BCAM Calibration

David Daniels*, Kevan Hashemi, Jim Bensinger
Brandeis University

September 15, 2000

Abstract

The Brandeis CCD Angle Monitor (BCAM) has been selected by the ATLAS collaboration to perform the polar alignment of the endcap muon system. This note describes the procedure for obtaining the nine calibration parameters which characterize any BCAM. This work was performed at Brandeis University.

1 Description of Prototype BCAM Camera

The operation of the BCAM camera is described in Atlas Note "The BCAM Camera"[1]. It is essentially a lens–CCD camera with the addition of an infrared (IR) filter in front of the lens, an aperture immediately behind the lens, and two (front-surface-reflecting) mirrors between the lens and CCD that bend the light path back upon itself to keep the overall package smaller (see Figure 1).

Figure 1: These are pictures of a BCAM with its cover (and IR filter) taken off and electronics removed. In the left picture, the light proceeds through the lens at the lower right, reflects off the mirrors at the top, and is recorded by the CCD at the bottom left. In the right picture, one may see the aperture behind the lens at the upper left, the two mirrors at the lower left and right, and the front of the CCD at the upper right. The screw holes on the optical table are placed at 1 inch (25.4 cm) intervals to give a sense of scale.

The BCAM is uniquely placed onto a set of three balls by a kinematic mount machined into the base of each camera. The mount consists of a conical shaped receptacle (cone), an elongated conical shaped

---

*Email: ddaniels@brandeis.edu
[1]ATLAS note in preparation.

groove (slot), and a flat milled area (plane) (see Figure 2). One ball sits in the cone hole and specifies a unique point about which the BCAM may revolve. The second ball rests in the slot, and together with the cone ball this specifies a unique axis about which the BCAM may still rotate. The third ball locates the camera on the axis of rotation, fixing it rigidly in space. A single screw provides the contact force holding the camera onto its three balls so that it may be mounted in any orientation (even upside-down). One advantage of the three ball kinematic mount is that a specific local coordinate system may be defined

Figure 2: The underside of a BCAM, showing the mounts for the three balls: cone at upper center, slot at lower left, and plane at lower right.

for each BCAM simply by measuring the positions of the three balls. It is therefore not important to place the balls precisely, but only to measure their locations precisely. The positions of the three balls can be measured by a micron-precision coordinate measuring machine (CMM), and this specifies the local coordinate system for each BCAM subsequently placed on them.

The BCAM local coordinate system may be defined in numerous ways, and the results of the calibration are insensitive to the particular definition. In this work, the BCAM coordinate system is defined as follows. The origin is placed at the center of the "cone ball". The $\hat{z}$ direction is initially defined as the direction from the middle of the "slot ball" to the origin. The middle of the "plane ball" lies in the xz plane, and the y axis completes a right-handed coordinate system. We then rotate these coordinate axes about the y axis by an angle approximately equal to one-half the angle that the balls make (rotation angle is $0.5596/2$ radians) so that the z axis runs roughly from the midpoint between the slot and plane balls towards the cone ball. This roughly corresponds to the direction in which the BCAM is viewing.

During the calibration several coordinate systems are used. The global coordinate system is the one into which all others are placed. It is roughly identified with the optical table on which the calibration apparatus is placed, and is therefore sometimes called the table coordinate system. The CMM measures ball centers in its own local coordinate system, and the BCAM which sits on the balls has yet another local coordinate system attached to it. Finally, the CCD which records the image positions has its own (2-dimensional) coordinate system.

The function of calibration is to define a particular set of parameters for each BCAM so that, given the positions of the centers of the three balls on which it sits in global coordinates, and the position of an image on the CCD in CCD local coordinates, one may construct a ray (point and direction) in global coordinates which points back toward the source.

2   Definition of Calibration Parameters

2.1   Simple Lens–CCD Camera Example

A simple lens–CCD camera schematic is shown in Figure 3. The pivot point (PP) of a camera is the point through which all light rays pass undeflected; in this case it corresponds to the center of the lens, $\vec{L}$. The line connecting the origin (center) of the CCD $\vec{C}$ and the pivot point $\vec{L}$ has a direction $\hat{n}$. The
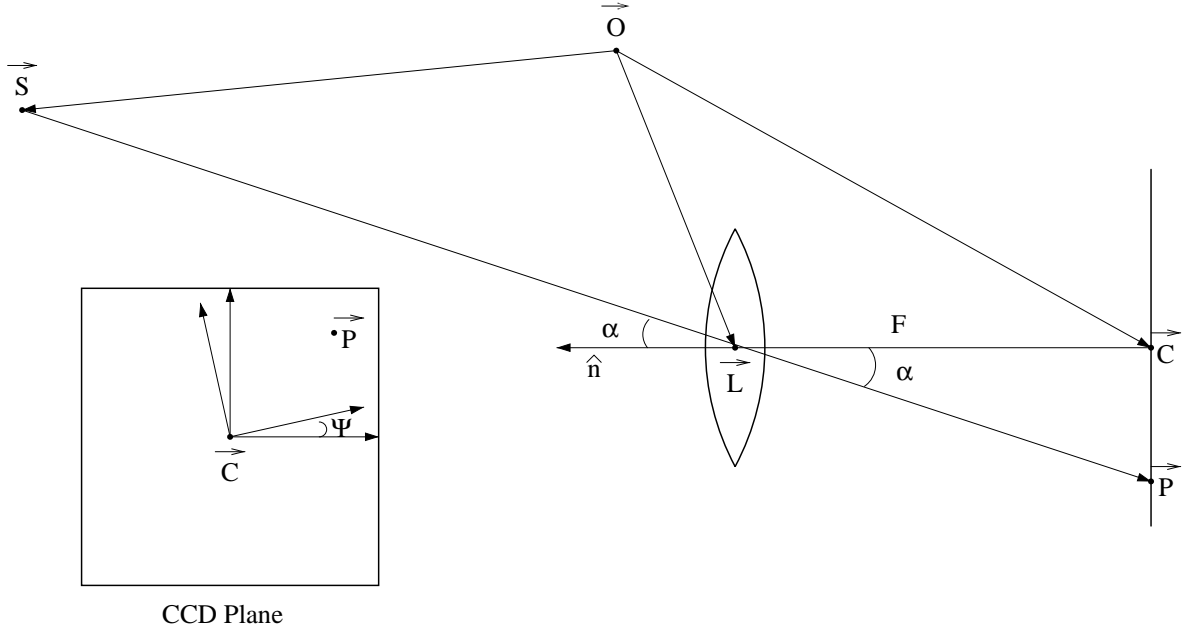
Figure 3: Schematic of a simple lens–CCD camera in an arbitrary coordinate system with origin $\vec{O}$.

distance between $\vec{L}$ and $\vec{C}$ is F. Light from a source at position $\vec{S}$ will form an image on the CCD at point $\vec{P}$. The CCD may be rotated by an angle Ψ about the BCAM z axis (the longitudinal axis). With these definitions, one may reconstruct the source bearing (ray pointing at $\vec{S}$) from a CCD image position $\vec{P}$ as follows. Rotating the image position in CCD coordinates about the longitudinal axis by $-\Psi$ and adding the result to the position $\vec{C}$ yields the image position in BCAM coordinates. The source $\vec{S}$ is then found by tracing the ray backward from the image point $\vec{P}$ through the pivot point $\vec{L}$.

A simple lens–CCD camera can thus produce a ray (point and direction) pointing towards a source given the image position and a suitable choice of seven parameters. In the example above we used the pivot point (L_x, L_y, L_z), a point on the CCD (C_x, C_y, C_z), and the angle Ψ. For the actual BCAM calibration we have chosen to represent the point $\vec{C}$ by the length F and the two direction cosines (t_x, t_y) which specify the direction $\hat{n}$, but this is equivalent to (C_x, C_y, C_z).

Additional elements in the optical description of the BCAM may affect the interpretation of these seven parameters, but they should not require any new parameters. For example, the addition of an aperture changes both the pivot point position and the CCD center, but the new system can still be described by a set of seven parameters. (See Appendix I.) It is important to keep in mind, therefore, that these parameters do not correspond directly to the positions of physical objects, i.e. "the position of the CCD plane" becomes a euphemism for the set of parameters (F, t_x, t_y) and does necessarily correspond to the physical location of the CCD.

2.2 Cylindrical Lensing Effects

The two mirrors should have no effect either, as they just bend the light path around in a 'U' shape. Unfortunately, the mirrors were not perfectly flat (their specifications state that they are flat to within four wavelengths over their surface of approximately 1.56×1.06 mm), which caused them to act as very long focal-length lenses. Since we cannot assume that the effect is purely spherical in nature, we model each mirror as a lens with spherical and cylindrical components, with the axis of the cylindrical component unknown. The combination of two spherical lenses is optically equivalent to a third single spherical lens, and two cylindrical lenses are optically equivalent to a spherical lens plus a cylindrical lens, so we can model the effect of the BCAM main lens plus the two mirrors as a single spherical lens plus a single cylindrical lens. The spherical lens is already accounted for in the seven calibration parameters, and the extra cylindrical component adds two more parameters.

3

The effect of a cylindrical lens is to move points on the CCD either closer together or further apart, depending on whether the cylindrical lens is focusing or defocusing. A focusing lens, however, is indistinguishable from a defocusing lens turned at right angles plus an additional overall scaling factor. The ambiguity is resolved by requiring the two additional BCAM parameters to describe a *defocusing* cylindrical lens with the overall scaling factor, if any, accounted for by a redefinition of the main lens focal length.

To first order, the displacement of the image on the CCD due to cylindrical distortion is linear in the perpendicular distance to the cylindrical axis. We therefore parameterize the cylindrical effect in terms of a two-dimensional vector $\vec{D}$ whose length represents the strength of the effect, and whose direction on the CCD dictates the direction of expansion (i.e. the correction to each image with position $\vec{T}$ is $\left(\vec{T} \cdot \vec{D} \right)$ in the direction of $\vec{D}$, or $\left(\vec{T} \cdot \vec{D} \right) \hat{D}$).

The BCAM can thus be parameterized by nine constants: the position of the pivot point (P_x, P_y, P_z), the distance F and direction cosines (t_x, t_y) of the vector from the CCD center to the pivot point, the angle ψ that the CCD makes about the z-axis, and the cylindrical distortion vector (D_x, D_y).

The IR filter can also contribute to the interpretation of the calibration parameters. If the glass were perfectly flat, it would simply move the pivot point slightly closer to the filter because of refraction. Planar non-uniformities (wedge angles in the glass, e.g.) could cause the pivot point to move in all three directions as well as affecting the values of the CCD plane parameters. If the glass has lensing effects similar to the mirrors then it could also contribute to the cylindrical distortion vector. Thus, even if the mirrors were removed from the BCAM design, the cylindrical parameters in principle still could be non-zero.

3   Calibration Apparatus

3.1   The BCAM Roll Cage

The BCAM was mounted on a set of three balls that were set into a brass roll cage (see Figure 4). The cage itself had cone–slot–plane (CSP) receptacles on four of its outer sides so that it could be mounted in four different orientations, corresponding to rotations about the BCAM longitudinal axis (see Figure 5). The roll cage was mounted on a brass base plate (see Figure 4) containing three mounting balls and two reference balls, and the base plate was fixed to an optical table (see Figure 6).

3.2   The Source Block

The source block was a 10.5×10×9.5 cm aluminum piece with 8 holes drilled into its front face, arranged in concentric rectangles (see Figure 7). The outer holes were not all visible by the BCAM at the ranges used for calibration, so only the inner four holes were instrumented with LED's. The source holes were conical in shape, with the "small" end facing the BCAM to limit reflections off the walls of the hole. The small ends of the inner four openings facing the BCAM had 2-mm diameters. Directly behind the face plate of the source block was a layer of opal glass diffuser, which was held in place by a backing plate bolted onto the face plate (see Figure 8). The backing plate also held the light sources themselves. The light sources were the HSDL4230 LED's from Hewlitt-Packard, which transmit 35 mW at 875 nm. The range of a source was defined as the distance from the BCAM pivot point to the front surface (small end) of the conical hole.

3.3   The Optical Table

The BCAM calibration measurements were performed on an optical table approximately 3.6m long by 1.5m wide (see Figure 6). After fixing the roll cage and placing the source block far enough away from the BCAMs to permit all sources to be in view in all four BCAM orientations, we were left with a usable dynamic range of 1.0–2.7m.

Figure 4: Brass roll cage (right) and base (left) being measured by the coordinate measuring machine (CMM). The BCAM views through a hole in the end of the roll cage which can be seen clearly in this picture.

A straightedge was clamped to the surface of the table, and the roll cage base was pressed up against this for location. The source block was then slid up and down along the straightedge to obtain different ranges.

The global (table) coordinate system was defined with respect to the roll cage base, since this could be measured by the CMM and placed on the optical table repeatably. The origin of this system was the corner of the roll cage base touching the surface of the table and the straightedge and closest to the source block. The longitudinal axis from the camera toward the sources along the straightedge was the z axis; the y axis pointed toward the ceiling, and the x axis pointed across the table away from the straightedge, completing a right-handed coordinate system.

3.4   The Coordinate Measuring Machine

All precision mechanical measurements were made with our in-house coordinate measuring machine (CMM), seen in Figure 4. This is a Brown & Sharpe model "microxcel pfx 454" fitted with a Renishaw MIP probe head and ruby-tipped stylus to measure positions. Its stated accuracy is 7.6μm, with 2.2μm repeatability. The measuring surface is a granite slab, and the dynamic range of the instrument is 510mm×458mm×406mm (L×W×H).

4   Calibration Procedure

4.1   Calibrating the Source Block

In order to calibrate a BCAM, we must begin with a set of sources with known *relative* positions. We define the position of an LED source as the location of the centroid of the distribution of light which is
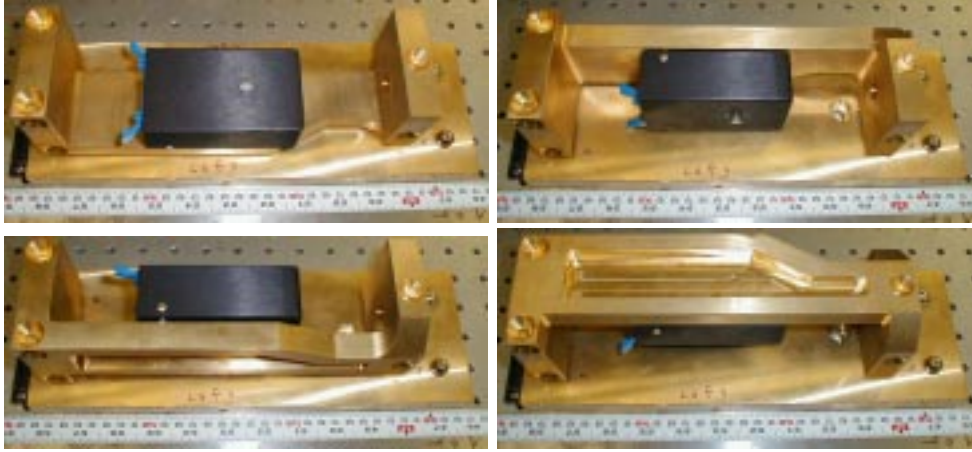
Figure 5: The BCAM mounted in the roll cage in its four orientations. Clockwise from upper left these are 0°, 90°, 180°, and 270°.

emitted by it. Since the light distribution may not be perfectly uniform, this position is not equivalent with e.g. the geometrical center of the exit hole in the face plate of the source block. We must therefore perform a relative calibration of the source positions before we can calibrate BCAMs.

4.1.1 Distances between sources

We begin by placing a single source on a micrometer stage at some set range and viewing it with an *uncalibrated* BCAM. As we move the source across the field of view of the BCAM in steps of known size, we get the image magnification at that range. Remounting the BCAM at 90° and repeating the process gives the magnification in the orthogonal direction on the BCAM's CCD. Recall that these magnifications may not be identical because of cylindrical distortions.

Replacing the single source and micrometer stage with the source block at the same range, we take images of the four sources we will use for BCAM calibration. Knowing the magnification of this BCAM (in both directions on the CCD) at that particular range yields the distances between all sources.

We used BCAM #44 to calibrate the source block. We placed the sources at a range of approximately 1.23m from the camera and moved them 10mm horizontally on a micrometer stage in steps of 0.5mm. With the BCAM in the 0° orientation (CCD x direction) the magnification was 0.126308±0.000006, and in the 90° orientation (CCD y direction) it was 0.125667±0.000005. Using these magnifications and the CCD spot positions, we can calculate the source positions shown in Table 1.

Source	x (mm)	y (mm)
5	17.842	14.923
6	9.824	14.912
7	9.842	8.921
8	17.869	8.912

Table 1: Source positions relative to a line from the upper left corner of the CCD through the pivot point of the BCAM, in a coordinate system with x and y axes parallel to the CCD's axes (i.e. there is an additional unknown global offset and rotation yet to be accounted for).

The errors on the relative source positions are the order of a micron, and these are dominated by the uncertainty in the micrometer stage calibration. The motion of the micrometer stage was calibrated with an interferometer, and was found to be accurate to 0.02%. If the distance between sources is on the order of 10mm, this gives a 2μm error on the distance between sources.

Figure 6: Picture of the optical table on which the calibration data were taken. A BCAM is mounted in the roll cage in the 0° orientation in the foreground, the source block is middle of the picture, and the laser for the interferometer may be seen at the upper left.

4.1.2   Source block rotation

We now know the distances between the four sources in the source block, but since the BCAM is uncalibrated, we do not know the absolute position of the array of sources nor their overall rotation about the longitudinal axis. There are two methods we could use for determining the source block rotation. We could either determine the rotation of the source block at each of the two ranges we use for BCAM calibration, or we could determine the apparent (false) rotation introduced by the BCAM calibration and subtract that afterwards. Both methods are described in this note, but only the latter method was used in the actual calibration.

Both methods involve the use of a flat granite bar on which are placed both the source block and the BCAM in its brass roll cage. Figure 9 shows a picture of the setup. Sliding a source across the bar in the BCAM field of view marks out a horizontal line (in global coordinates) on the CCD. For the first method, we may then replace the BCAM and source block on the optical table and take images of the four sources at each of the source block's ranges. Knowing the direction of the global horizontal (x axis) on the CCD, we may reconstruct the global rotation of the sources about the longitudinal axis at each range. These rotations should in general be very similar, but may not be the same if for instance the optical table is not perfectly flat.

For the second method, we first calibrate a single BCAM (described in Sections 4.2 and 4.3) assuming that there is no rotation of the source block relative to the BCAM. This will result in a bias in the CCD rotation angle parameter ψ, which can be measured and subtracted from all future calibrations. The advantage of this approach is that it accounts for all possible sources of rotation at once, by simply measuring their combined effect on the BCAM calibration.

Figure 7: Front view of the source block showing the rectangular array of conical source holes.
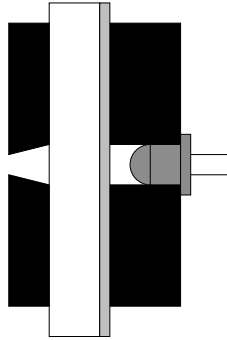

Figure 8: Cross section of one LED source in the source block. From right to left we have the LED fixed into a hole in the backing plate, the opal glass diffuser, and the source block face plate showing a conical hole. The BCAM views this source from the left.

Moving a source on the granite bar in 10 steps of 1–2mm at a range of approximately 2m yields an effective BCAM–source block rotation of 3.6mrad (after initial BCAM calibration). Using this angle and the distances calculated in Table 1, we can recalculate source positions in global coordinates such that they cancel this net rotation for all future BCAM calibrations (see Table 2). There is still an unknown transverse offset to the source block, but we can proceed with the BCAM calibration without this information.

4.2   Data Collection

If we know the relative positions of the sources to good precision, we can obtain the calibration parameters as follows. To get the cylindrical distortion vector, we look at the spacing between all four source images at a single range. For the CCD rotation angle ψ, we need the image positions for two sources at a single range. To get the longitudinal parameters P_z and F (see Section 2.2), we use two images at two different ranges. Finally, for the transverse parameters P_z, P_z, t_x, and t_y, we need a single source seen from two camera rotations (about the longitudinal axis $\hat{z}$). With nine unknown parameters, we need at least nine independent measurements per BCAM.

For each camera, we took (x,y) images of all four sources at two ranges and at four BCAM orientations, for a total of 64 measurements. This gives us more than enough data and allows for some redundancy in the calibration.

Figure 9: Picture of the granite bar and BCAM setup used to obtain the relative rotation of this uncalibrated BCAM about the longitudinal axis of the global coordinate system.

Source	x (mm)	y (mm)
5	54.568	72.920
6	46.550	72.880
7	46.589	66.889
8	54.616	66.909

Table 2: Source positions in global coordinates, up to a small transverse offset. Relative positions and overall rotation is correct.

The measurement procedure was performed by taking four images, one at a time, at the 0° BCAM orientation and "near" source range. The BCAM was then rotated to the 90° position in the roll cage and another set of four images was taken. After all four BCAM orientations were taken like this, the sources were moved backward to the "far" range, the BCAM was placed back in its 0° orientation, and the process was repeated.

4.3 Data Analysis Methods

Two software reconstruction methods were used on these data to obtain the calibration parameters for the BCAMs, and these programs served as checks on each other. The first calibration method used a series of steps to identify groups of calibration constants, and the second method used MINUIT to perform a multivariate χ^2 minimization in the space of the calibration parameters.

4.3.1 Analytical approach

The analytic approach is probably the more intuitive of the two reconstruction techniques used. Since groups of calibration parameters are obtained in semi-independent stages, each stage may be more easily debugged before moving on to the next. In addition, it is easier to see directly which data affect which calibration parameters using this approach.

In the first stage, the cylindrical distortion vector is obtained from the distances between image centroids on the CCD. One does not need to know the exact rotation of the BCAM nor the range used. For four positions, there are 6 pairs (Figure 10a), and though these are not independent, they may all be used to construct the cylindrical distortion vector. If the images form a rectangle, yet the sources are known to lie on the corners of a square, then there are two orthogonal possible distortions: one expanding,
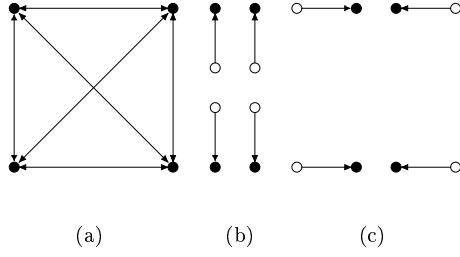
Figure 10: a) Lengths between image spots used for determining the cylindrical distortion vector. Orthogonal b) expansions and c) contractions are impossible to tell apart if the overall scale factor is not known a priori. Open circles represent undistorted spot positions, closed circles reprsented distorted positions.

the other contracting (Figure 10b,c). Since the magnification is unknown (because the pivot point position and PP-to-CCD distance have not yet been determined), either of these options is allowed. We arbitrarily choose to accept the expanding solution, shown in Figure 10b; if the distortion is actually contracting, this will appear as an expanding distortion in the orthogonal direction plus an additional overall demagnification. In addition, since the distortion moves the spot outward from the center of the CCD along the *unsigned* direction of the vector $\vec{D}$, both $\vec{D}$ and $-\vec{D}$ are adequate descriptors of the distortion. We (again arbitrarily) choose the solution with $D_x > 0$ so that we may have a unique solution.

In the analytical approach, a method was chosen that did not rely on knowing the relative source positions.[2] This decoupled the source block calibration (Section 4.1) from the cylindrical distortion vector determination. For any pair of spot centroids on the CCD $\vec{T_1}$ and $\vec{T_2}$, corresponding to two sources at a given BCAM orientation and source range, the vector between them could be called $\vec{V_1}$. In the absence of any cylindrical distortion, this vector would be the ideal $\vec{U_1}$, where $\vec{V_1} = \vec{U_1} + \left(\vec{U_1} \cdot \vec{D}\right)\hat{D}$. Given the observable $\vec{V_1}$, we can approximate $\vec{U_1}$ (assuming $\left|\vec{D}\right| \ll \left|\vec{U_1}\right| \approx \left|\vec{V_1}\right|$) by $\vec{V_1} - \left(\vec{V_1} \cdot \vec{D}\right)\hat{D}$. After a 90° BCAM rotation, the same two sources make a new vector, $\vec{V_2}$. Rotating $\vec{U_1}$ by 90° as well, we get $\vec{U_1'}$, and $\vec{U_1'} + \left(\vec{U_1'} \cdot \vec{D}\right)\hat{D} \equiv \vec{U_1''}$ should be the same as $\vec{V_2}$. The difference $\delta = \vec{U_1''} - \vec{V_2}$ can then be minimized using a successive approximation approach in order to solve for the components D_x and D_y.

Since there are 6 pairs of source images for each BCAM orientation, we can calculate δ differences for each of them for each pair of 90° BCAM rotations. There are 4 pairs of BCAM rotations (0°/90°, 90°/180°, 180°/270°, and 270°0°), so we have a total of 24 δ distances. (We only use the images from the source "near" range.) In the successive approximation, we minimize the sum of these 24 δ's, added in quadrature. The step size in the approximation was effectively 0.02% for each component of the vector $\vec{D}$, so this serves as a lower bound for the precision of this approach. The cylindrical distortion vector accounts for two of the nine calibration parameters.

With the cylindrical distortion obtained, we can get the rotation of the BCAM by comparing the orientation of the distortion-corrected images with the known orientation of the source block from the source block calibration. In the analytical approach, this rotation angle ψ is estimated from each pair of source images, and the results are averaged over all pairs. ψ represents the third BCAM calibration constant.

The longitudinal parameters P_z and F (Section 2.1) are derived by comparing two sources at the near and far ranges. Figure 11 shows a schematic of the setup. The objects, represented by the two large arrows in the figure, are the distances between a pair of sources at the near ($z_1 - P_z$) and far ($z_2 - P_z$) ranges. The magnification m is defined by the ratio of image "height" (h_i) to object "height" (h_o). By similar triangles, the magnification is also given by the ratio F/r where r is the range. The difference in

---

[2]For an alternate method of obtaining the cylindrical distortion vector using a single BCAM orientation and the relative source positions, see Appendix II.
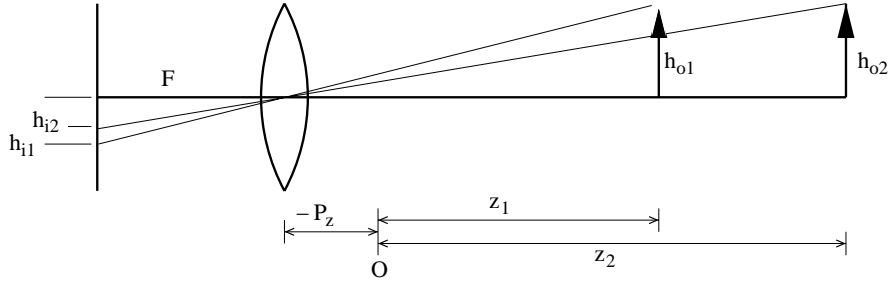
Figure 11: Schematic of the layout for determining the longitudinal calibration parameters. $\vec{O}$ represents the global origin, which is arbitrary.

ranges is then

$$r_2 - r_1 = z_2 - z_1 \equiv \Delta z = \frac{F}{m_2} - \frac{F}{m_1} = F \frac{m_1 - m_2}{m_1 m_2}$$

$$\Rightarrow F = \frac{m_1 m_2 \Delta z}{m_1 - m_2}. \tag{1}$$

Then, given that $m_1 = F/(z_1 - P_z)$, e.g., we could insert F to find P_z, and we would have two more BCAM calibration parameters.

We can also derive P_z independently of the source block calibration. Given that the ratio of image sizes equals the ratio of the ranges by similar triangles, we have

$$\frac{h_{i1}}{h_{i2}} = \frac{r_2}{r_1} = \frac{z_2 - P_z}{z_1 - P_z}$$

$$(z_1 - P_z) \frac{h_{i1}}{h_{i2}} = z_2 - P_z$$

$$P_z \left(1 - \frac{h_{i1}}{h_{i2}}\right) = z_2 - z_1 \frac{h_{i1}}{h_{i2}}$$

$$P_z = \frac{z_2 - z_1 \frac{h_{i1}}{h_{i2}}}{1 - \frac{h_{i1}}{h_{i2}}}.$$

The final four BCAM parameters are the transverse variables P_x, P_y, t_x, and t_y. They represent the transverse components of the pivot point and the CCD center, respectively. A central axis may be drawn through the CCD center and the pivot point, such that the position of an image on the CCD with respect to the CCD origin is proportional to the distance of the corresponding source from this central axis. This is similar to the optical axis of a lens.

If we rotate the camera about the longitudinal axis of the BCAM ($\hat{z}$) by 180°, the central axis will also flip. The difference in the CCD positions of a source viewed from both orientations will be proportional to the spread of the central axis at the source range before and after the flip (see Figure 12). Since we have determined the longitudinal parameters, we know the ranges r_1 and r_2, and the PP-CCD distance F. This gives us the magnification of the CCD images, which gives us the two separations of the axes (flipped and unflipped) at two ranges. Projecting back to the longitudinal positions of the pivot point and the CCD, we get the axis separations there. One-half of the separation represents the offset, and this is the parameter we are interested in. Examining the x and y components of the image point separately yields the x and y components of the PP and CCD parameters.

4.3.2 Global minimization method

The second method uses MINUIT to perform a multivariate χ^2-like minimization in the space of the calibration parameters. Given a set of BCAM parameters and source positions, one can trace the light path from source to CCD as follows. First, transform the source position from global coordinates into
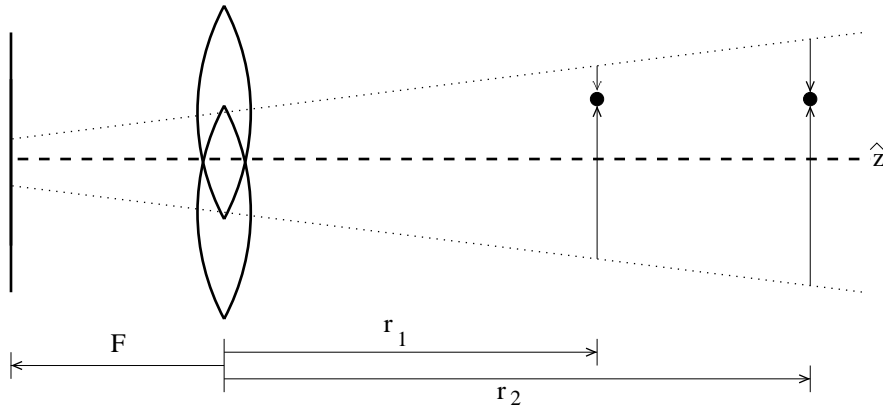
Figure 12: Schematic of the layout for determining the transverse calibration parameters.

BCAM local coordinates. Next, trace the ray straight from the source through the pivot point to the CCD. Transform that point from BCAM local coordinates into CCD local coordinates. If we designate this projected CCD point as $\vec{P_i}$ for source i and the actual BCAM reading for the same source as $\vec{B_i}$, then the function to be minimized is $\sum \left(\vec{P_i} - \vec{B_i} \right)^2$, where the sum is over all sources. This function should be minimized by the best-fit nine BCAM parameters.

Actually, two separate functions are minimized, but they share the form described above. One function fits the shape of the source array, and this defines the cylindrical distortion vector $\vec{D}$; the other fits for the remaining 7 BCAM parameters. Since the two set of solutions are not entirely decoupled, the fits are repeated iteratively until they converge on a single set of 9 parameters. This typically involves two minimizations for each function, and order (which function is minimized first) does not seem to matter.

If we know the absolute position of the sources (which we do not), then this minimization procedure can fit for all 9 calibration parameters from a single BCAM orientation at two ranges. Since we have four BCAM orientations, we get four semi-independent sets of results. Unfortunately, while we know the relative source positions, including the overall rotation, we do not know the absolute position of the source block with a high degree of precision. This affects the transverse position of the pivot point, since if the source block is placed too high, e.g., the fit will try to compensate by moving the PP down by the same amount. The other parameters should not be affected by such a source block translation. Fortunately, opposite pairs of BCAM orientations (0° and 180°, or 90° and 270°) make equal and opposite errors in the PP parameters. Averaging opposing pairs cancels the errors and yields two semi-independent sets of BCAM parameters. For simplicity, let the average of the 0° and 180° orientations be the vertical solution, and the average of the 90° and 270° orientations be the horizontal solution. These should be the same, and their average is the quoted calibration result for a BCAM. A rough estimate of the precision of each parameter is taken to be 1/2 the difference between the vertical and horizontal solutions for that parameter. See Table 3 for an example output from the parameter fit.

	Pivot Point			CCD Origin			Rotation	Cyl. Distortion	
	P_x (mm)	P_y (mm)	P_z (mm)	t_x (mrad)	t_y (mrad)	F (mm)	ψ (mrad)	$D_x \times 10^3$	$D_y \times 10^3$
0°	10.1291	17.3544	-15.6343	-0.7877	-0.3643	153.8669	-18.5586	1.9592	-1.3063
90°	10.3705	16.9231	-15.1128	-0.8477	-0.2393	153.8311	-18.3217	1.9056	-0.1560
180°	9.9389	16.6503	-17.9022	-0.6848	-0.1978	153.9796	-18.6213	3.2526	-1.1536
270°	9.6584	17.1263	-15.3619	-0.6789	-0.3933	153.8366	-18.2988	1.8792	-0.2394
Vertical	10.0340	17.0023	-16.7683	-0.7363	-0.2810	153.9232	-18.5900	2.6059	-1.2300
Horizontal	10.0144	17.0247	-15.2373	-0.7633	-0.3163	153.8338	-18.3102	1.8924	-0.1977
Average	10.0242	17.0135	-16.0028	-0.7498	-0.2987	153.8785	-18.4501	2.2492	-0.7138
Error	0.0098	0.0112	0.7655	0.0176	0.0176	0.0447	0.1399	0.3568	0.5161

Table 3: Example output from the global minimization analysis. The vertical and horizontal averages are semi-independent, and the row marked "Average" represents the average of those two.

5 Calibration Results

5.1 Precision

The calibration precision is composed of two parts: software and measurement precision. The difference between the two reconstruction programs on the same set of data represents the software precision. Repeating the calibration with new data for a given BCAM represents the measurement precision.

5.1.1 Agreement between the two reconstruction methods

Table 4 shows the calibration constants from 8 BCAMs calibrated with the analytical approach; Table 5 shows the calibration constants for the same 8 cameras (same data) analyzed with the global fit method. The difference between the calibration constants using the two methods is shown in Table 6. One can see that the differences between the two methods are small with respect to the uncertainties in the parameters themselves.

Camera	Pivot Point			CCD Origin			Rotation	Cyl. Distortion	
Number	P_x (mm)	P_y (mm)	P_z (mm)	t_x (mrad)	t_y (mrad)	F (mm)	ψ (mrad)	$D_x \times 10^3$	$D_y \times 10^3$
41	10.024	17.015	-16.065	-0.750	-0.299	153.892	-18.423	2.200	-0.600
42	9.980	17.017	-15.612	1.771	-0.947	153.537	0.864	8.000	0.200
44	10.006	17.041	-14.775	-0.873	-0.742	153.783	4.341	4.600	0.600
45	9.967	17.038	-13.976	-1.027	1.483	153.484	0.141	9.600	4.000
46	10.024	16.991	-16.289	-1.047	0.306	153.360	2.173	9.400	1.200
49	9.943	16.987	-16.376	-1.228	0.339	153.120	41.301	5.600	2.600
50	10.019	17.026	-14.977	-0.416	0.862	153.823	-0.907	9.000	1.400
51	10.026	17.043	-15.797	1.264	-0.531	151.836	-1.181	10.600	-6.000
Error	0.004	0.003	0.730	0.020	0.020	0.071	0.107	—	—

Table 4: Calibration constants for 8 BCAMs using the analytical construction. The error line at the bottom represents the standard deviation of each calibration constant for a given camera; it is not the standard deviation of the distribution of constants from all 8 BCAMs, as they do not have a common mean. The error on the cylindrical distortion parameters has not yet been calculated in the analytical approach.

Camera	Pivot Point			CCD Origin			Rotation	Cyl. Distortion	
Number	P_x (mm)	P_y (mm)	P_z (mm)	t_x (mrad)	t_y (mrad)	F (mm)	ψ (mrad)	$D_x \times 10^3$	$D_y \times 10^3$
41	10.024	17.014	-16.003	-0.750	-0.299	153.879	-18.450	2.249	-0.714
42	9.980	17.016	-15.604	1.771	-0.947	153.556	0.591	7.791	0.234
44	10.006	17.040	-14.832	-0.874	-0.742	153.800	4.160	3.427	0.492
45	9.967	17.037	-14.041	-1.027	1.483	153.548	-0.132	8.942	3.741
46	10.023	16.991	-16.299	-1.047	0.306	153.369	2.005	9.358	1.158
49	9.943	16.987	-16.340	-1.228	0.340	153.159	41.271	4.996	2.618
50	10.018	17.026	-15.046	-0.416	0.862	153.859	-1.134	8.621	1.578
51	10.025	17.042	-15.678	1.264	-0.532	151.840	-1.313	10.608	-5.964
Error	0.005	0.008	0.971	0.018	0.015	0.084	0.140	0.621	0.587

Table 5: Calibration constants for 8 BCAMs using the global fit construction (MINUIT). The error line at the bottom represents the standard deviation of each calibration constant for a given camera; it is not the standard deviation of the distribution of constants from all 8 BCAMs, as they do not have a common mean.

Camera	Pivot Point			CCD Origin			Rotation	Cyl. Distortion	
Number	P_x (mm)	P_y (mm)	P_z (mm)	t_x (mrad)	t_y (mrad)	F (mm)	ψ (mrad)	$D_x \times 10^3$	$D_y \times 10^3$
41	0.000	0.001	-0.062	0.000	0.000	0.013	0.027	-0.049	0.114
42	0.000	0.001	-0.008	0.000	0.000	-0.019	0.273	0.209	-0.034
44	0.000	0.001	0.057	0.001	0.000	-0.017	0.181	1.173	0.108
45	0.000	0.001	0.065	0.000	0.000	-0.064	0.273	0.658	0.259
46	0.001	0.000	0.010	0.000	0.000	-0.009	0.168	0.042	0.042
49	0.000	0.000	-0.036	0.000	-0.001	-0.039	0.030	0.604	-0.018
50	0.001	0.000	0.069	0.000	0.000	-0.036	0.227	0.379	-0.178
51	0.001	0.001	-0.119	0.000	0.001	-0.004	0.132	-0.008	-0.036
Average	0.000	0.001	-0.003	0.000	0.000	-0.022	0.164	0.376	0.032
Std. Dev.	0.000	0.001	0.067	0.000	0.000	0.024	0.097	0.420	0.131

Table 6: Difference in calibration constants as computed by the analytical approach and the global fit method (analytical−global fit).

5.1.2 Repeatability

In order to test the measurement precision, we repeated the calibration data-taking for each of the 8 BCAMs. The measurement precision involves placement of the BCAM in the roll cage (once per calibration), the placement of the roll cage on the brass base in each of the BCAM orientations(8x per calibration), the placement of the source block (twice per calibration), and the inherent resolution of the camera.

Table 7 shows the results of the repeatability test on the calibration parameters. The standard deviation

Camera	Pivot Point			CCD Origin			Rotation	Cyl. Distortion	
Number	P_x (mm)	P_y (mm)	P_z (mm)	t_x (mrad)	t_y (mrad)	F (mm)	ψ (mrad)	$D_x \times 10^3$	$D_y \times 10^3$
41	0.009	0.000	-0.413	-0.005	-0.005	0.007	-0.149	0.200	0.000
42	-0.008	0.003	-1.386	0.002	-0.008	0.126	0.215	0.000	0.000
44	-0.010	0.003	-0.480	-0.001	-0.002	0.078	-0.055	0.000	0.000
45	-0.003	-0.008	0.383	0.001	-0.003	-0.036	0.124	0.000	0.000
46	0.001	-0.012	-0.579	-0.013	-0.011	0.027	0.022	0.200	0.000
49	-0.005	-0.007	0.454	0.003	-0.002	-0.059	-0.028	0.000	0.200
50	-0.006	-0.001	1.473	0.002	-0.002	-0.195	-0.202	0.200	0.000
51	-0.014	0.011	1.099	0.005	-0.004	-0.125	-0.067	0.000	0.000
Average	-0.005	-0.001	0.069	-0.001	-0.005	-0.022	-0.017	0.075	0.025
Std. Dev.	0.007	0.007	0.952	0.006	0.003	0.105	0.137	0.104	0.071

Table 7: Difference in calibration constants between two sets of calibration data, analyzed by the analytical method.

represents the measurement precision, and it is of the same order as the software precision.

5.2 Accuracy: Freiburg Test

Although the BCAM calibration appeared to be internally consistent (precise and repeatable), there was no way to test the accuracy of the devices without additional equipment. With this in mind, a set of six calibrated BCAMs were sent to Freiburg to take advantage of the large CMM available there. The Freiburg test is described in detail in ATLAS Note @@@. (This note may or may not exist, sent an email to Karen asking her about it.)

5.2.1 The 4m test

There were two setups used at Freiburg. In the first, four cameras were arranged at three stations roughly equidistant over a total span of 4.4m (see Figure 13). The outer two stations (EI and EO) each contained
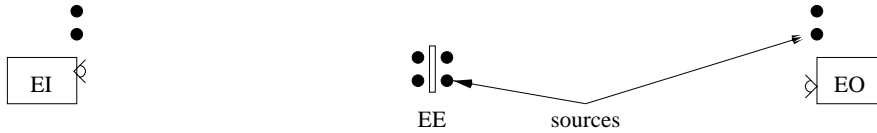
Figure 13: The Freiburg 4m test setup. Figure not drawn to scale.

a BCAM and two light sources (LED's), and these faced inward toward the middle station (EE). The middle station had four sources, two facing each end station. From each end station, at least one source at each of the other two stations could be seen.

The BCAMs sat on three ball mounts as described in Section 1, and the sources were mounted independently. The positions of the three ball mounts for the BCAMs as well as some position information for the sources were obtained from the Freiburg CMM *in situ*. Then, light centroids were obtained for each source that could be viewed from each BCAM. With this information, two different measurements of sagitta could be derived.

The sagitta angle was defined as the angle between the two rays from a single BCAM pointing toward two sources, one source nearer and the other farther away. We obtained two measurements of the sagitta angle. The "mechanical" sagitta angle used the CMM measurements of the source positions and BCAM three ball mount, our own independent source calibration, and the pivot point position from the BCAM calibration. The "optical" sagitta angle was just the difference in bearing between the two rays obtained from the BCAM measurements of the sources. The difference between these two sagitta angle measurements we took to be our sagitta measurement error.

Each end BCAM could only see one EE source and both of the other end sources, which were given the suffixes L and R to distinguish them. So, for instance, the EI BCAM might view one EE source and one EO source. The sagitta in this case might be EI-EE-EOL or EI-EE-EOR. Neglecting the EE BCAMs, there were thus four different sagittas in the Freiburg 4m test (EI-EE-EOR, EI-EE-EOL, EO-EE-EIR, and EO-EE-EIL). In addition, the EE station was placed in four different vertical positions and also rotated by three angles about the vertical axis (perpendicular to the surface of the granite table of the CMM), and both CMM and BCAM measurements were taken in each configuration. Some configurations were duplicated for consistency, and including these, there were nine different configurations measured in the 4m test.

Each BCAM took around 16 measurements of each spot and averaged them for better precision. The optical sagitta was then calculated from these spot positions and compared to the mechanical sagitta. The difference between these two definitions of sagitta gives some measure of the accuracy of the BCAM calibration. This is actually an overestimation of the BCAM calibration accuracy because the source calibration accuracy is also involved, and this is unknown.

In the nine configurations we measured a total of 34 different sagittas, both optical and mechanical.[3] The RMS difference between the 34 sagitta angles was 7.2μrad. This gives an upper limit on the calibration accuracy of the BCAMs. In ATLAS, however, duplicate measurements of the same sagitta would be averaged to obtain a more accurate assessment. Averaging over the four different optical sagittas for each test configuration and comparing the result to the average mechanical sagitta yields a RMS sagitta angle of 5.6μrad.

The absolute accuracy of the BCAMs could be estimated by comparing the bearing of a ray from the BCAM to a given source, both optically and mechanically, as for the sagitta comparison. Averaging over all BCAM–source pairs in all nine configurations yields an absolute position RMS error of 42μrad.

5.2.2   The 16m test

The other test at Freiburg was conducted over the longer baseline of 16m. In this test, there were four stations (see Figure 14). Again, because of mirror calibration problems, only the outermost two BCAMs

---

[3]In one configuration the EI BCAM could not see both sets of sources, so we only had two sagitta measurements for that configuration.
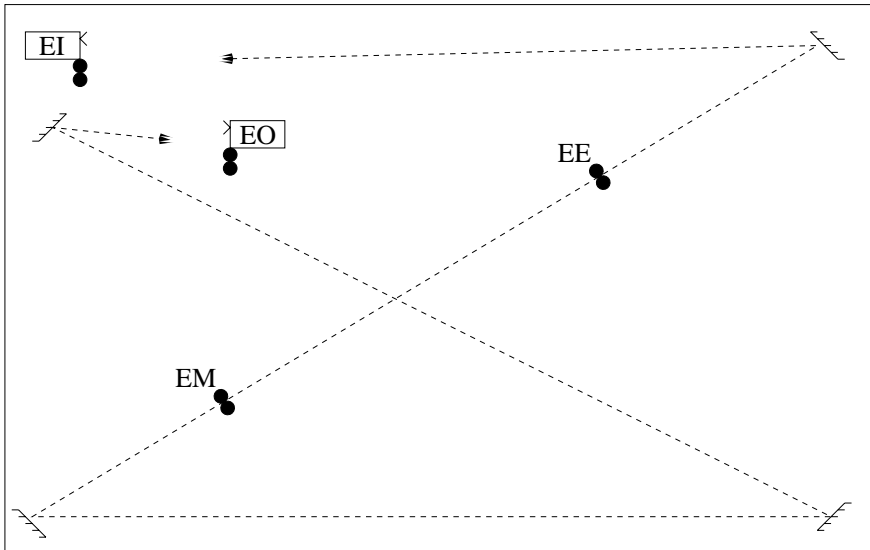
Figure 14: The Freiburg 16m test setup. Figure not drawn to scale.

were used in this test, and they each observed the sources at the other three stations. The distances from station EI to the other three were approximately 4.15m, 6.55m, and 16.34m. As in the 4m test, one station (EE) was moved up and down and rotated for a total of 7 different configurations. The RMS error between the measurements of the stationary sources from run to run was around 2μrad. The RMS error in the optical measurement as compared to the mechanical measurement for the moving sources was also around 2μrad for the EO BCAM (12m from EO to the sources on EE), and around 6μrad for the EI BCAM (4m from EI to EE). This is consistent with the RMS sagitta error observed in the 4m test.

6 Conclusion

We have established a calibration procedure for the BCAMs used in the Atlas endcap muon polar alignment system, we have performed this procedure on a sample of BCAMs, and we have evaluated the accuracy of the calibration at Freiburg. The calibration involves nine parameters per BCAM, and the two methods of deriving these parameters agree with each other. The test at Freiburg shows that the upper limit on the absolute position accuracy is 42μrad, and the tracking or relative accuracy has an upper limit of 7μrad.

Appendix I

This appendix describes the effect of a small aperture on the seven parameters which describe a camera in the absence of cylindrical distortions. First, we consider the image of a point source with known location on a calibrated camera with no aperture. Figure 15 shows a simple lens–CCD camera. The dark horizontal line connects the pivot point (the center of the lens) with the center of the CCD. An object (or light source) at a range r a distance h_o from this center line will appear as a focused image on the image plane, a transverse distance (height) h_f from the center line. If the CCD is inserted a distance α behind the lens, then an (out-of-focus) image will appear on the CCD with a centroid height h_c from the center line. By similar triangles, it is apparent that

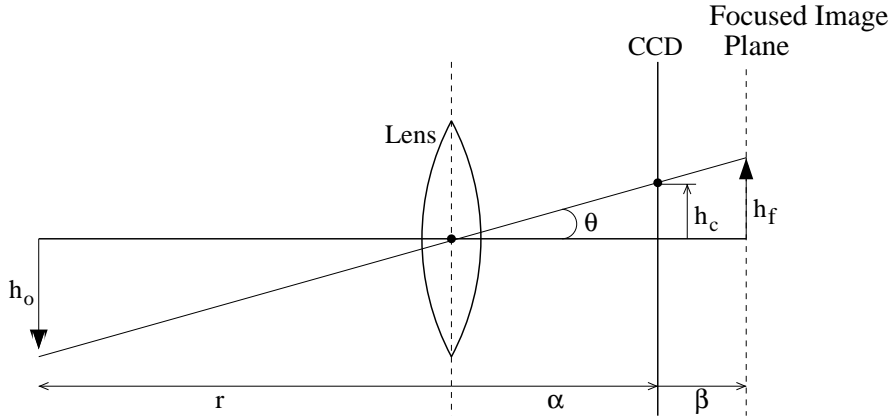$$\frac{h_c}{\alpha} = \frac{h_o}{r} = \frac{h_f}{\alpha + \beta}. \tag{2}$$

16

Figure 15: Simple lens–CCD camera.

$h_c/\alpha = h_o/r = h_f/(\alpha + \beta)$. The location of the center line and the distance α are given from the calibration constants directly, and h_o comes from the location of the source with respect to the center line.

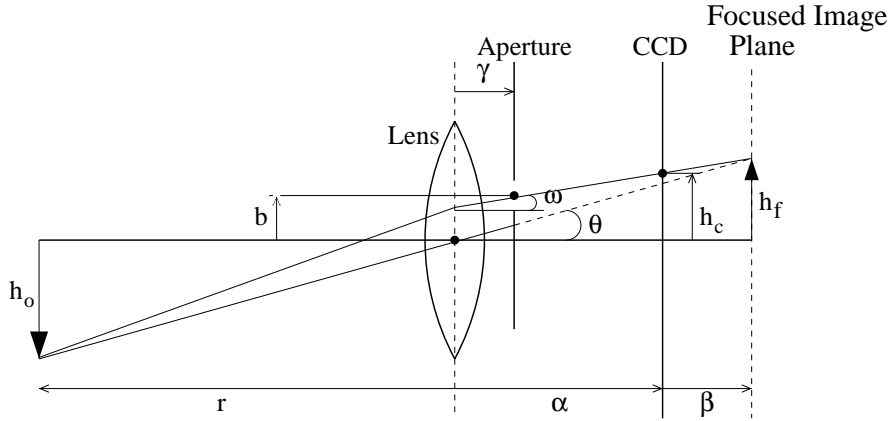If we add a small aperture, we have the situation diagrammed in Figure 16. The aperture is displaced

Figure 16: Lens–CCD camera with an aperture.

from the lens by a longitudinal distance γ and a transverse distance b. Since all rays from the object will converge at the image plane, we obtain the angle ω:

$$\omega = \frac{h_f - b}{\alpha + \beta - \gamma} \quad \text{for} \quad \omega \ll 1. \tag{3}$$

Then, the image on the CCD will appear at a height

$$h_c = b + \omega\,(\alpha - \gamma). \tag{4}$$

Using Equation 3 to remove the ω from Equation 4, we get

$$h_c = b + (h_f - b)\,\frac{\alpha - \gamma}{\alpha + \beta - \gamma}. \tag{5}$$

Removing h_f by the use of Equation 2 yields

$$h_c = b + \left[\frac{h_o}{r}\,(\alpha + \beta) - b\right]\frac{\alpha - \gamma}{\alpha + \beta - \gamma}. \tag{6}$$

17

From the thin lens approximation of a lens with focal length f, we have the familiar lensmaker's formula

$$\frac{1}{r} + \frac{1}{\alpha + \beta} = \frac{1}{f} \tag{7}$$

which we can use to remove the β from Equation 6:

$$h_c = b + \left[\frac{h_o}{r} - \frac{b}{\alpha + \beta}\right] \frac{\alpha - \gamma}{1 - \gamma/(\alpha + \beta)}$$

$$h_c = b + \left[\frac{h_o}{r} - b\left(\frac{1}{f} - \frac{1}{r}\right)\right] \frac{\alpha - \gamma}{1 - \gamma\left(\frac{1}{f} - \frac{1}{r}\right)}$$

$$h_c = b + \left[\frac{h_o + b}{r} - \frac{b}{f}\right] \frac{\alpha - \gamma}{1 - \frac{\gamma}{f} + \frac{\gamma}{r}}$$

$$h_c = b + \left[\frac{h_o + b - br/f}{r + \gamma/\left(1 - \frac{\gamma}{f}\right)}\right] \frac{\alpha - \gamma}{1 - \frac{\gamma}{f}} \tag{8}$$

Making the substitutions $h_o' = h_o + b$, $r' = r + \gamma/\left(1 - \gamma/f\right)$, and $\alpha' = (\alpha - \gamma)/\left(1 - \gamma/f\right)$, we get

$$h_c = \frac{h_o' \alpha'}{r'} + b\left(1 - \frac{r\alpha'}{r'f}\right).$$

Writing this entirely in terms of the primed coordinates, we have

$$h_c = \frac{h_o' \alpha'}{r'} + b\left(1 - \frac{r'\alpha'}{r'f} + \frac{\gamma\alpha'}{r'f\left(1 - \gamma/f\right)}\right)$$

$$h_c = \frac{h_o' \alpha'}{r'} + \frac{b\gamma\alpha'}{(f - \gamma)r'} + b\left(1 - \frac{\alpha'}{f}\right).$$

If we define $h_o'' \equiv h_o' + b\gamma/\left(f - \gamma\right) \equiv h_o + b\left[1 + \gamma/\left(f - \gamma\right)\right]$, then we have

$$h_c = \frac{h_o'' \alpha'}{r'} + b\left(1 - \frac{\alpha'}{f}\right). \tag{9}$$

Notice that the first part of this equation looks similar to Equation 2, except that the variables are primed. Notice also that the primed variables depend only on their unprimed counterparts, the aperture parameters b and γ, and the focal length f. Since the unprimed variables are directly related to the camera calibration constants, this means that the aperture effects can be completely absorbed in a reinterpretation of the calibration constants. For instance, as long as the longitudinal axis maintains its direction, the change in h_o can be interpreted as a transverse shift in the pivot point position. Likewise, the change in r can be attributed to a shift in the longitudinal position of the pivot point. The change in α can be interpreted in a change in the pivot point to CCD distance. Finally, the second term in Equation 9 can be absorbed in a CCD transverse shift. In other words, a camera with a small aperture is optically equivalent to another lens–CCD camera with no aperture, so the seven parameters which describe that camera should apply equally well to the camera with the aperture.

So far, we have assumed a point source, and we have stated without proof that the centroid of the out-of-focus image follows a straight line from aperture to focus. Let us argue for the second statement first. Assuming a point source, the light distribution across the aperture will be uniform for a distant source as long as the aperture is fully illuminated. Since all rays follow straight lines to the focus spot, the light distribution from aperture to focus follows some sort of a cone shape (Figure 17). The centroid of the distribution at the aperture is therefore at the geometric middle of the aperture, and the centroid of the focused spot is of course at the spot position itself. At any plane cutting through the cone of light which is parallel to the aperture (i.e. the CCD), the centroid will be in the middle of the circle (or oval) of light on that plane; thus, the centroid follows a straight-line path from the center of the aperture to the focus (dotted lines in Figure 17).
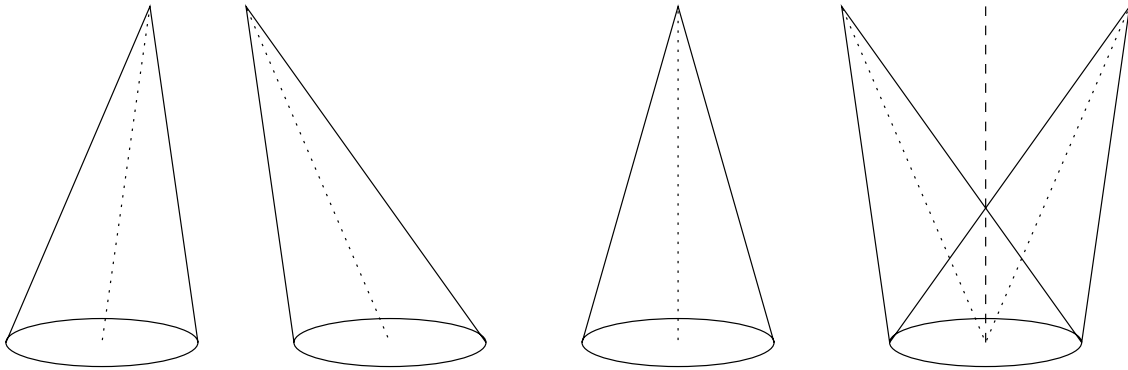
Figure 17: Light distributions from various point sources from apertures (bottom) to foci (top). The three leftmost images show single point sources at different source positions but a common range. The dotted line follows the centroid of the light distribution on a plane parallel to the aperture. The right image shows the light distribution for a combination of two point sources. The centroids of the two sources individually are marked by the dotted lines, and the composite centroid is given by the dashed line.

What about non-point sources? Any non-point source can be composed of up to a countably infinite number of point sources. How would this complex source behave? Consider a combination of two point sources of equal intensities, but different positions. The light distributions at the aperture are both uniform, but their cones converge on different spots at the image plane (Figure 17, right). The centroids of each separate cone can be traced from the center of the aperture to their focal points, and the aggregate centroid of both point sources at any point between aperture and image plane is just the midpoint between the two separate centroid lines (dashed line in Figure 17). At the image plane, for eaxmple, the aggregate centroid is the midpoint between the two image points. Since the sum of the two centroid lines is itself a line, any two point sources create a centroid that behaves like that of a third, carefully chosen point source. If the sources do not have equal intensities, then the aggregate centroid should be some linear combinations of the two individual centroids, which still describes a straight line. The unfocused image on the CCD may be ugly (for two point sources it might look like a fuzzy figure-eight), but its centroid behaves linearly.

This concept is extendable to any non-uniform source, composed of any countable number of point sources. The centroid at the aperture should be in the geometrical center of the aperture, and the centroid at the focused image is predictable from the source configuration. Between them the centroid is composed of a linear superposition of distinct point source centroids, and this superposition should map out a straight line from center of aperture to image plane. Since that is exactly the same behavior as the centroid of some carefully constructed point source, which obeys Equation 9, then we can safely say that the same behavior holds for a generic source.

Appendix II

This appendix shows an alternate method of deriving the position and direction of the cylindrical distortion vector. The method described in Section 4.3.1 has the advantage that it does not use the (imperfect) knowledge of the relative source positions; however, the method described in this appendix uses only a single BCAM orientation.

The distance between two images $\vec{T_1}$ and $\vec{T_2}$ on the CCD compared to the distance between corresponding sources in the source block $\vec{S_1}$ and $\vec{S_2}$ shows the effect of the cylindrical distortion plus the overall magnification. The largest ratio $\left|\vec{T_1} - \vec{T_2}\right| / \left|\vec{S_1} - \vec{S_2}\right|$ of the six possibilities (cf. Figure 10a) happens for the pair that lies closest in direction to the distortion vector. Since the sources are arranged roughly in a square, one can always find two other pairs which lie at $\pm45°$ from the vector $\vec{T_1} - \vec{T_2}$. Using the relative

lengths of these three vectors on the CCD, it is possible to reconstruct both the direction and magnitude of the distortion vector $\vec{D}$.

The length L_{ab} of a vector on the CCD between points $\vec{T_a}$ and $\vec{T_b}$ is

$$L_{ab} = \left| \vec{T_a} - \vec{T_b} + \left[\left(\vec{T_a} - \vec{T_b} \right) \cdot \vec{D} \right] \widehat{D} \right| \tag{10}$$

where the usual notation holds that a vector $\widehat{A}$ represents the unit vector in the direction of vector $\vec{A}$. Dividing the length L_{ab} by the distance between the corresponding sources $\vec{S_a}$ and $\vec{S_b}$ yields the ratio

$$R_{ab} = \left| \frac{\vec{T_a} - \vec{T_b}}{\left| \vec{S_a} - \vec{S_b} \right|} + \left[\frac{\vec{T_a} - \vec{T_b}}{\left| \vec{S_a} - \vec{S_b} \right|} \cdot \vec{D} \right] \widehat{D} \right| \tag{11}$$

$$R_{ab} = \left| m \widehat{T_a - T_b} + m \left[\widehat{T_a - T_b} \cdot \vec{D} \right] \widehat{D} \right|$$

$$R_{ab} \approx m \left(1 + D \cos \theta_{ab} \right) \tag{12}$$

where m is the magnification, which is unknown, D is the length of vector $\vec{D}$, and θ_{ab} is the angle between the vector $\vec{T_a} - \vec{T_b}$ and $\vec{D}$. If R_0 represents this ratio for the vector lying closest to $\vec{D}$, and the others are denoted by $R_\pm$, then we may write

$$R_\pm = m \left[1 + D \cos \left(\theta_{ab} \pm 45° \right) \right]$$

$$R_\pm = m \left[1 + D \cos \theta_{ab} \cos 45° \mp D \sin \theta_{ab} \sin 45° \right]$$

$$R_\pm = \frac{m}{\sqrt{2}} \left[\sqrt{2} + D \cos \theta_{ab} \mp D \sin \theta_{ab} \right] . \tag{13}$$

Adding R_+ and R_-, we get

$$R_+ + R_- = \sqrt{2} m \left[\sqrt{2} + D \cos \theta_{ab} \right] \tag{14}$$

and subtracting, we have

$$R_+ - R_- = -\sqrt{2} m D \sin \theta_{ab}. \tag{15}$$

Subtracting $1/2$ of Equation 14 from Equation 12 yields

$$m D \cos \theta_{ab} \left(1 - \frac{1}{\sqrt{2}} \right) . \tag{16}$$

Adding the appropriately weighted squares of Equations 15 and 16 yields the product $(mD)^2$. Since m and D are both positive, we have mD. This may be inserted into Equation 15 to get $\sin \theta_{ab}$, and since we know by construction that $\theta_{ab} < 45°$, we take the arcsine to get θ_{ab}. Since we know both the product mD and the angle θ_{ab}, we can solve Equation 12 for the magnification m. Finally, since we know both the product mD and m, we solve for D directly. With the magnitude of $\vec{D}$ known, we can then solve Equation 11 (or 10) for the direction of the distortion vector.

If the three vectors are separated by some known angles other than 45°, this method may still be used, but the algebra is harder since all the $\sqrt{2}$ factors get replaced with sines and cosines of known angles, and the cancellation is not so nice in general. In a computer program, of course, one would use the sines and cosines of the known angles, rather than the factors of $\sqrt{2}$ everywhere, so this is not a problem.