Optical recipes for light field microscopes

Stanford Computer Graphics Laboratory Technical Memo 2006-001

Marc Levoy
Computer Science Department
Stanford University

June 20, 2006
(revised June 28 and August 9)

Introduction

In our SIGGRAPH 2006 paper "Light Field Microscopy" [Levoy 2006], figure 2 gives the optical layout of our prototype. While this figure and its accompanying caption are not incorrect, they also do not provide enough detail to set up and adjust a light field microscope. In particular, since microscope objectives do not contain focusing reticles, is not clear what is meant by positioning the camera sensor so that "each microlens records an in-focus image of the objective." In this brief technical memo, we show a more detailed optical layout, with emphasis on how to vertically position and focus each element in the optical train. We show this first for microscopes having non-infinity-corrected optics (as in our SIGGRAPH paper). We then describe an arrangement suitable for microscopes employing infinity-corrected optics, which is the more common case.

The figures on the succeeding two pages are standard thin-lens ray diagrams, drawn more or less to scale, but with the customary gap between the first and second principal planes of each lens or lens subsystem omitted, thereby reducing each to a single plane. Although the horizontal dimensions in these diagrams are exaggerated relative to a real microscope, the vertical distances between elements are scaled correctly relative to one another, in accordance with the dimensions shown on the right side of each figure. Moreover, to make the two optical arrangements more easily comparable, the tube lengths f_2 have been set equal. This is nearly literally true for Zeiss microscopes. In addition, the widths of each microlens, their focal lengths f_3, and the exit pupil sizes (equal to the open diameter of the telecentric stop) and focal lengths f_1 of the objectives have been set equal.

Regarding the unusual convention of drawing colored bundles of rays, on each page figures (a) and (c) show bundles that converge to points of the specimen and intermediate image plane, while figures (b) and (d) show bundles that converge to points on the telecentric stop and the sensor. Thus, figures (a) and (c) may be compared to drawings of the "field conjugates" in a standard textbook on microscopy (for example [Inoué 1997, p. 24-25]), and figures (b) and (d) to drawings of the "aperture conjugates". In keeping with this convention, the dashed lines in figures (a) and (b) represent the same ray within the diagrams on that page, as do the dashed lines in figures (c) and (d). Of course, these rays are different between the two pages.

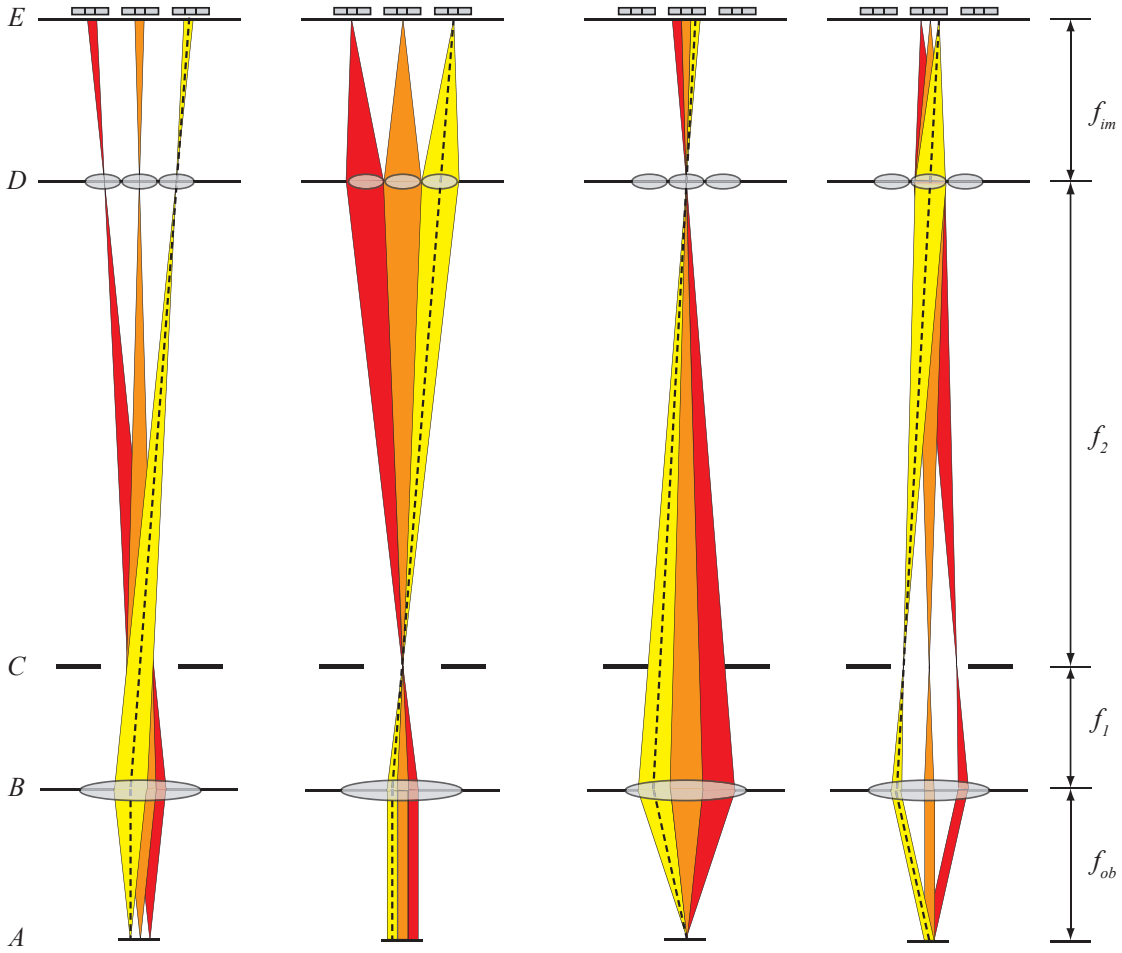

(a) central pixel and two neighbors, axial view, field conjugate

(b) same as (a), but aperture (and light field plane) conjugate

(c) central pixel, axial and two oblique views, field conjugate

(d) same as (c), but aperture (and light field plane) conjugate

Case 1: non-infinity-corrected optics. The specimen A is imaged by an objective B through a telecentric stop C onto intermediate image plane D. A microlens array placed at D forms multiple subimages on sensor plane E. Three such subimages are shown, represented diagrammatically by three non-abutting rectangles. Each subimage contains three pixels, denoted by divisions within the rectangles. The gap between these subimages is made larger than is recommended by our SIGGRAPH paper, to keep the figure uncluttered.

If the microscope is focused normally, the intermediate image plane lies $f_1 + f_2$ behind the second principal plane of the objective, where f_2 is the tube length (typically 160mm) and f_1 is the focal length of the objective. The telecentric stop is placed f_1 behind this second plane, thereby making the cones connecting A and B in figure (a) (or equivalently the shafts in figure (b)) parallel, which in turn makes the microscope orthographic. The distance f_{ob} from the specimen to the first principal plane of the objective is

$$1 \,/ \left(\frac{1}{f_1 + f_2} - \frac{1}{f_1} \right)$$

If the microlenses are positioned so that they form on the sensor plane a focused image of the telecentric stop as shown in (b) and (d), or of any conjugate of it, such as the condenser diaphragm if the microscope is adjusted for Köhler illumination, then the distance f_{im} from the microlenses to the sensor is

$$1 \,/ \left(\frac{1}{f_2} - \frac{1}{f_3} \right)$$

where f_3 is the focal length of a microlens. In this configuration, the central pixel in each subimage will be axial views of the specimen as shown in (a) and (b), and the peripheral pixels in each subimage will be oblique views as shown in (c) and (d).

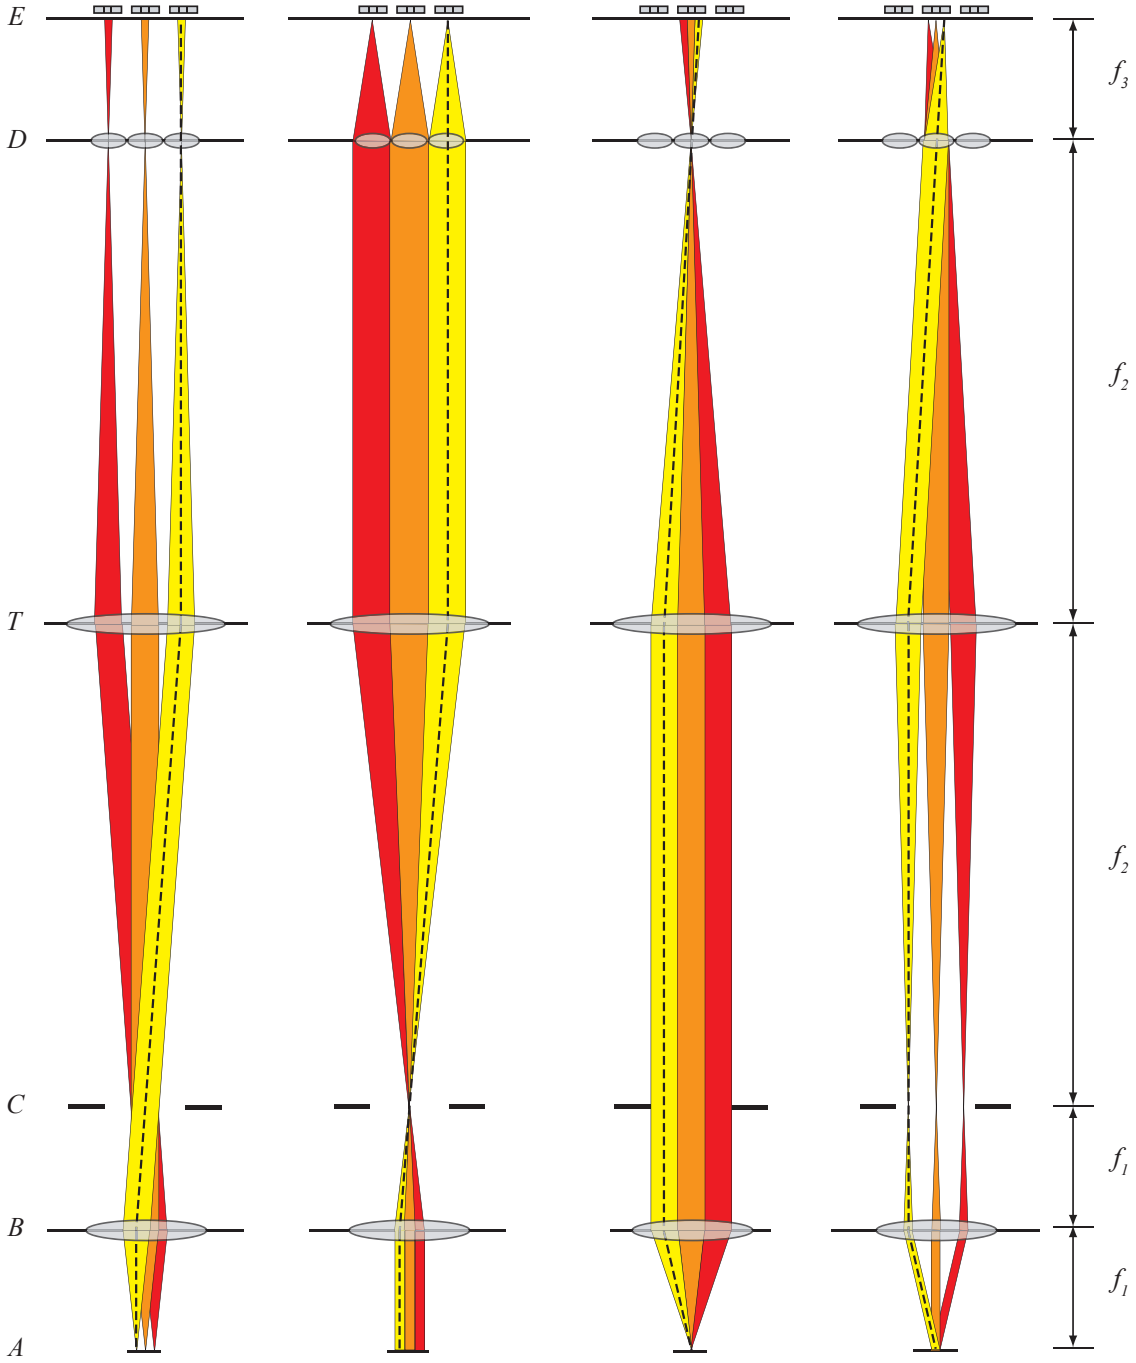

(a) central pixel and two neighbors, axial view, field conjugate

(b) same as (a), but aperture (and light field plane) conjugate

(c) central pixel, axial and two oblique views, field conjugate

(d) same as (c), but aperture (and light field plane) conjugate

Case 2: infinity-corrected optics. The specimen A is imaged by an objective B through a telecentric stop C and a tube lens T onto intermediate image plane D. A microlens array placed at D forms multiple subimages on sensor plane E. Three such subimages are shown, as before.

If the microscope is focused normally, the distance from the specimen to the first principal plane of the objective, and from its second principal plane to the telecentric stop, are both equal to the focal length f_1 of the objective. In addition, the distance from the second principal plane of the tube lens to the intermediate image plane is f_2, the focal length of the tube lens (also called the tube length in an infinity-corrected microscope). These two constraints cause the rays leaving each point on the specimen to form parallel bundles between B to T as shown in (a) and (c). This is the so-called "infinity section" of

the microscope. Its length is normally $f_1 + f_2$, although this length can be modified if care is taken not to cause vignetting at the tube lens. More will be said in the next section about the effect of changing this length.

If the microlenses are positioned so that they form on the sensor plane a focused image of the telecentric stop as shown in (b) and (d), or of any conjugate of it as before, then the distance from the microlenses to the sensor is f_3, the focal length of a microlens. It is worth noting that this distance is slightly shorter than f_{im} in the non-infinity case. Also, the bundles of rays impinging on the sensor plane from adjacent points on the specimen are parallel and vertical in the infinity case (see figures (a) and (b)), but angled in the non-infinity case. These two factors make the subimages slightly smaller in the infinity case, and the gaps between them smaller, than in the non-infinity case.

In practice these differences are very small - much smaller than they appear in the horizontally exaggerated diagrams in this memo. Moreover, the fundamental design rule described in our SIGGRAPH paper - of matching the F-number of the microlenses to the image-side F-number of the microscope objective, still holds. On the other hand, since the chief rays of these subimages are vertical in the infinity case, they may suffer less from aberrations. Also, the consistent alignment between microlenses and pixels in the infinity case suggests that permanently bonding a microlens array to a sensor may be more practical in this case.

Microscopes with focusing nosepieces

Some infinity-corrected microscopes have a focusing nosepiece, meaning that the objective moves up and down, rather than the stage, but the tube lens does not move. Among Nikon's products, their physiology microscopes (e.g. the FN1) and inverted microscopes (e.g. the TE2000) employ this design. On these microscopes, refocusing the microscope is tantamount to changing the length of the infinity section (as described earlier), although in this case only slightly.

The effect of changing the length of this infinity section is to move outward or inward the lateral positions of the beams from the peripheral pixels (red and yellow beams in figures (a) and (b)) where they pass through the tube lens. Following these changes upward through the optical train, these peripheral beams will thereby be rotated around fixed points on the intermediate image plane. Because these points are fixed, the lateral magnification of the image would not change if this were an ordinary microscope. However, in a light field microscope these beams continue upward to the sensor plane. On that plane the location and width of the subimages will change.

Fortunately, these changes are very small indeed. For example, for a 40× objective with f_1 = 4mm, f_2 = 160mm, f_3 = 2.5mm, and a field half-width W of 10mm at the intermediate image plane, if we refocus the microscope by 1 micron in Z, thereby moving the exit pupil at B by 1 micron in Z, this will cause a lateral shift of the centroid of the yellow beam at T of 1 micron / (f_2/W) = 1/16 micron. This will in turn cause a lateral shift (in the other direction) of the centroid of the yellow beam when it reaches E of 1/16 micron / (f2/f3) = 1 nanometer! Even moving the objective 100 microns in Z will shift the subimages laterally by only 100 nanometers. This is two orders of magnitude smaller than the size of a typical sensor pixel.

Discussion

As in any complex optical train, it is sometimes difficult to determine from a ray diagram which element limits the widths of ray bundles. In the foregoing diagrams, in figures (a) and (c) the width of each bundle is limited by the spatial extent of a pixel on the sensor, while in figures (b) and (d) their width is limited by the diameter of a microlens. Similarly, in figures (a) and (b) the width of the entire collection of ray bundles is limited by the width of the microlens array (or sensor if that is smaller) - three microlenses in this toy example, while in figures (c) and (d) the corresponding limit is the diameter of the telecentric stop. These limits fit our intuition that the spatial extent of a microscope image is determined by constraints at one or more field conjugates, while the angular extent is determined by constraints at one or more aperture conjugates.

Regarding application of this optical arrangement as a light field microscope, we remind the reader briefly of a few points made in our SIGGRAPH paper. First, the spatial resolution of the light field, hence the number of pixels in images computed from it, is determined by the number of microlenses - three in this example. Second, the angular resolution of the light field, i.e. the number of unique oblique views that can be computed from it, is determined by the number of resolvable spots in each microlens subimage - three in this example. As proved in our paper, the latter number also gives the number of slices with non-overlapping depths of field in focal stacks computed from the light field. Third, although it is not obvious from the foregoing figures, all rays that leave the specimen and are not vignetted by the objective or tube lens are indeed captured by the microlens array and sensor. In other words, aside from a negligible amount of scattering from light that strikes the seams between adjacent microlenses, no light is lost in the conversion from a standard microscope to a light field microscope.

Since publication of our SIGGRAPH paper, people have asked us whether light fields recorded by our microscope, because they are discretely sampled, contain aliasing. Yes they do. The property (discussed above) of not losing any rays is necessary but not sufficient to prevent aliasing in a sampled light field. In our situation, such aliasing may occur either because insufficient pre-filtering is applied before partitioning rays among microlenses, or because insufficient pre-filtering is applied before partitioning rays among pixels on the sensor. The latter type of aliasing is easily ameliorated by employing a camera with high pixel resolution, an anti-aliasing filter in front of the sensor, and less-than-perfect focusing in the camera's relay lens system if there is one. The former type of aliasing could potentially be ameliorated by placing a diffuser in front of the microlens array; however, we have not done this, for fear of losing resolution in all four dimensions of the light field.

The presence of aliasing also suggests that super-resolution methods may be applicable, for example by shifting the microlens array by fractions of a microlens width and capturing multiple light fields. This may provide a path for recovering some of the spatial resolution lost by introducing the microlens array in the first place. However, our ability to increase spatial resolution while retaining angular resolution is limited by the amount of information contained in the light field, which is in turn limited by diffraction. Thus, these techniques are likely to work well mainly for microscope objectives having low magnification and high numerical aperture.

References

[Inoué 1997] Inoue, S. Spring, K.R., *Video Microscopy*,

[Levoy 2006] Levoy, M., Ng, R., Adams, A., Footer, M., Horowitz, M., ‘‘Light field microscopy,’’ *ACM Transactions on Graphics* (Proc. SIGGRAPH 2006), Vol. 25, No. 3, 2006.