Calibration of an Array Camera

*Sean Curry** and *Sheldon Baumrind*
Craniofacial Research Instrumentation Laboratory, University of California, San Francisco, CA 94143
J.M. Anderson
Department of Civil Engineering, University of California, Berkeley, CA 94720

ABSTRACT: The Craniofacial Research Instrumentation Laboratory at the University of California, San Francisco, has begun to employ an array camera to automate their biostereometric data acquisition procedures. Data from close-range stereo x-rays and photographs have been utilized for several years in the analysis of morphological changes in the cranial region during orthodontic treatment and after orthognathic surgery. The use of an array camera will help to automate the routine digitizing of these images.

The array camera presently being studied has a charge injection device (CID) sensor with a square array of 128 by 128 pixels. This camera was calibrated to determine the sensor and lens characteristics. A test field composed of a movable grid plate was imaged under varying conditions. Image processing techniques were employed to locate the target points to sub-pixel resolution. A bundle block adjustment with added parameters for camera calibration was developed to perform the actual calibration.

Although the present camera is limited to the imaging of small objects, the results indicate that the bundle block calibration procedure can be successfully applied to an array camera. The techniques developed will be utilized to calibrate larger array cameras in the future.

INTRODUCTION

THE CRANIOFACIAL RESEARCH INSTRUMENTATION LABORATORY (CRIL) at the University of California, San Francisco, has for a number of years been developing instrumentation and software packages for the acquisition and analysis of biostereometric data. Stereo photographs and stereo x-rays of the head are used to analyze growth and treatment effects during various orthodontic treatment programs. In addition, data are obtained from stereo photographs of dental study casts and merged with the coordinate data obtained from the other records (Baumrind, 1984).

These systems have been developed in conjunction with the Civil Engineering Department of the University of California, Berkeley, and are now in routine use at UCSF. Variations of these techniques are also being used to monitor hip prosthesis loosening and vertebral movement following spinal fusion surgery.

One problem that characterizes all of these data analysis tasks is the large volume of data that must be manually acquired. For each patient studied with the full integrated three-dimensional system, there are pre-treatment and post-treatment stereo x-ray and stereo photographic pairs, and stereo photographs of the study casts. Although only a small set of points on all of these physical records is analyzed, there is still a large volume of manual digitizing. Even with the present on-line digitizing software, this portion of each study takes a substantial amount of time.

Therefore, the CRIL has begun to experiment with methods to reduce the digitizing effort. Although we do not expect to completely automate our procedures within the immediate future, we have begun to develop methodologies now, in view of the following factors:

- Array cameras are becoming more readily available. Array sizes are increasing and prices are dropping. Techniques that are developed now will be transferable to the next generation of hardware.
- Direct digital radiography is making its advent in hospitals. Once it becomes widespread, we will have a direct digital source of data for the x-ray portions of our analyses.
- For small objects, such as dental study casts, we can acquire direct digital data now with sufficient resolution, if sub-pixel point location techniques are employed.
- Resolution for larger images such as stereo cranial x-rays can be increased by moving a lower resolution camera over the larger film surface, thus dividing the image into a series of sub-images.
- The identification and tracking of anatomic features (as opposed to discrete implant and reference markers) will require several years of development (Curry, 1984). This important part of our work is proceeding experimentally, and will be a powerful tool when combined with new hardware.

With these factors in mind, the CRIL has begun to study the problem of automating its data acquisition. An array camera has been interfaced to a minicomputer, and software has been written for digital picture acquisition and storage. The first major task undertaken was the calibration of the array camera.

*Presently with Vexcel Corporation, 2905 Wilderness Pl., Boulder, CO 80301.

PHOTOGRAMMETRIC ENGINEERING AND REMOTE SENSING,
Vol. 52, No. 5, May 1986, pp. 627–636.

0099-1112/86/5205-627$02.25/0

HARDWARE AND SOFTWARE SYSTEM

A prime consideration in developing the prototype system was the use of low-cost, readily available hardware. We pay a penalty in slower processing speeds and lower resolution. However, once our initial system testing is completed, a second generation of cameras and processors (higher resolution and faster) will be acquired on the basis of what has been learned in this phase.

The camera employed at this point is a GE TN2200 CID (charge injection device) solid state camera (Figure 1) with an array of 128 by 128 pixels. It is controlled by a PN2200 Automation Interface, which supplies power and timing signals to the camera, and converts the analog outputs from the camera to 8 bits of digital data, equal to 256 grey levels. Frames from the camera are stored in a Poynting Products Frame Buffer, and passed to a DUAL 83/80 computer. Pictures are displayed through a Poynting Products video buffer onto a standard TV monitor. Two stepper motors are also interfaced to the DUAL to allow object rotation under the camera.

All software is written in the C language, and runs under the UNIX operating system. This software combination is widely transportable, compact and fast, and easy to maintain.

CALIBRATION

A radiometric and geometric calibration procedure for the TN2200 camera was developed and tested. Radiometric criteria evaluated included detector response linearity, system noise, and system stability over time. Geometric calibration consisted of the evaluation of array regularity, principal point and principal distance, and symmetric radial lens distortion coefficients. A full bundle block adjustment program with added parameters for camera calibration was implemented.

As Gruen (1983) has noted, both radiometric and geometric calibrations are important, and are in fact interconnected. For example, nonlinearity in detector response or high levels of system noise could affect subsequent point location operations. Generally, radiometric corrections should be applied before any geometric operations are performed.

FIG. 1. TN2200 array camera.

RADIOMETRIC CALIBRATION

Sensor Linearity. The charge injection device (CID) array in the TN2200 camera is a semiconductor device. Its spectral response peaks in the 0.7-μm wavelength region, the boundary between red and the near infrared. Pixel illumination is converted to an output voltage by the array, and then to a digital value in the range of 0 (black) to 255 (white) by the analog to digital (A/D) converter. Ideally, the entire conversion process should be linear.

A Kodak Grey Step Card (log linear to $\approx$ 0.5 percent) was imaged under a number of different lighting conditions, and with a number of different camera speeds. From each resulting digital image, the average intensity of a 3 by 3 pixel region in the center of each density step was computed. These average intensities were plotted on a log linear scale, and found to be within 1 percent of a straight line response. This level of sensor linearity was satisfactory for our purposes.

System Noise. System noise is defined as fluctuations in pixel grey levels caused by random and systematic perturbations in the CID sensor and A/D converter. Systematic noise can be removed by subtracting an image of the noise from each picture, while random noise must be reduced through several other techniques.

Systematic noise is seen as a repeated spurious response by a sensor element. A spurious response is generally defined as an output of more than 10 percent of the saturation voltage, in comparison with the pixel's nearest neighbors. The saturation voltage is the maximum usable signal output voltage from the sensor. Blemished pixels are determined by running the camera in the dark (lens cap on) and measuring the output from each pixel. There will always be a minimum signal called the dark signal, defined as the output signal caused by thermally generated electrons. It is a linear function of integration time and an exponential function of chip temperature (Fairchild, 1981). Responses above this dark current are spurious, and those pixels must be rejected from subsequent images if the response cannot be corrected.

The TN2200 camera was run at several clock rates with the lens cap in place. At each rate, an image of the dark signal was stored, and the process repeated after approximately 30 minutes. The first row was always unacceptable and was rejected from all images. Other minor pixel values above the dark current were fairly stable. A mask was generated with these stable spurious pixels, and those pixels were subtracted from all subsequent images.

Random noise results from imperfections in the

sensor manufacture, electron leakage between sensor photosites, variations in the analog to digital (A/D) conversion, and other factors. Levels of random noise can be approximately quantified by running the camera at various clock rates, which affects pixel integration times, and under different lighting conditions, while viewing a neutral grey surface. A Kodak Grey Card (Part # R-27) was chosen for the grey surface. Images were stored as single frames, and also as the average of five and ten consecutive frames. Comparison images (single frames and averaged frames) were then acquired an hour later under exactly the same conditions. The two like images made at different times were compared after removal of systematic bad pixels. Under perfect conditions, with a repeatable sensor and A/D converter, each pair would be expected to compare perfectly, pixel by pixel. A difference function (Rosenfeld, 1982) was used to compute the maximum absolute pixel differences, the sum of the absolute pixel differences, and the sum of the squares of the absolute pixel differences.

Under both lighting conditions, the overall best performance was achieved at the 2.5-MHz clock rate (15 frames per second), especially when considering the maximum absolute differences. At that clock rate, the lowest sum of squares of differences was achieved with the average of ten frames, but there was no improvement in the maximum absolute differences when using the average of five frames compared to using ten frames. These tests were performed at a temperature of approximately 25 degrees Celsius. It was subsequently determined that, at 21 degrees, the random noise diminished even further.

Blooming was evident at all frame rates and temperatures. Blooming occurs when a bright region in the frame locally overloads the sensor, causing overflow into adjacent pixel rows and columns beyond the actual bright spot. Although the CID sensor is less sensitive to this process than a Charge Coupled Device (CCD) camera, it is still a serious problem in high contrast scenes.

Another basic measure of the image quality properties of a lens/detector system is the modulation transfer function (MTF). It can be determined by imaging a square bar target array, or by imaging a sharp edge, determining the resulting line spread function (LSF), and taking the Fourier transform of the LSF. Due to time and equipment constraints, the MTF of the TN2200 camera was not determined.

Sensor Stability. After the camera and associated electronics are switched on, the sensor and controller undergo thermal and other changes. A test was devised to check the internal stability of the sensor readout over a period of eight hours. A single discrete point was placed in the field of view of the camera, and an image ($n=5$) stored every hour for eight hours. The first image was acquired immediately after system start-up. No attempt was made to control the ambient temperature, which varied between 22 and 26 degrees Celsius.

Using the sub-pixel location techniques described later, the point location in x and y was computed to 0.01 pixels least count. The test was repeated twice, on two different days. The sensor was very stable, although there was a slight drift immediately after the initial power-up.

Recommendations for Image Acquisition with the TN2200 Camera.

(1) Compute an average dark current image, and subtract it from all subsequent images.
(2) Discard the first row of every image.
(3) Run the camera at 2.5MHz, with the temperature as low as possible.
(4) Store each image as the average of five consecutive frames.
(5) Avoid extreme low light conditions.
(6) Avoid high contrast scenes which cause blooming.
(7) Turn the camera system on 1 hour before use.

Geometric Calibration

The array camera can be calibrated as can any other non-metric photogrammetric camera. The CID sensor can be treated as a small, low-resolution photographic plate, and the camera principal point location, principal distance, and lens distortion should all be evaluated. The solid state sensor has certain advantages and limitations when compared to a film based camera.

Advantages

- There are no film distortions due to film buckling and emulsion problems.
- Any sensor geometry irregularities are theoretically stable and can be mapped.
- The "film" is always in perfect registration, allowing multiple exposures.
- There are no fiducials. The image coordinate system is inherent in the array itself.

Limitations

- The present sensor is of low resolution, which limits the area that can be imaged, and necessitates extra image processing.
- When photons strike the sensor surface, there is a stochastic process determining which photo site actually registers the photons striking near pixel boundaries.
- There may be other sensor output irregularities, associated with the A/D conversion process, that will distort the image geometry. Additional preprocessing is required to reduce their impact.

Array Geometry. The CID array is fabricated as a silicon semiconductor device with pixel centers 45.72 μm apart. The pixels, or photo sites, exist as electron wells in the silicon substrate, and as such are not entirely discrete elements. The masks used in the sensor manufacture are computer generated at a large scale and are photo reduced onto the silicon surface. Manufacturing tolerances are extremely high, as the device will not function correctly if the photo sites are not regularly spaced. According to GE engineers,

the photon indeterminacy error will be of more significance than any pixel geometry error.

The sensor was placed under a measuring microscope at 100× magnification. A ruled stage with 1 division equal to 20 μm, and estimation to 2 μm, was used to measure the outside dimensions of the sensor and the center coordinates of randomly selected pixels. The measurements resulted in a mean pixel separation of 45.54 μm over the sensor surface, within 0.2 μm of the manufacturer's specifications. It was not determined whether this discrepancy could be attributed to the ruled stage, or to the sensor. In view of its small magnitude, the manufacturer's specified pixel separation was employed in all subsequent computations.

Calibration Design. A full bundle block adjustment program with added parameters for camera calibration was implemented, and the camera was calibrated at two fixed principal distances. One was chosen to allow direct close-range imaging of study casts and other small objects, while the other was designed for direct skull and cranial x-ray imaging. It should be noted that there are no inherent constraints on the principal distances that can be chosen or the sizes of objects that can be imaged. Practically, however, the present camera has low resolution and should be restricted to the imaging of smaller objects for the sake of obtaining reasonably accurate data. All calibration procedures developed and described here are general in nature and can be applied to larger format cameras.

The close-range calibration technique developed by Curry *et al.* (1981) was adapted for use with the array camera. A two-dimensional test field consisting of a planar target array was moved a precise distance parallel with itself along the Z axis to provide the third dimension of the test field. This method is particularly suitable for use with an array camera because multiple exposures are unlimited, and as many Z planes can be created as are necessary.

A number of questions were proposed in conjunction with the calibration of the array camera:

- How significantly would multiple positions of the test field improve the calibration results?
- Is lens distortion measurable in this low resolution camera?
- How significant is the sub-pixel determination of target points?
- Is the image coordinate weighting scheme significant?

These issues are examined further in the following sections.

Image Acquisition for Calibration. The small test field (Figure 2), for extreme close-range calibration, consisted of a Kodak micro-flat glass plate with a regular array of 100 holes drilled in the emulsion with a Wild PUG 80-μm drill. The larger test field (Figure 3) was constructed from a flat aluminum plate coated with black paint and containing an array of 63 targets consisting of 0.5-mm diameter white dots. The small plate was measured in a Kern MK2 monocomparator, with a least count of 1 μm. All

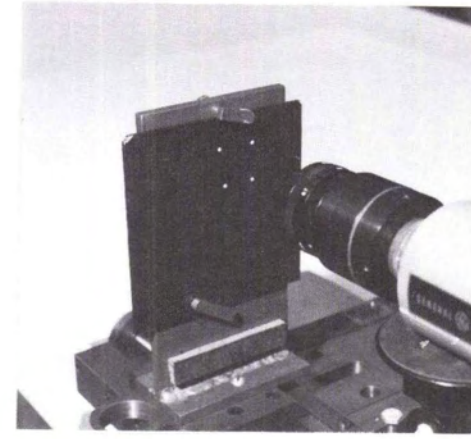

FIG. 2. Glass plate test field for extreme close range calibration. Note that individual targets are not visible.

points were measured three times, and the mean X and Y values were recorded. The large plate was measured on a Wild A5 coordinatograph table with a least count of 10 μm.

In both cases, the plates were mounted on an adjustable plate holder constructed from a lead-screw tape comparator. The plate could then be moved in the Z direction with a precision of 25.4 μm, with estimation to 2.5 μm. In addition, the plates could be shifted along the X axis. This movement was equivalent to displacing the camera in X, and provided a base for subsequent check point computations.

A theodolite was used to align the camera sensor and the target array using auto-reflection techniques (Figure 4). The camera was first aligned to the theodolite by reflecting the cross-hair image from the glass cover plate on the sensor. Then, the target array plate was placed in position and aligned with a first surface mirror. This technique allowed fixing the ω and ϕ rotation angles, which are highly correlated with the x and y coordinates of the principal point in the bundle adjustment.

At each position of the target array, an image was

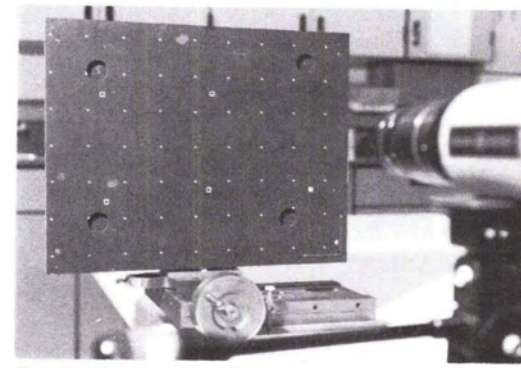

FIG. 3. Aluminum test field for close range calibration.

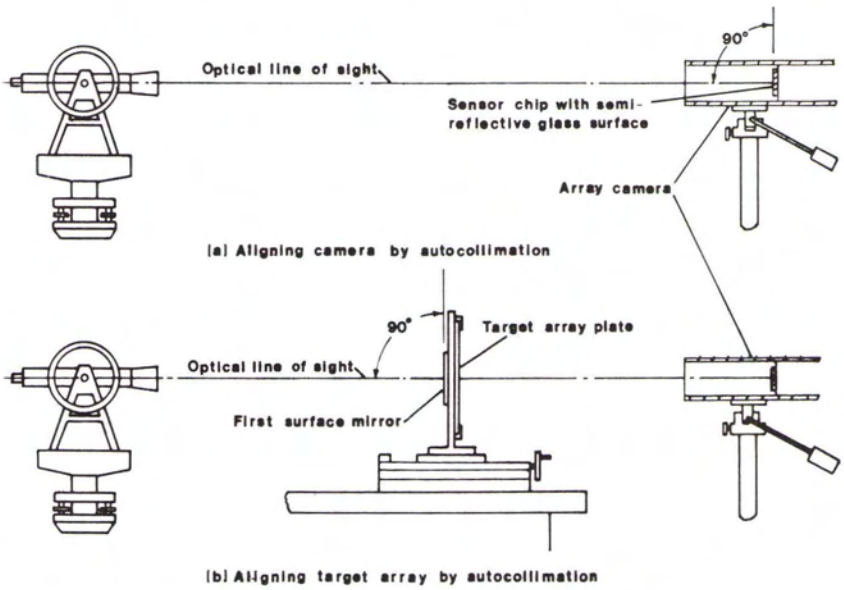

FIG. 4. Alignment of array camera and test field.

stored as the average of five consecutive frames. For the small array, the plate was located at both extremes of the depth of field, at object distances of 30 mm and 42.5 mm, and also in the middle of the range, for a total of three plate positions. The plate holder was then shifted approximately 15 mm in X to provide a base, resulting in a total of six images. Three plate positions were also acquired for the larger target array, at object distances of 300 mm, 338 mm, and 376 mm. It was displaced approximately 75 mm in X for the second group.

In addition, for each array, two additional experimental images were added. It was hypothesized that target point images that fell near pixel boundaries might be less accurate than those nearer the pixel centers, due to inter-pixel boundary problems. Therefore, for plate positions at both extremes of the depth of field, the target plate was shifted, in Z, the amount that would offset most targets about one half a pixel in x and y on the sensor surface. These images were called "delta" images, and were an attempt to ensure that, on one image out of the pair, each target would land near a pixel center, leading to improved sub-pixel point determination.

This type of experiment was made possible by the perfect multiple image registration that can be obtained with an array camera. In fact, it would be possible to approximate a "continuous" calibration, where the test field was as dense in Z as in X and Y.

After the rejection of bad pixels and the correction for dark current (see Radiometric Calibration), each frame was enhanced by convolving it with a Laplacian filter. The enhanced images were passed to a modified version of the Sequential Similarity Detection Algorithm (SSDA) (Barnea and Silverman, 1972). This algorithm locates potential target points in an efficient manner, using a method equivalent to correlation. At this stage, targets were reported only to integer pixel row and column values. Generally, 98 percent to 100 percent of all targets were found on the enhanced images. The next major step was the sub-pixel determination of the target centers.

Sub-Pixel Point Location. The data produced by the modified SSDA algorithm located potential targets to integer row and column pixels. Various investigators have examined the potential of locating targets in digital images to subpixel accuracy. Given that the target producing the image invokes an output signal from the detector covering more than one pixel, the spatial and intensity distribution of the signal can be modeled and the center of the target computed. The output signal, or image function, is the convolution of the input signal and the point spread function (PSF). Thurgood and Mikhail (1982) experimented with the sub-pixel locations of idealized targets with both rectangular and Gaussian point spread functions.

A preliminary investigation of the calibration data produced from images of the target array indicated that neither the triangular, rectangular, nor Gaussian distributions would be satisfactory, and two different methods were developed for the determination of the sub-pixel target centers. Although the two methods were approximate, they suited the low resolution of the data at hand.

Method 1 employed an empirical approach to modeling the image function of the targets. When the intensities were plotted along the x and y axes, they appeared to form parabolas of various

configurations. Software was developed to fit a parabolic curve to each target's output, independently on each axis, and to compute a peak on each axis. Frobin and Hierholzer (1983) used a somewhat similar approach in the analysis of their rasterstereographic data, except that they fit a parabola to the first derivative of the amplitude of the signal. Initial tests indicated that a parabolic curve fit directly to the signal amplitudes was more satisfactory in the calibration system.

The solution was determined separately in x and in y. This method assumed no correlation between x and y. If the target spanned two rows or two columns, two independent solutions were performed and the mean of the two peaks was computed. A scanning algorithm was developed to determine the starting and ending pixel of each target, and to determine how many rows and columns were covered. No specific parabola shape was enforced, because parabola shape could vary according to the exact target illumination.

The residuals from the parabolic fits indicated that the parabolic curve was a reasonable approximation. The sum of the squares of the residuals was passed to the calibration program to be used in an image coordinate weighting experiment.

In the second approach to sub-pixel determination, the image function was modeled in the x and y directions concurrently. An average parabolic curve was initially determined, based on the mean of the parabolae computed using the first method. This parabolic shape was then held fixed in all subsequent matching. A region of 3 by 3, 4 by 4, or 5 by 5 pixels surrounding the peak pixel was chosen, depending on the shape and size of the affected region. The average background signal in the region was computed and subtracted from the central 2 by 2 pixel sector of the selected patch. Then, the fixed parabolic curve was fit to the x and y image functions simultaneously, using a least squares solution.

The second approach had several advantages when compared to the first. By restricting the computations to the 2 by 2 central target region, errors due to rapid signal attenuation in the outer pixels were reduced. The joint determination of the x and y peak also strengthened the solution. However, to keep the number of variables to a minimum, the parabolic shape was not allowed to vary.

The image coordinates from sub-pixel location consisted of the pixel row and column location of all target images, expressed as tenths of pixels.

Point Matching. In order to proceed with the bundle adjustment, the image coordinates generated by the sub-pixel location software were matched with the control points on the target arrays. Ordinarily, point numbers would be assigned by a human operator during the digitizing phase. The point numbers would serve to match digitized points with the control file. In the present system, there was no human intervention in the point digitizing procedure, and point numbering and matching was performed automatically. Each point, as identified by the SSDA routine, and located to sub-pixel accuracy by the subsequent processing, was known only by its pixel coordinates and a point "type" identifier.

Unique identification points (ID) were placed on each of the target plates (see Figures 2 and 3). These points were sufficiently distinct in appearance from the other control points and from each other that the SSDA algorithm could identify and separate them from the output stream. A data file was assembled for each target plate position position containing the plate number, its Z coordinate at the time of image acquisition, the point type of the four ID points, the point number of each ID point, and their X and Y plate coordinates as determined in the comparator.

Software was written to open this file and match the four ID points in the file with the four ID points that the SSDA algorithm had identified in the input image coordinate file. If all four points were present and matched in point type, the four image points were transformed to the four plate points with an eight parameter transformation. The parameters from this transformation were then used to transform all remaining image points to the XY plate coordinate system. Each point's nearest neighbor was located on the plate and, if it was within a specified tolerance, the image coordinate was assigned the point number of the plate point. Suitable warnings were generated for points matched more than once, or not at all, but virtually 100 percent of the image points could be matched to plate points with a 1 or 2 pixel tolerance. Those that did not match were generally the result of noise in the image, and were discarded.

This procedure was also used to generate the object point coordinate files used as input to the bundle adjustment. Each plate point was assigned the Z coordinate of the plate position at which matching was being performed, and written to a file. In this fashion, a single XY coordinate date base for the target array was transformed into a three-dimensional data set covering all plate positions.

Bundle Adjustment. A general bundle adjustment program was written in the C language. The development was patterned after Anderson (1975) and Brown (1969). Parameters were added for principal point location, principal distance, and symmetric radial lens distortion. The added parameters were assumed to be block invariant. Symmetric radial lens distortion was modeled with orthogonal polynomials (Van Roessel, 1969).

The weight matrix for this system of equations contained the weights for image coordinates, object space coordinates, camera station parameters, and camera calibration parameters. All initial estimates for these values could be fixed by the assignment of large weights, or allowed to float by the assignment of small weights.

A number of calibration tests were performed with varying control configurations and image coordinate weights, with the objective of testing the hypotheses described earlier. Table 1 summarizes the results of

TABLE 1. CALIBRATION TEST RESULTS, GROUP 1

9-mm Lens plus Extender—Small Calibration Array								
Test	Code	c (mm)	xp (mm)	yp (mm)	σ_o	Deg. of freedom	RMSE XY (μm)	RMSE Z (μm)
1	2-v-s1-c	9.48	2.84	2.81	0.019	436	11.7	68.8
2	2-v-s1-cc	9.40	2.81	2.82	0.017	433	12.8	62.5
3	2-u-s1-c	9.41	2.82	2.84	0.014	436	9.4	60.0
4	2-u-s1-cc	9.42	2.82	2.84	0.014	433	9.4	60.0
5	2-v-i-c	9.72	2.85	2.80	0.034	436	27.5	106.0
6	2-v-i-cc	9.51	2.79	2.82	0.030	433	27.5	108.8
7	3-v-s1-c	9.49	2.82	2.80	0.025	643	17.4	57.5
8	3-v-s1-cc	9.41	2.80	2.81	0.023	640	18.5	62.5
9	3-u-s1-c	9.33	2.81	2.83	0.017	640	15.0	54.9
10	3-u-s1-cc	9.34	2.81	2.83	0.017	640	15.1	55.1
11	3-v-i-c	9.73	2.87	2.77	0.038	643	25.8	79.3
12	3-v-i-cc	9.56	2.82	2.81	0.033	640	25.2	91.3
13	4-v-s1-c	9.46	2.88	2.84	0.019	859	11.9	68.7
14	4-v-s1-cc	9.44	2.89	2.84	0.018	856	12.3	64.5
15	4-u-s1-c	9.52	2.85	2.88	0.013	859	8.9	52.3
16	4-u-s1-cc	9.53	2.85	2.87	0.013	856	8.9	52.3
17	4-v-i-c	9.31	2.89	2.85	0.035	859	22.1	70.6
18	4-v-i-cc	9.32	2.89	2.84	0.035	856	23.2	74.3

TABLE 2. CALIBRATION TEST RESULTS, GROUP 2

9-mm Lens—Large Calibration Array								
Test	Code	c (mm)	xp (mm)	yp (mm)	σ_o	Deg. of freedom	RMSE XY (μm)	RMSE Z (μm)
19	2-v-s1-c	8.89	2.87	3.01	0.012	308	10.3	48.2
20	2-v-s1-cc	8.88	2.87	3.01	0.012	305	10.0	48.2
21	2-v-i-cc	8.86	2.86	2.98	0.023	305	27.5	93.7
22	2-u-s2-c	8.89	2.84	2.97	0.019	306	12.7	50.3
23	2-u-s2-cc	8.90	2.83	2.97	0.019	303	12.7	50.3
24	2-u-s1-cc	8.81	2.85	2.99	0.018	305	11.4	56.5
25	3-v-s1-c	8.91	2.87	3.00	0.014	459	11.3	49.0
26	3-v-s1-cc	8.91	2.86	3.01	0.014	456	10.8	49.0
27	3-v-i-cc	8.82	2.91	2.98	0.027	456	25.3	82.8
28	3-u-s2-c	8.88	2.83	2.97	0.019	460	8.7	50.3
29	3-u-s2-cc	8.88	2.83	2.97	0.019	457	8.7	50.8
30	3-u-s1-cc	8.82	2.84	2.99	0.022	456	11.0	56.6
31	5-v-s1-c	9.02	2.82	3.04	0.021	587	12.5	50.3
32	5-v-s1-cc	8.97	2.83	3.04	0.020	584	11.0	50.3
33	5-v-i-cc	8.82	2.79	3.01	0.037	584	27.0	96.3
34	5-u-s2-c	8.80	2.86	2.96	0.010	597	8.0	46.8
35	5-u-s2-cc	8.80	2.86	2.96	0.010	594	8.0	47.0
36	5-u-s1-cc	8.84	2.85	2.97	0.017	584	9.74	49.9

the tests using the 9-mm lens with extender for extreme close-range work, while Table 2 contains the results obtained with the 9-mm lens alone.

Within each table, the tests are identified with a code as follows:

First field = number of control plate positions.
Second Field = image coordinate weighting scheme.
v = weights proportional to variances from parabolic fitting (sub-pixel method 1 only).
u = unit weights.
Third field = sub-pixel location method.
1 = independent x and y parabolas.
2 = simultaneous x and y parabolas with average background subtracted.
i = integer x and y pixel locations.
Fourth field = number of camera calibration parameters.
c = principal distance, principal point.
cc = as above, plus symmetric radial lens distortion.

In all subsequent discussions, the extreme close-range group of calibrations is termed "Group 1" and the medium close-range set is termed "Group 2." Not all possible combinations of tests were performed. The image coordinate weighting tests were used only to compare unit weighting and the variance-related weights from the first sub-pixel computation method. The second sub-pixel method was tested only with Group 2. The results of Group 1 made clear that the unit pixel locations (no sub-pixel computation) were substantially less accurate, as would be expected. Thus, only a few of the unit pixel results for Group 2 are reported here.

Ten randomly selected check points were withheld from the calibration and used to estimate system performance in X, Y, and Z. The root mean square error on position (RMSE XY) and height (Z) are reported in micrometres at the sensor scale. Note that one pixel equals 45.7 μm. It should also be noted that the base/height ratio is poor for both groups of calibrations due to constraints on plate movement (ΔX) within the plate positioner. It is approximately 0.4 for the Group 1 calibrations, and 0.22 for the Group 2 tests, when computed using the distances to the middle position of the target array. Therefore, the RMSE Z values will be somewhat higher than they would have been if the base between plate positions had been free of constraints.

For each calibration, the reference variance was computed. Assuming that the adjustments are all independent, these reference variances can be statistically tested for homogeneity and equivalence. The reference variances should not be homogeneous within groups, because the calibration configurations were deliberately chosen to test the different hypotheses described earlier. However, the non-homogeneity must be established before further testing is performed. Bartlett's test allows the homogeneity of the variances to be tested (Hamilton, 1965). At the 0.01 significance level, the calibration reference variances were not homogeneous.

Given that the sample reference variances are s_1^2 and s_2^2, from two normal populations with variances σ_1^2 and σ_2^2, and independent random samples of size n_1 and n_2, then the F test can be used to test the significance of differences in the computed reference variances. The degrees of freedom are $v_1 = n_1 - 1$ and $v_2 = n_2 - 1$, and the F statistic is

$$F_{v_1 v_2} = \frac{s_1^2}{s_2^2}$$

To simplify testing, the values were sorted so that s_1^2 was always greater than s_2^2, and a one-sided F test was employed. The null hypothesis was $H_0 : \sigma_1^2 = \sigma_2^2$, and the alternative hypothesis was $H_1 : \sigma_1^2 > \sigma_2^2$. The null hypothesis was rejected if $F_{v_1 v_2} > F_{\alpha, v_1, v_2}$. The results of these tests are arranged in Tables 3 and 4. The tables should be read as follows:

(1) Select the first calibration across the top row of the table. The calibration numbers refer to Tables 1 and 2.
(2) Select the comparison calibration down the first column of the table.
(3) If the table symbol is ".", there is no significant difference between the reference variances of the two calibrations. If the symbol is ">", the first calibration has a significantly greater reference variance; if it is "<", it is significantly less. All tests were performed at the 0.01 significance level.

The reference variance test matrices in Tables 3 and 4 can be used to compare individual calibrations within a group, or an entire row or column of the matrix can be examined to determine how a particular calibration ranked overall. For example, in Group

TABLE 3. CALIBRATION SIGNIFICANCE TESTS—GROUP 1

Reference Variance Significance Matrix—Group 1

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
1		.	<	<	>	>	>	>	.	.	>	>	.	.	<	<	>	>
2	.	.	.	>	>	>	>	.	.	>	>	.	.	<	<	>	>	
3	>	.	.	>	>	>	>	.	.	>	>	>	>	.	.	>	>	
4	>	.	.		>	>	>	>	.	.	>	>	>	>	.	>	>	>
5	<	<	<	<		.	<	<	<	<	.	.	<	<	<	<	.	.
6	<	<	<	<	.		<	<	<	<	>	.	<	<	<	<	.	.
7	<	<	<	<	>	>		.	<	<	>	>	<	<	<	<	>	>
8	<	<	<	<	>	>	.		<	<	>	>	<	<	<	<	>	>
9	.	.	.	.	>	>	>	>		.	>	>	.	.	<	<	>	>
10	.	.	.	.	>	>	>	>	.		>	>	.	.	<	<	>	>
11	<	<	<	<	.	<	<	<	<	<		.	<	<	<	<	.	.
12	<	<	<	<	.	.	<	<	<	<	.		<	<	<	<	.	.
13	.	.	<	<	>	>	>	>	.	.	>	>		.	<	<	>	>
14	.	.	<	<	>	>	>	>	.	.	>	>	.		<	<	>	>
15	>	>	.	.	>	>	>	>	>	>	>	>	>	>		.	>	>
16	>	>	.	.	>	>	>	>	>	>	>	>	>	>	.		>	>
17	<	<	<	<	.	.	<	<	<	<	.	.	<	<	<	<		.
18	<	<	<	<	.	.	<	<	<	<	.	.	<	<	<	<	.	

TABLE 4. CALIBRATION SIGNIFICANCE TESTS - GROUP 2

Reference Variance Significance Matrix - Group 2

	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36
19		.	>	>	>	>	.	.	>	>	>	>	>	>	>	<	<	>
20	.		>	>	>	>	.	.	>	>	>	>	>	>	>	<	<	>
21	<	<		<	<	<	<	<	>	<	<	.	.	.	>	<	<	<
22	<	<	>		.	.	<	<	>	.	.	.	.	.	>	<	<	.
23	<	<	>	.		.	<	<	>	.	.	.	.	.	>	<	<	.
24	<	<	>	.	.		<	<	>	.	.	>	.	.	>	<	<	.
25	.	.	>	>	>	>		.	>	>	>	>	>	>	>	<	<	>
26	.	.	>	>	>	>	.		>	>	>	>	>	>	>	<	<	>
27	<	<	<	<	<	<	<	<		<	<	<	<	<	>	<	<	<
28	<	<	>	.	.	.	<	<	>		.	.	.	.	>	<	<	<
29	<	<	<	.	.	.	<	<	>	.		.	.	.	>	<	<	<
30	<	<	.	.	.	<	<	<	>	.	.		.	.	>	<	<	<
31	<	<	.	.	.	.	<	<	>	.	.	.		.	>	<	<	.
32	<	<	.	.	.	.	<	<	>	.	.	.	.		>	<	<	.
33	<	<	<	<	<	<	<	<	<	<	<	<	<	<		<	<	<
34	>	>	>	>	>	>	>	>	>	>	>	>	>	>	>		.	>
35	>	>	>	>	>	>	>	>	>	>	>	>	>	>	>	>		>
36	<	<	>	.	.	.	<	<	>	>	>	>	>	.	>	<	<	

1, the reference variance for Test 5 is significantly larger than all tests except 6, 11, 12, 17, and 18 (see Column 5, Table 3). Referring to Table 1, one notes that test 5 was performed using integer row and column pixels and that the variance of this test is equivalent only to other tests using row and column pixels. Thus Tables 3 and 4 can be used to provide statistical verification of the results given in Tables 1 and 2.

In general, the calibrations were successful and stable. In the best case, in Group 2, the RMSE in position was as low as 0.2 pixel on the sensor surface. When the results are analyzed in the light of the hypotheses posed earlier, several points become apparent.

- The sub-pixel determination of target location is very important. The use of integer pixel locations results in significantly poorer results in all cases, as would be expected. In Group 2, sub-pixel method number 2 improved the results when compared to method number 1, but only as the number of plates increased.
- The addition of symmetric lens distortion coefficients improves the results only slightly, and never significantly at the 0.01 significance level. The most improvement is evident on the integer pixel location plates. Evidently, the resolution of the camera is too low for the determination of lens distortion.
- The use of unit weights for image coordinates is significantly better in nearly all instances. The variance related weighting model is wrong, and should not be employed. The goodness of fit in the parabolic fitting process is not a measure of the accuracy of point location. In fact, this result may indicate that the entire parabolic fitting algorithm needs improvement.
- The inclusion of a single additional control plate position in the middle of the object space range does not improve the results significantly in most instances. However, the addition of the "delta" frames improves the results in most cases, especially when employed in conjunction with the second sub-pixel location algorithm. This supports the hypothesis stated earlier and verifies that targets located near pixel centers have higher sub-pixel accuracy.

In summary, the best results are obtained with unit image coordinate weights, using the second sub-pixel method, with two "delta" positions of the control field. The calibrated camera parameters from these runs for both Group 1 and Group 2 were used in all subsequent data acquisition and analysis. Future work will involve the development of improved sub-pixel location algorithms and improved image point weighting schemes. Also, the addition of parameters to the bundle adjustment will be investigated, to determine whether some of the array camera parameters can be more directly modeled. Data from dental study casts are now being collected and analyzed to determine system performance with real objects.

ACKNOWLEDGMENTS

This project was supported by NIH National Institute of Dental Research Grants DE03598 and DE03703. Sean Curry also acknowledges the Wild Heerbrugg Company for their fellowship support of his graduate studies, and support from the UC Berkeley Patent Fund.

REFERENCES

Anderson, J.M., and C. Lee, 1975. In-Flight Calibration of an Aerial Camera, *Photogrammetric Engineering and Remote Sensing*, Vol. 61, No. 11, pp. 1337–1348.

Barnea, D.T., and H. Silverman, 1972. A Class of Algorithms for Fast Digital Image Registration, *IEEE Transactions on Computers*, Vol. C-21, pp. 179–186.

Baumrind, S., and S. Curry, 1984. Merging of Data from Different Records in Craniofacial Research and Treat-

ment, *Close Range Photogrammetry and Surveying Symposium*, American Society of Photogrammetry, pp. 35–46.

Brown, D., 1969. *Advanced Methods for Calibration of Metric Cameras*, 1969 ASP Symposium on Computational Photogrammetry, Syracuse, NY.

Curry, S., J.M. Anderson, and F.H. Moffitt, 1981. Calibration of Extreme Close-Range Camera. *ASP Proceedings, pp. 311–324.*

Curry, S., and S. Baumrind, 1984. Analysis of Stereo Cranial X-Rays using Digital Images, *Close Range Photogrammetry and Surveying Symposium*, American Society of Photogrammetry, pp. 35–46.

Fairchild CCD Imaging, 1982. *CCD—The Solid State Imaging Technology*, Fairchild Instrument Co.

Frobin, W., and E. Hierholzer, 1983. Automatic Measurement of Body Surfaces Using Rasterstereography, *Photogrammetric Engineering and Remote Sensing*, Vol. 64, No. 10, pp. xxx.

Gruen, A.W., 1983. *A Test Strategy for High Resolution Image Scanners*, Report # 350, Ohio State University, Columbus, Ohio.

Hamilton, W.C., 1965, *Statistics in the Physical Sciences*, Ronald Press.

Rosenfeld, A., and A. Kak, 1982, *Digital Picture Processing*, Academic Press, N.Y., N.Y.

Thurgood, J.D., and E.M. Mikhail, 1982, *Subpixel Mensuration of Photogrammetric Targets in Digital Images*, Report CE-PH-82-2, Purdue University.

Van Roessel, J.W., 1969, *Calibration of a Metric Camera*, MS Thesis, University of California, Berkeley.

(Received 1 July 1985; revised and accepted 21 January 1986)

ANNOUNCEMENT AND CALL FOR PAPERS

International Colloquium on Network Design Optimization in Non-Topographic Photogrammetry

Baltimore, Maryland
31 March – 2 April 1987

This Colloquium — organized by ISPRS WG V/1 and sponsored by the American Society for Photogrammetry and Remote Sensing (ASPRS) and the International Society for Photogrammetry and Remote Sensing (ISPRS) — will have as its theme "Network Design Optimization in Non-Topographic Photogrammetry." Papers are invited in the following areas which correspond to the terms of reference of WG V/1:

- Prediction and assessment of reliability and precision of non-topographic systems
- Design of optimal systems considering both technical and economic aspects
- Accuracy improvement of close-range photogrammetric systems
- Enhancement of software capabilities
- Application aspects of on-line point positioning

The Colloquium will be conducted as a component of the 1987 ASPRS-ACSM Annual Convention.

Those wishing to present a paper should submit an abstract by *1 September 1986* to

H. M. Karara
University of Illinois
2215 Newmark Civil Engineering Laboratory
208 North Romine Street
Urbana, IL 61801