdot \left(x_{\text{desired}} - x_{\text{actual}}\right)^2 \quad (2)$$

$$T_x^{\max} := \lambda_{x_{\max}}^2 \cdot \max\left\{0, \left(x_{\text{actual}} - x_{\text{desired}_{\max}}\right)\right\}^2 \quad (3)$$

$$T_x^{\min} := \lambda_{x_{\min}}^2 \cdot \max\left\{0, \left(x_{\text{desired}_{\min}} - x_{\text{actual}}\right)\right\}^2 \quad (4)$$

Figure 5: *Objective function for the nonlinear optimization. The individual objectives are weighted by user-controlled parameters $\lambda\bullet$. For the objectives with the $T\bullet$ terms, any deviation from the desired value incurs a penalty; for the others, only a deviation beyond ($T_\bullet^{\max}$) or below ($T_\bullet^{\min}$) the limit increases the cost function.*

3.3. Interactive mirror design

The terms of the objective function are individually weighted with user-chosen parameters $\lambda_\bullet$, initialized to 1. As there is no intrinsic best way of defining the trade-off between the parameters, we leave the choice to the user, providing instant feedback in an interactive application to support exploration of the solution space (Figure 6 and Video S1).

A user defines the desired properties and weights and promptly observes the actually achieved statistics of the current solution. He/she can also define the variables that are not subject to the optimization: number of views/segments and sensor dimensions. The user interface colours objectives which either have been exactly met (min/max limits) or are within 1% of the desired value green, red otherwise.

Behind the scenes, at each iteration the software computes the 3-D mirror shape and generates the statistics for the objective function. Surface area, bounding box, and focal length constraint terms are immediately given by the resulting geometry. Sensor coverage is estimated by tracing 10,000 rays from the camera towards the mirror and counting missed intersections. We model the angular spread as the maximum angle at the scene centre between any pair of virtual camera positions (the reflections of the real camera's optical centre in the mirror facets), and express uniformity of the light field as the ratio between the maximum and minimum angles between neighbouring virtual camera positions.

The observable scene extent can be expressed as follows: we consider a plane containing the centre point of the scene and intersect the reflected rays of the camera's projection centre to the edge midpoints of all mirror segments with this plane. The minimum distance of these intersection points to the scene centre gives a radius of scene points that can be observed by all virtual cameras. We let the optimizer solve for the optimal orientation of the scene. As the optimizer does not take rays that miss the scene into account, it can artificially under-report the usable scene volume. However, for

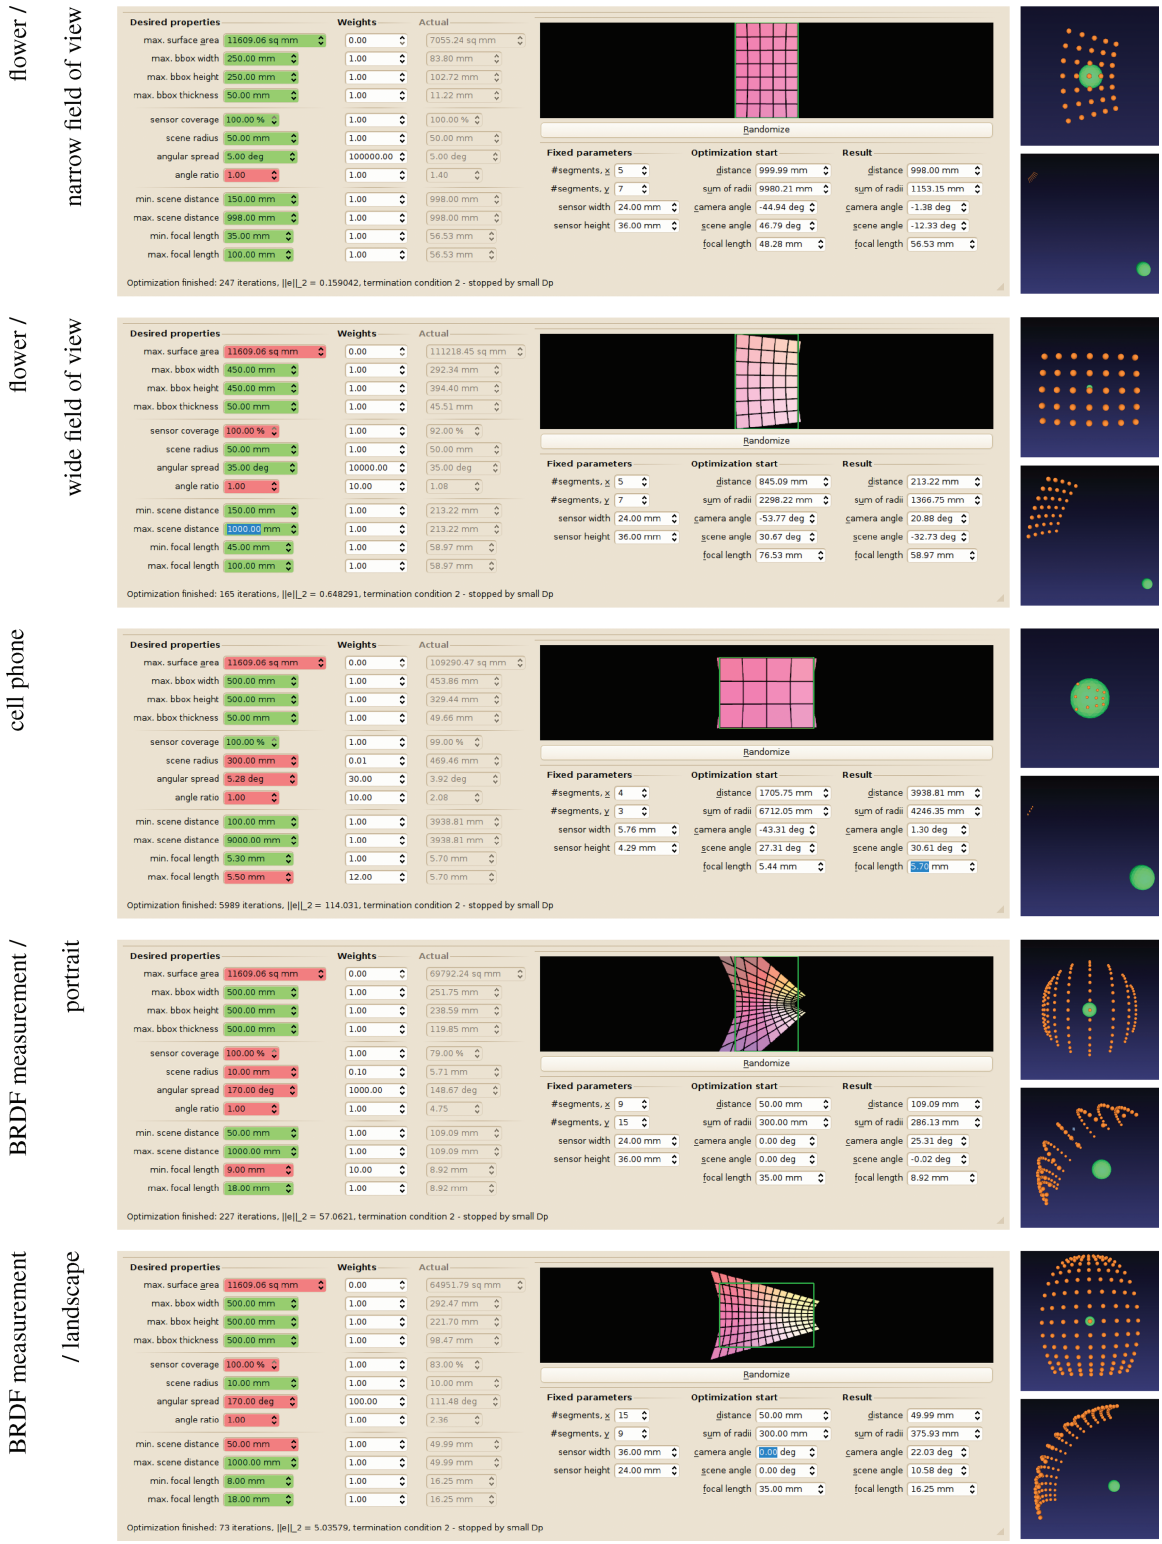

Figure 6: *Optimization program: objectives and achieved properties in the end state. An OpenGL window shows an instant visualization of the mirror shape in comparison to the picture visible in the camera (green rectangle). In the rightmost column: virtual camera locations (orange) compared to the observable scene centre (green). The diameter of the green sphere corresponds to the diameter of the green measurement line in the ray diagrams from Figure 8.*

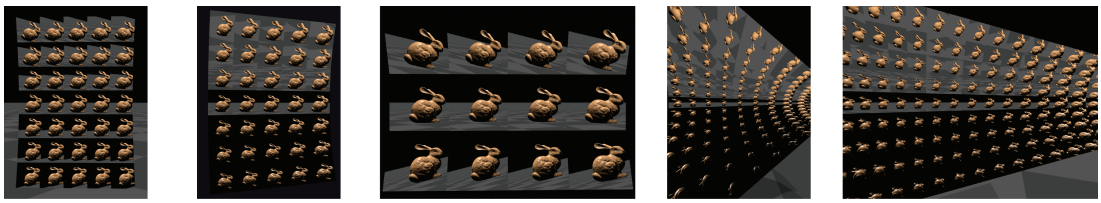

Figure 7: *Ray tracings for the mirror designs used in the case study: the Stanford bunny is rendered floating atop a checkerboard as seen from the camera. The bunnies have been individually scaled to fill the field of view; their absolute size is comparable within the pairs of 'macro' and 'material measurement' examples.*

the parameter it finds, it guarantees that the reported scene volume can be completely observed.

The convergence speed varies with the number of active constraints and the error landscape; our implementation typically converges in less than one second on a commodity PC.

4. Example Mirror Shapes

In order to show the flexibility provided by our interactive mirror designer, as well as the ease with which a user can find an appropriate design, we consider a number of example designs. The scenarios involve different imagers, different working volumes, and vastly different goals. In addition to the discussions in this section, we will provide results for a physical realization of the macro/narrow field of view design in Section 7 and Figures 1 and 15.

For every individual mirror design, we generate a set of evaluations: in addition to the statistics provided during the optimization process, we perform a ray casting in 2-D to show the spatial geometry of the setup and the distribution of rays in space (Figure 8). Inside this diagram, we define a line segment (green) centred at the scene centre position, and plot the locations and directions of all intersecting rays in an epipolar plot (Figure 9). For a camera with infinite resolution, the individual dots would connect to lines from the bottom left to the upper right; the number of lines corresponds to the number of views, or the number of mirror segments in the plane containing its symmetry axis.

For additional reference, we show ray traced simulations of the resulting mirror shapes in Figure 7.

Macro/narrow field of view: First, we consider a macro photography scenario for the capture of scenes 50 mm in diameter (see Figure 6 for a screenshot of the optimizer with all parameters and properties). We design the mirror for a full-frame DSLR and aim at a refocusing application. This requires the virtual camera positions to be close together, implying a small angular spread of 5 degrees. The resulting mirror is nearly planar, with rather shallow angles between facets. Constraints are almost all satisfied, the only blemish being angle ratio (at 1.4). This is the design we manufactured and used for Figure 1 and the results shown in Figure 15.

Macro/wide field of view: We modify the design to record the same scene type, but with a multiview stereo reconstruction application in mind (which requires a greater range of angles), increasing the angular spread to 35°, while relaxing some of the manufacturing constraints. The optimization results in a design that leads to a highly uniform distribution of virtual cameras trading in some sensor coverage (92% left) and a slightly higher visual distortion (Figure 7, second from left) compared to the previous design.

Cell phone: The next scenario investigates a design that employs a small sensor typical for cell phone built-in cameras (5.76 × 4.29 mm) for a scene diameter of 300 mm. This results in a comparatively large mirror (454 × 330 × 50 mm). In this case a 'perfect' design is hard to achieve, requiring a compromise among several desirables (e.g. the resulting scene radius is 469 mm). As the resolution of such a camera is lower than that of a DSLR, fewer segments must suffice.

Material measurement: Given the prior use of ellipsoidal mirrors for BRDF measurement, it is natural to use our interactive design program to find a construction that provides many views under a large set of angles. In contrast to mirrors with continuous curvature, our segmented design enables the observation of an extended scene region (10 mm × 10 mm) which enables the capture of a surface light field [WAA*00] of a small material patch; combined with controlled, changing illumination, this could be used to capture a spatially varying BRDF.

This design, which aligns the symmetry axis of the ellipsoid with the shorter image edge (*portrait* mode), provides the largest angular coverage among those we explored, at the cost of the greatest distortion (Figure 7, fourth from the left) and greatest inefficiency with respect to usage of the sensor area (79%). Note that the distortions on the right-hand side of Figure 7 (third from the left) reduce the volume that can actually be observed by all facets considerably.

To amend the deficiencies of the above configuration, we also investigate a design which rotates the camera back to *landscape* orientation, aligning the larger image edge with the symmetry axis, and reduce the weight of the angular coverage penalty term. As could be expected, this results in a lower angular spread (111° vs. 149° in the portrait case), but also brings a considerable improvement with respect to distortions (Figure 7).

5. Effect of Photographic Parameters

In a conventional recording scenario, in which a scene is acquired by a camera with a single sensor and a single lens, several photographic

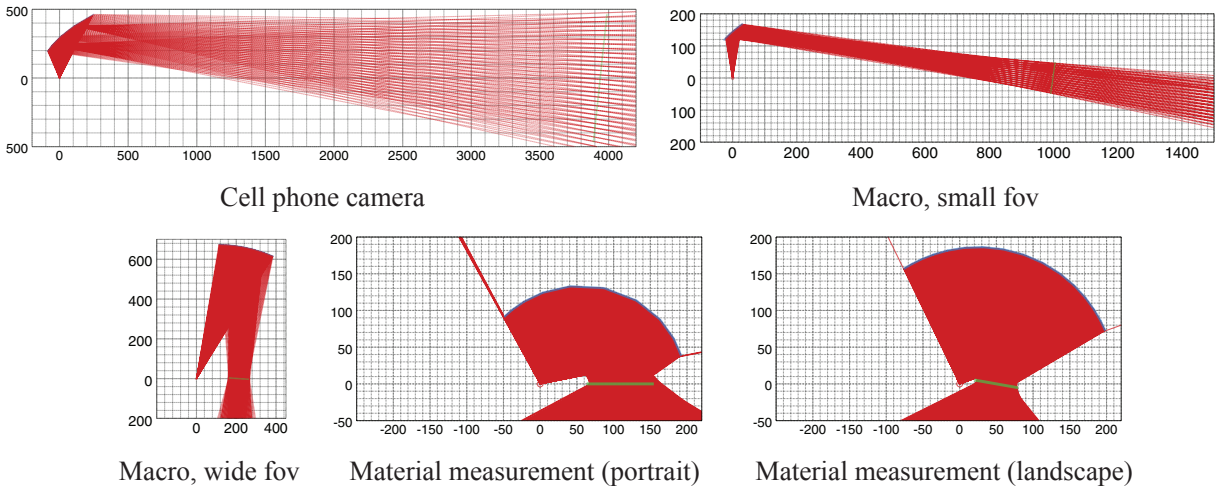

Figure 8: *2D Ray diagrams visualize the path of rays in the symmetry plane of the mirror: rays (red) arrive in the camera eye point (red circle) after reflection off the mirror segments (blue). The green line segment represents the scene diameter. These plots visualize the relative size, position and orientation of scene, mirror segment and camera; when printed in 1:1 scale, they are also useful for establishing correct alignment. All distances are in mm.*

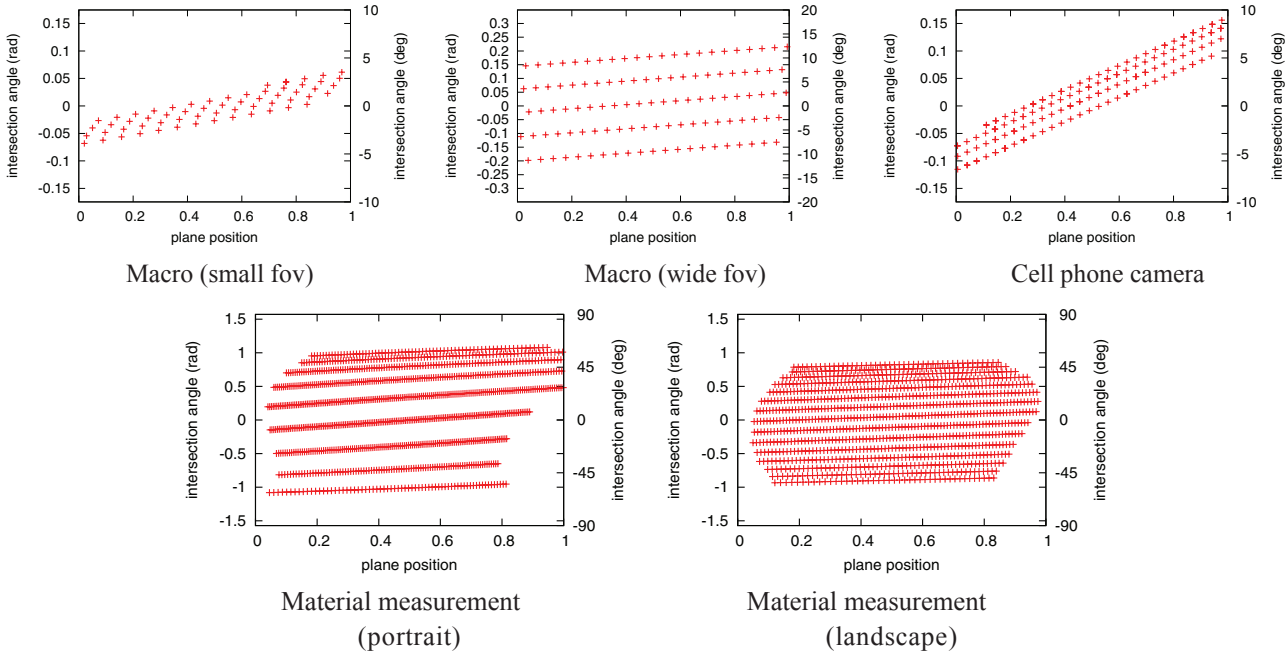

Figure 9: *Epipolar plots: the ray diagrams shown in Figure 8 provide little insight into the achievable light field quality, as the distribution of rays is hard to interpret. Therefore, we provide epipolar plots that show the angle and position for ray intersections with the scene diameter (green line segment in Figure 8). In actual use, the total number of plotted points corresponds to the number of camera pixels observing the mirror. Vertically, the sampling density is controlled by the number of line segments, horizontally the sampling density increases with camera resolution.*

parameters can be varied; among them are camera position with respect to the scene, the aperture of the entrance pupil, and the focal length of the camera lens. The same is true for a faceted mirror array setup; however, the impact of these parameters is radically different. In this section, we will discuss their influence with geometrical considerations, demonstrated by means of ray diagrams and ray traced imagery.

The *focal length* of the lens conventionally facilitates a trade-off: as the focal length increases, the field of view narrows, losing

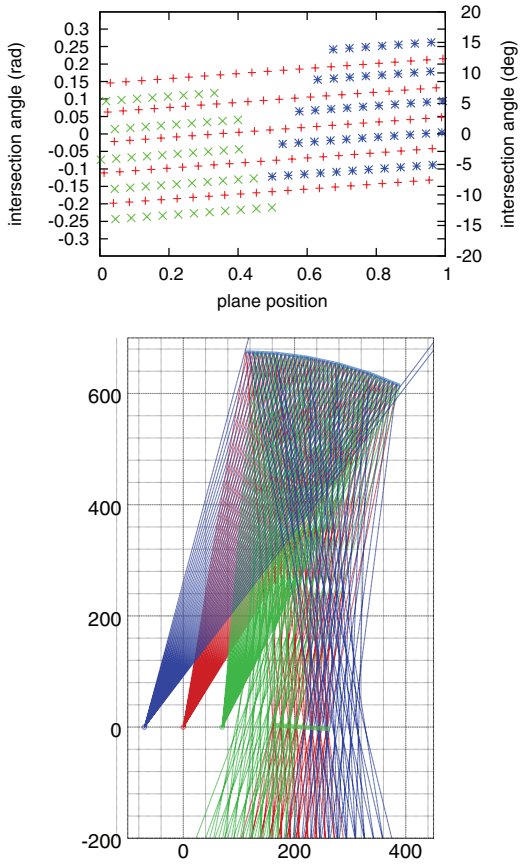

Figure 10: *Translating the camera reveals in-between views (each colour represents a different viewpoint). This figure shows the 'macro, wide field of view' mirror design, which has not been optimized to exploit this effect, thus leaving a gap in the set of observable rays.*

observed working volume in return for gained spatial resolution. In our method, while even a single mirror facet remains in full view, increasing the focal length maintains the observed working volume; here, resolution is gained in exchange for the number of views corresponding to total angular coverage of the scene.

A zoom objective makes changing the focal length quick and effortless and permits to construct a setup which quickly changes between recording single views with high resolution or many views with reduced resolution—or any desired setting in-between. For a parfocal lens, this is equivalent to a crop of the recorded image while increasing the resolution; for a varifocal lens, refocusing may be required between subsequent recordings.

In changing the *camera position*, a traditional camera can be used to change the working volume which is observed; with many exposures at varying positions [GGSC96, LH96], a full light field can be eventually assembled.

With a segmented, ellipsoid mirror, each of the virtual cameras contributes a picture of its own; for moderate changes in position, these pictures are inserted in-between views recorded in previous exposures (Figure 10). The sampling gap in the centre of the scene would be avoidable, if a smaller focal length would be used in conjunction with a larger mirror that would need to provide one more facet on each side. This creates another trade-off: a larger mirror with correspondingly higher manufacturing cost enables a more flexible use of the setup.

Perhaps most interesting is the effect of an extended *aperture* which is responsible for out-of-focus blur. In a conventional camera, the aperture size and shape control how much scene silhouettes are blurred, and which shape the blur kernel assumes. In light field photography, this out-of-focus blur can be generated in software (after appropriate view interpolation, see Figure 15).

As the aperture of the recording camera opens, it, too, will in due course affect the picture of the scene in the same way as in a conventional camera and lose high frequencies in the recorded light field. However, as the optical path from the camera to the scene is considerably longer than the path from the camera to the mirror, long before a large aperture impacts the sharpness of the recording, another, interesting effect happens: while the contours of mirror segments blur against the mirror background, the observed scene stays sharp; instead of blurring its contours, views start to blend into neighbouring views (Figure 11).

The ray diagram in Figure 12 explains the effect: finite ray bundles which hit the centre of a mirror segment are wholly reflected and focus again on a single point in the scene. Ray bundles which hit a boundary, though, are split up into two or more sub-bundles that

Figure 11: *The Stanford bunny reflected by a mirror in front of a checkerboard for reference. As we increase the simulated aperture diameter, the mirror contours blur, while the bunny stays in focus for a long time. As a result of the finite aperture, the pictures of the bunny form a blend of neighbouring views, thus extending the observed volume — eventually, the bunny appears fully visible, even though it is too large to fit into a single view under pinhole aperture conditions.*

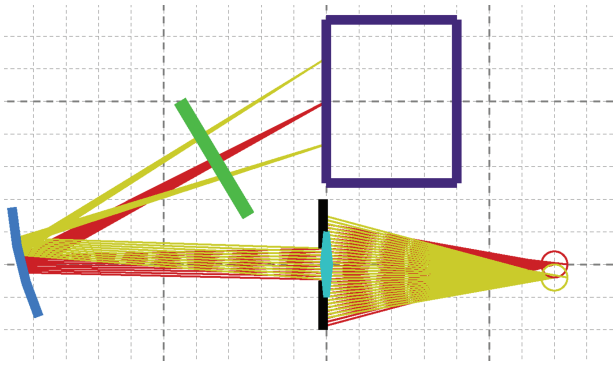

Figure 12: *With a finite aperture, centre ray bundles (red) get focused on a single point in the scene. Beams that hit a fragment boundary (yellow) split up into distinct parts. The weights for the split follow from the mirror geometry.*

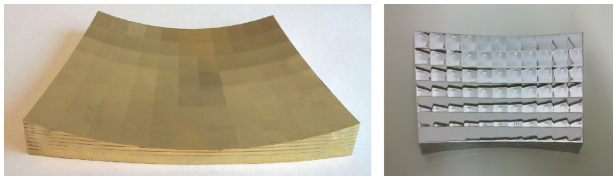

Figure 13: *Mirror construction from electroplated brass. The mirror shape D is cut out from a solid brass block with a CNC mill. Polishing and electroplating with a bright nickel finish creates the desired mirror surface.*

individually converge on different points on the scene—hence, the scene discontinuities are not blurred, but images of virtual views add up discretely in a linear combination with weights that vary from pixel to pixel, depending on what fraction the beam is split in hits a scene point. The weights of this linear combination are only dependent on the mirror geometry and location, and can therefore be known beforehand. This suggests that, with several exposures of varying aperture, the working volume might be extendable for static scenes by solving a set of linear equations which separate the individual contributions. However, due to the challenges arising from occlusions and non-Lambertian surfaces, we can only observe this effect here — its application is left for future work.

6. Manufacturing the Mirror

One of the major motivations for our automated mirror design pipeline was the widespread availability of computer-controlled rapid-manufacturing systems capable of high precision and low fabrication cost. The flexibility of these systems permits relatively few constraints in the design optimization process, allowing any final design to be turned into a physical prototype within a few hours. We have explored two manufacturing strategies, offering different tradeoffs between overall geometric accuracy and surface finish quality.

6.1. Plated brass

We used a milling machine under computer numerical control (CNC) to mill the surface geometry into a solid brass block (Figure 13). Because this process leaves residual grooves along the path the drilling tool took, machine-assisted polishing is used to create planar surfaces. Bright nickel plating then creates a thin coating on the surface, providing a durable mirror finish (Figure 13). The mirror in the illustration was approximately $20 \times 13 \times 5$ cm, with 11×7 facets. Including raw materials, milling, polishing, and plating, the total build cost was approximately USD 750.

The clear advantages of this approach are its reliance on standardized rapid manufacturing processes to achieve a precise alignment. Indeed, we observe that the facets are oriented correctly, to a high degree of accuracy. In addition, manufacturing cost is essentially independent of the number of facets (depending almost entirely on the volume and area of the piece), leading to practical fabrication of mirrors with many dozens of facets.

A major drawback of this process, however, is that it is difficult to make the facets perfectly flat. Marks left by the milling tool remain visible even with a substantial amount of polishing. In addition, the polishing process smooths out the corners between facets, leading to an unusable area around the edges of each facet.

These tradeoffs are illustrated in Figure 14, which shows the reflection of a checkerboard target in the manufactured mirror described above. The precise alignment of the marked red checkerboard squares in the rectified images shows that the facets are pointed in the expected direction. However, some waviness is visible when zooming into the image, and there is substantial distortion (caused by corner smoothing) around the edges of each facet.

As in this design, the geometry of the mirror is very precisely known, it itself may be used to calibrate the parameters of the capturing camera by aligning a rendered version of the mirror geometry against the observed picture.

6.2. Individual stainless steel facets

An alternative manufacturing process is to separate the construction of the mirroring surface of each facet from its alignment. In this process (Figure 1) we mill a support structure from acrylic, then glue on individual facets that are cut from a sheet of mirror-polished stainless steel.

One major advantage of this design is lower manufacturing cost. The plating process, with its high setup cost, is avoided, and the support structure may be milled from inexpensive acrylic rather than brass. While both approaches scale with surface area, as the surface area needs to be milled out of the support structure, the milling process for the acrylic is faster and thus less expensive, as the acrylic is ablated more easily, and the surface can be cut with a smaller tool path density (residual grooves would be covered by the steel facets). Cutting the facets is more efficient, as the cutting time and cost scale with the length of the edges as opposed to surface area.

One disadvantage is that, after gluing the facets on the mirror support, their alignment is less precise than in the electroplated

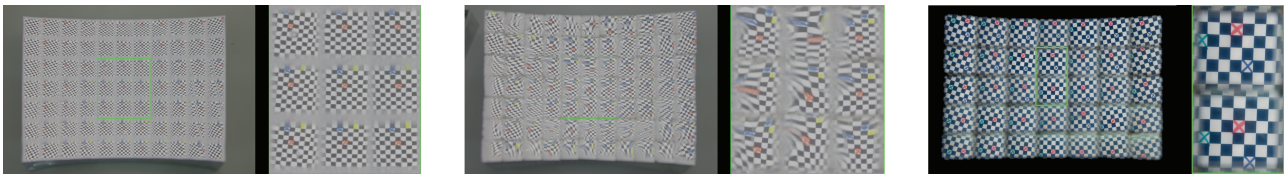

Figure 14: *Comparison of manufacturing alternatives. Left: Electroplated brass mirror as seen by the camera and a magnification of the centre region: several surface sections show blemishes which are reminders of the tool path that were not polished out, and artefacts introduced in the coating process. Centre: Polished steel plates, manually cut out of a 0.762-mm-thick steel sheet. This is a fast and inexpensive means of producing the mirror, but the distortions induced by mechanical stress during the cutting process render this particular configuration unusable. Right: Precise laser cutting avoids these distortions.*

Figure 15: *Refocusing results for light fields recorded with a laser cut mirror setup after upsampling to 1617 images (brightness/contrast manually adjusted after recording). The top row shows the result of simulating different aperture shapes for artistic bokeh control, the bottom row demonstrates refocusing for a short-exposure recording of liquid splashing off a spoon. The rightmost column shows refocusing using only the original 35 views (top) and the rectified input views (bottom). Please refer to the supplemental video for more examples.*

construction; we deal with that by having the user click on a single visible scene point in each segment and optimizing for individual rotation of the segments so that the corresponding rays converge at some point in space.

For cutting the facets, we investigated several technologies; CNC milling a 0.762-mm-thick steel sheet induced distortions (Figure 14) due to mechanical stress. For a 1.524-mm-thick sheet, they were reduced; we achieved the best results with laser cut facets, as seen in Figure 1 and 14. They are mostly free of distortion, but the material used for them has suboptimal mirroring qualities, inducing a small loss of contrast which can be easily compensated by tonal curve edits in any photo processing program.

7. Discussion

Example recordings: We have experimented with several scenes recorded with the mirror design from Figure 1, which enables refocusing applications by means of projective texturing of a virtual plane, as demonstrated in Figure 15 and in the supplemental video. With only 35 input views, gaps can be seen; however, the sampling is sufficiently dense to permit effective view interpolation with optical flow [ZBW11]; after upsampling to 1617 views, most gaps are closed. As we use a single sensor with globally synchronized shutter for recording, we can acquire quickly changing scenes with short exposures, such as a splash of liquid against a spoon.

Comparison to existing systems: The design choices of the proposed catadioptric setup for light field recording set it apart from existing techniques as follows: in contrast to a *moving camera*, it enables single exposure recording. This is also possible with *camera arrays*, which have the additional benefit of being able to extend the number of views with almost linear cost by adding more cameras. As an advantage, though, our approach does not require synchronization, and, the mirror being a compact, single piece of equipment, it can be calibrated globally.

This is also the main difference to previous *assemblies of mirror segments*: as a consequence of the use of precise machining tools during the construction, the position of the mirror facets is either an intrinsic result of the construction process (in the plated brass version), or easy to do by hand (when aligning pre-cut mirror facets on acrylic support).

In addition, we make use of an integrated design pipeline, which relieves the user from manually constructing mirror geometry, and supports the setup construction throughout — from the definition of design goals over visualizations of mirror shapes and results up to the final construction.

Camera-internal modifications with *refractive elements* are more compact and, once assembled, easier to transport. They are limited to applications that can make do with a small range of observed viewing angles. While they enable, for instance, refocusing, they are limited by the entrance pupil of the camera — our method can provide solutions for much more diverse applications.

Limitations: While the proposed method automatically navigates a rich design space of mirror geometries, it is restricted to planar mirror facets. This creates views for virtual cameras which are easy to model and render with, but as they effectively share the field of view of the recording camera among them, their individual fields of view are limited as the number of desired views increases.

Future work: Most beneficial extensions to the presented research will include broadening the space of mirror shapes that is considered in the optimization process. Extensions to non-planar facet geometry could do away with the minimal focal lengths in the virtual views, as convex facets would facilitate almost arbitrarily wide fields of view. With irregular mirror geometry, different sampling densities may be achievable which could make super-resolution possible.

The manufacturing process can be further improved to deliver higher precision in the local surface geometry, and mass production techniques may enable the inexpensive duplication of the same mirror shape from a single mold.

In conclusion, we have presented a comprehensive method for mirror-based light field recording which leverages algorithmic control throughout the entire pipeline from design to construction on the path towards true post-hoc-photography.

Acknowledgements

We thank Lawrence McIntyre for his support in the construction of the prototypes, Nicholas Johnson for fruitful discussion regarding the manufacturing processes, Yiming Liu for his help with some of the setup, Anette Müller for her support while recording the result pictures, Matt Pharr and Greg Humphreys for the PBRT software, Paolo Cignoni and colleagues for MeshLab[CC], Lena Gieseke for contributing illustrations, and Helen Tauc for the voice-over of the video. This work was supported by a fellowship within the Postdoc-Programme of the German Academic Exchange Service (DAAD) and US NSF grant #CCF-1012147.

References

[AVR10] Agrawal A., Veeraraghavan A., Raskar R.: Reinterpretable imager: Towards variable post capture space, angle and time resolution in photography. *Computer Graphics Forum (Proc. Eurographics special issue) 29* (2010), 763–772.

[CC] Cignoni P., Colleauges: MeshLab. http://meshlab.sourceforge.net/, Visual Computing Lab -ISTI -CNR.

[CNR08] Cossairt O., Nayar S. K., Ramamoorthi R.: Light field transfer: Global illumination between real and synthetic objects. *ACM Transactions on Graphics (Proc. SIGGRAPH special issue) 27*, 3 (August 2008), 57.

[Dan01] Dana K. J.: BRDF/BTF measurement device. In *ICCV Proceedings of Eighth IEEE International Conference on Computer Vision* (Vancouver, British Columbia, July 2001), vol. *2*, IEEE, pp. 460–6.

[FKR12] Fuchs M., Kächele M., Rusinkiewicz S.: Design and fabrication of faceted mirror arrays for light field capture. In *VMV 2012: Vision, Modeling & Visualization*, Goesele M., Grosch T., Preim B., Theisel H., Toennies K., (Eds.), Eurographics Association, Magdeburg, Germany (2012), pp. 1–8.

[GGSC96] Gortler S. J., Grzeszczuk R., Szeliski R., Cohen M. F.: The lumigraph. In *SIGGRAPH '96: Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 1996), ACM Press, pp. 43–54.

[GHAO10] Ghosh A., Heidrich W., Achutha S., O'Toole M.: A basis illumination approach to BRDF measurement. *International Journal of Computer Vision 90* (2010), 183–197. doi: 10.1007/s11263-008-0151-7.

[GZC*06] Georgiev T., Zheng K. C., Curless B., Salesin D., Nayar S., Intwala C.: Spatio-angular resolution trade-offs in integral photography. In *Rendering Techniques 2006 (Proc. EGSR)* (Nicosia, Cyprus, 2006), Eurographics Association, pp. 263–272.

[IMG00] Isaksen A., McMillan L., Gortler S. J.: Dynamically reparameterized light fields. In *SIGGRAPH '00: Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 2000), ACM Press/Addison-Wesley Publishing Co., pp. 297–306.

[IWH10] Ihrke I., Wetzstein G., Heidrich W.: A theory of plenoptic multiplexing. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conf.* (San Francisco, CA, USA, 2010), IEEE, pp. 483–490.

[LCV*04] Levoy M., Chen B., Vaish V., Horowitz M., McDowall I., Bolas M.: Synthetic aperture confocal imaging. *ACM Transactions on Graphics (Proc. SIGGRAPH special issue), 23* (August 2004), 825–834.

[LCWT06] Lanman D., Crispell D., Wachs M., Taubin G.: Spherical catadioptric arrays: Construction, multi-view geometry, and calibration. In *Third International Symposium on 3D Data Processing, Visualization and Transmission* (Chapel Hill, NC, USA, 2006).

[LH96] Levoy M., Hanrahan P.: Light field rendering. In *SIGGRAPH '96: Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 1996), ACM Press, pp. 31–42.

[LHG*09] Levin A., Hasinoff S. W., Green P., Durand F., Freeman W. T.: 4D frequency analysis of computational cameras for depth of field extension. *ACM Transactions on Graphics (Proc. SIGGRAPH special issue) 28* (July 2009), 97:1–97:14.

[Lou04] Lourakis M.: levmar: Levenberg-marquardt non-linear least squares algorithms in C/C++, July 2004. http://www.ics.forth.gr/∼lourakis/levmar/ (accessed on 15 Jan. 2011. Version used: 2.5).

[MP04] Matusik W., Pfister H.: 3D TV: a scalable system for real-time acquisition, transmission, and autostereoscopic display of dynamic scenes. *ACM Transactions on Graphics (Proc. SIGGRAPH 2004) 23*, 3 (2004), 814–824.

[MSY07] Mukaigawa Y., Sumino K., Yagi Y.: High-speed measurement of brdf using an ellipsoidal mirror and a projector. In *Computer Vision and Pattern Recognition, 2007*. CVPR '07. IEEE Conference (2007), IEEE, pp. 1–8.

[MTK*11] Mukaigawa Y., Tagawa S., Kim J., Raskar R., Matsushita Y., Yagi Y.: Hemispherical confocal imaging using turtleback reflector. In *Computer Vision – ACCV 2010* (2011), Kimmel, R., Klette, R., Sugimoto, A. (Eds.), vol. 6492 of Lecture Notes in Computer Science, Springer, Berlin-Heidelberg, pp. 336–349.

[Ng06] Ng R.: *Digital Light Field Photography*. PhD thesis, Stanford University, 2006.

[NLB*05] Ng R., Levoy M., Brédif M., Duval G., Horowitz M., Hanrahan P.: *Light Field Photography with a Hand-held Plenoptic Camera*. Tech Report, Stanford University Computer Science, April 2005.

[PHD06] Peers P., Hawkins T., Debevec P.: A Reflective Light Stage. *ICT Technical Report ICT-TR-04.2006*, 2006.

[TARV10] Taguchi Y., Agrawal A., Ramalingam S., Veeraraghavan A.: Axial light field for curved mirrors: Reflect your perspective, widen your view. In *Computer Vision and Pattern Recognition (CVPR)*, 2010 IEEE Conf. (San Francisco, CA, USA, 2010), IEEE, pp. 499–506.

[TAV*10] Taguchi Y., Agrawal A., Veeraraghavan A., Ramalingam S., Raskar R.: Axial-cones: Modeling spherical catadioptric cameras for wide-angle light field rendering. *ACM Transactions on Graphics (Proc. of SIGGRAPH Asia special issue) 29*, 6 (2010), 172.

[UWH*03] Unger J., Wenger A., Hawkins T., Gardner A., Debevec P.: Capturing and rendering with incident light fields. In *Rendering Techniques 2003 (Proc. EGSR)* (Leuven, Belgium, June 2003), Eurographics Association, pp. 141–149.

[VAR*08] Veeraraghavan A., Agrawal A., Raskar R., Mohan A., Tumblin J.: Non-refractive modulators for encoding and capturing scene appearance and depth. In *Computer Vision and Pattern Recognition, 2008*. CVPR 2008. IEEE Conference (Anchorage, AK, USA, 2008), IEEE, pp. 1–8.

[VRA*07] Veeraraghavan A., Raskar R., Agrawal A., Mohan A., Tumblin J.: Dappled photography: mask enhanced cameras for heterodyned light fields and coded aperture refocusing. *ACM Transactions on Graphics (Proc. SIGGRAPH 2007) 26*, 3 (2007), 69.

[WAA*00] Wood D. N., Azuma D. I., Aldinger K., Curless B., Duchamp T., Salesin D. H., Stuetzle W.: Surface light fields for 3D photography. In *SIGGRAPH '00: Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 2000), ACM Press/Addison-Wesley Publishing Co., pp. 287–296.

[War92] Ward G. J.: Measuring and modeling anisotropic reflection. In *SIGGRAPH '92: Proceedings of the 19th Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 1992), ACM, pp. 265–272.

[WIH10] Wetzstein G., Ihrke I., Heidrich W.: Sensor saturation in Fourier multiplexed imaging. In *Computer Vision and Pattern Recognition (CVPR)*, 2010 IEEE Conf. (San Francisco, CA, USA, 2010), IEEE, pp. 545–552.

[WJV*05] Wilburn B., Joshi N., Vaish V., Talvala E.-V., Antunez E., Barth A., Adams A., Horowitz M., Levoy M.: High performance imaging using large camera arrays. *ACM Transactions on Graphics 24* (July 2005), 765–776.

[Yan00] Yang J. C.: *A Light Field Camera for Image Based Rendering*. PhD thesis, Massachusetts Institute of Technology, 2000.

[YEBM02] Yang J. C., Everett M., Buehler C., McMillan L.: A real-time distributed light field camera. *Rendering Techniques 2002 (Proc. Eurographics Rendering Workshop)* (2002), 77–86.

[ZBW11] Zimmer H., Bruhn A., Weickert J.: Optic flow in harmony. *International Journal of Computer Vision 93*, 3 (2011), 368–388.

Supporting Information

Additional Supporting Information may be found in the online version of this article at the publisher's web site:

Video S1: Supplemental video.