by Phillip Reu

Introduction to Digital Image Correlation: Best Practices and Applications

INTRODUCTION

I am pleased to be able to write this new article series in *Experimental Techniques*. This journal is aimed at the *practicing* technologist, engineer, or researcher. It is my goal to make this series *practical*. I will do this by disseminating the best current Digital Image Correlation (DIC) research available and boiling it down to "what does it mean for me in the lab?" Along with consultation with the foremost practitioners of DIC and the developers who are writing the codes, I will also draw on my experience applying DIC to challenging measurement problems at Sandia National Laboratories. Applications at Sandia have covered a wide range of experiments including electron microscopes and stereo-microscopes and large-scale application with fields of view of tens of meters; at low speed and at high- and ultra-high speed (1 million frames per second). As with any experimental technique, there is an "art" to conducting the experiment correctly, for example speckling and lighting the object. As the article series progresses, I hope that it will become a guide to "Best Practices" in the application of DIC. This guide will hopefully aid both the long time users and those who are new to the field: graduate students, technologists, and others beginning to work in the field of experimental mechanics.

Typical outdoor DIC application at Sandia, showing interlaced DIC stereo-rigs.

[†]Early DIC References

1. Lucas, B.D., and Kanade, T., in *Proceedings of the 1981 DARPA Imaging Understanding Workshop*, pp. 121–130 (1981).
2. Peters, W., and Ranson, W., *Optical Engineering* **21**(3):427–431 (1982).
3. Sutton, M.A., et al. *Image and Vision Computing* **1**(3):133–139 (1983).
4. Bruck, H., et al. *Experimental Mechanics* **29**(3):261–267 (1989).

[‡]Sutton, M.A., Orteu, J.J., and Schreier, H.W., *Image Correlation for Shape, Motion and Deformation Measurements*, Springer, New York (2009).

DIGITAL IMAGE CORRELATION: A BRIEF HISTORY AND DEFINITION

DIC is the most important advance in experimental mechanics since the strain gage. Why? Because imaging has always been important to science. Inventions such as the telescope and the microscope that allowed scientists to "see" new things for the first time, invariably led to an explosion of scientific and engineering knowledge. This is true of DIC as well, where the deformation (strain) of an object can be visualized. Historically, the first practical form of DIC came about in the 1980s[†] because of the advent of the digital camera with a sensor plane that is uniform and the affordable processing power of the emergent personal computer. The DIC concepts, however, did not appear out of nowhere. They were a logical development of the earlier and related photogrammetry techniques used for aerial imagery, robotic vision, and microscopy. With the newly available digital cameras, it was a natural next step to apply the cameras in an experiment by first using coherent light, for holography and digital speckle pattern interferometry, and then later using an applied painted speckle pattern.[‡] Once this step was taken,

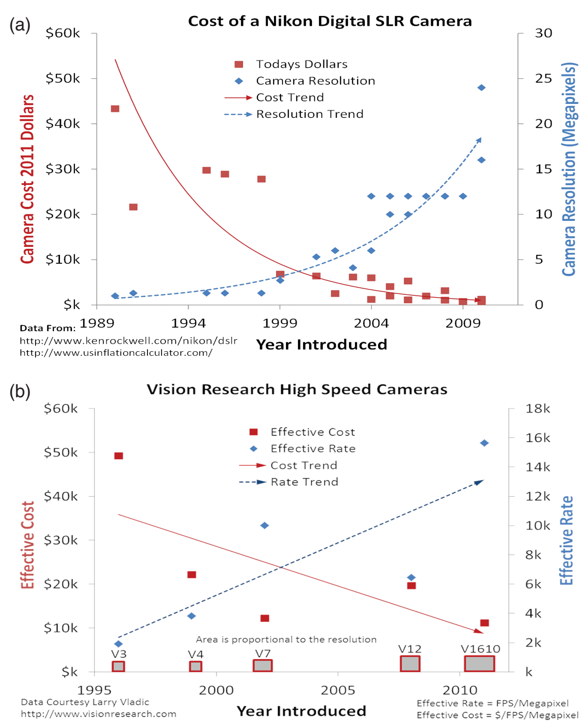

The cost for digital cameras has steadily decreased through time; both for machine vision (a) and high-speed cameras (b).

The Art and Application of Digital Image Correlation is written by Phillip L. Reu (Phillip.Reu.DIC@gmail.com). He received his PhD from the University of Wisconsin-Madison and is currently a Principal Member of Technical Staff at Sandia National Laboratory. Phillip specializes in the development of novel full-field measurement techniques for collecting data in previously un-measurable regimes. He began working with digital image correlation in 2004 and has applied it in the field to a wide range of experiments including: quasi-static, micro-scale, multi-system, large fields-of-view, and ultra-high speed. His image correlation research is focused on understanding the effect of the unavoidable compromises made in field measurements to the final DIC uncertainty.

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract No. DE-AC04-94AL85000.

doi: 10.1111/j.1747-1567.2011.00798.x

most of the familiar techniques used in DIC have followed naturally; subset matching, gray-level interpolation, and stereo camera calibration for triangulation. These have developed into the three main branches of DIC: for in plane, 2D-DIC; for x, y, and z data, 3D-DIC, and for measurements within a volume, V-DIC. The rapid adoption and spread of DIC has been greatly assisted by the phenomenal advances in imaging technologies over the last 20 years, including high-resolution cameras (14 megapixels are now common) and high-speed cameras (5000 to millions of frames per second). Furthermore, the great flexibility of the mathematical concepts underlying both 2D- and 3D-DIC enable it to be used with methods that are not imaging in the traditional optical sense (e.g., atomic force microscope [AFM] applications, where the image is normally a height map; scanning electron microscope [SEM] applications; and X-ray and computed tomography [CT]).

Commercial DIC codes

- ARAMIS—Trilion (GOM Gbmh)
- ISTRA 4D (Q-4xx)—Dantec Dynamics
- StrainMaster—LaVision Inc.
- Vic3D—Correlated Solutions

University codes (short list)

- MatchID—Catholic University College Ghent, KULeuven
- Opticist—The Catholic University of America
- Matlab—Karlsruhe Institute of Technology (KIT) and Johns Hopkins University

SOME CHALLENGES WITH WRITING THE ARTICLE

The goal is to make this series as general as possible; however, references to various commercial and university software codes will be necessary at times. I see the greatest hurdle being that many of the commercial codes are now functioning as “black boxes” where the images are put in and results are cranked out. This is valuable for the user, but without input from the software companies important aspects of the process (i.e., recommended calibration methods, subset sizes, and speckle patterns) can be difficult to discuss. I believe that many of the best practices will be universal to all the codes because all the underlying *principles* are similar. Therefore, a good speckle pattern for one code *should* be the same for all codes, whether subset based or not. At a fundamental level, this is true because the quality of the correlation is a function of the quantity of information contained in an image. An image with a large quantity of information will render better results than one with low information content. Information in this context is defined by the contrast and number of edges in an image.

Recently, a new type of DIC software based on a finite element approach is being developed that solves the entire field of view simultaneously,[†] rather than the subset approach used by the earliest researchers and the commercial codes. Notwithstanding, many of the principals involved here will also apply to full-field DIC methods (e.g., optimum speckle).

[†]Hild, F., and Roux, S., *Strain* **42**(2):69–80 (2006).

[‡]phillip.reu.DIC@gmail.com is my contact e-mail.

The “hidden components” of DIC

1. Calibration	(3D and maybe 2D)
2. Subset shape function	(both 3D and 2D)
3. Gray level interpolation	(both 3D and 2D)
4. Subset matching	(both 3D and 2D)
5. Triangulation	(3D only)
6. Post-processing	(both 3D and 2D)

CONCLUSION

With this introduction, I would like to invite you to join with me in developing “DIC Best Practices.” Please e-mail me article ideas and questions you would like answers to.[‡] I will consult with other experts in the field and attempt to present practical answers to your questions to improve everyone’s DIC capabilities. Together we can make this a community forum where great ideas can be shared to improve and extend the use of DIC.

I would also like to direct the reader’s attention to an excellent book on DIC published by Sutton, Orteu, and Schreier, titled *Image Correlation for Shape, Motion, and Deformation Measurements*. It is available from Springer press. Many concepts presented here will be drawn from this book.[‡]

The next articles in this series will introduce the “Hidden Components” of DIC.

Hidden Components of DIC: Calibration and Shape Function—Part 1

by Phillip Reu

Introduction

The general steps required to conduct digital image correlation (DIC) are often laid out in the commercial software in a logical flow: obtain calibration images, calibrate stereo-rig (3D), obtain speckle images, and run analysis. For the user, the greatest effort is setting up the stereo-rig, speckling the sample, optimizing the lighting, and calibrating the system. Buried in these mundane experimental steps are the ''hidden components'' of 2D- and 3D-DIC. These components are the software processes that take the experimental images and turn them into shape and displacement data. The hidden components are what answer the questions: What makes a great DIC image? When is my system calibrated? And is my data good? Because these hidden components are so important to optimizing the experiment, I will briefly introduce them in the next three articles. With the exception of triangulation and possibly calibration, all the components are common to both 2D and 3D-DIC and function in similar ways for both. Calibration is unique because it is often skipped in 2D when the preferred result is strain.† Strain is defined as a change in length divided by the original length (there are a number of different strain definitions, but this concept holds true for all of them). The units of strain (ε) are length/length, which makes strain unitless. Because it is unitless, any length dimension, whether millimeters or pixels, divides out making calibration unnecessary. However, if you want to measure the 2D displacement or the physical dimensions in units other than pixels, a camera calibration will be necessary to relate the pixel size in the camera to a physical length scale at the sample.

Calibration

A stereo-rig is defined as two or more cameras that are calibrated as a unit to take 3D-DIC measurements. Calibration is the method by which the DIC software orients and scales the camera images to the physical world. The calibration process does this by determining the *intrinsic* and *extrinsic* camera parameters. The intrinsic parameters are those that relate a single camera image to the world. The extrinsic parameters are what relate the cameras to one another and are required

The Art and Application of Digital Image Correlation is written by Phillip L. Reu. He received his PhD from the University of Wisconsin–Madison and is currently a Principal Member of Technical Staff at Sandia National Laboratory. He began working with digital image correlation in 2004 and is focused on understanding the influence of the unavoidable compromises made in field measurements on the final DIC uncertainty. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract No. DE-AC04-94AL85000.

General steps required for DIC

1. Obtain calibration images.
2. Calibrate stereo-rig.
3. Obtain speckle images.
4. Run analysis.

The ''hidden components'' of DIC

1. Calibration	**(3D and Maybe 2D)**
2. Subset shape function	**(Both 3D and 2D)**
3. Gray level interpolation	(Both 3D and 2D)
4. Subset correlation	(Both 3D and 2D)
5. Triangulation	(3D only)
6. Post-processing	(Both 3D and 2D)

†2D calibration is important if your camera does not have square pixels (becoming rare); or when the optical path contains lenses or viewports with large optical distortions.

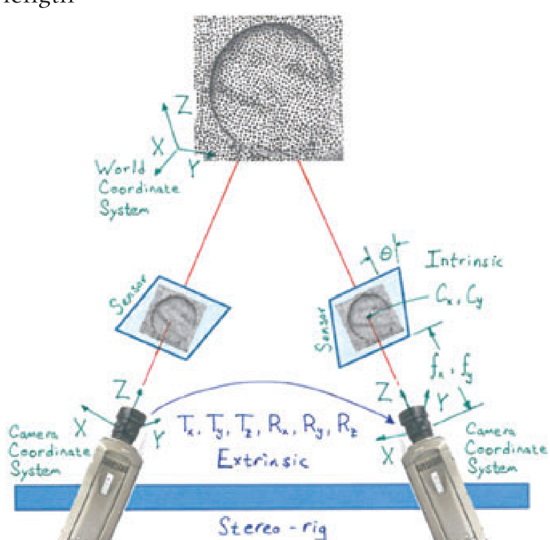

Stereo-rig indicating the calibration parameters.

Intrinsic Parameters (each camera)

1. Focal length (f_x, f_y)
2. Image plane center (C_x, C_y)
3. Image skew angle (θ)
4. Lens distortions ($K_{1\ldots n}$, radial only)

Extrinsic Parameters (stereo-rig)

1. Translation between the cameras (T_x, T_y, and T_z)
2. Angle between the cameras (R_x, R_y, and R_z)

for triangulation while leaving the results in the camera coordinate system. The intrinsic parameters include: the sensor center point, lens focal length, camera skew, and any lens distortions. The extrinsic parameters are the cameras relative orientations and *x*, *y*, and *z* positions in space. There are two important benefits of the independence of the intrinsic and extrinsic calibration parameters: The first is that the cameras can be calibrated individually for the intrinsic parameters; then oriented to each other in an additional calibration step. This is very helpful when measuring very large fields-of-view (meters). The second helps with the fairly common situation of having the cameras jostled between the calibration and the test time. If the intrinsic parameters remain unchanged, which is typical, the extrinsics can be recalculated using an orientation calibration.

As a historical note, the *bundle adjustment* calibration methods[a] are what transformed DIC from a laboratory curiosity into a commercial enterprise. Early DIC implementations required a tedious series of camera translations and measurements that made the calibration process too difficult for commercial success. The bundle adjustment process simplified the calibration by allowing the user to acquire a series of images of a calibration object (often a plate with dots or a grid pattern) in different poses. From these images it then calculates all of the camera parameters using a minimization method.

Subset Shape Function

Much like in finite element (FE) simulations, DIC software breaks the image into a number of subsets (or facets) for analysis. Unlike FE these subsets are independent and have no enforced continuity between them.[‡] These square subsets, defined in the reference image, are then sought in all the subsequent images of both cameras. The square subset defined in the reference image *will* be a different shape in the corresponding stereo camera image[*] and *will* change shape as the sample is deformed. To accommodate this shape change and motion a mathematical *shape function* is required. For most applications, the subset must be able to be moved for rigid body translation and warped sufficiently to match the deformed shape or the projection of the subset from one camera to the other. The affine shape function is the simplest function that meets these requirements and is the most commonly used shape function because of its computational efficiency. It does have limitations in that it can only represent a linear function: translation, stretch, and shear. Also, because it is linear only flat object shapes and linear deformation gradients can be perfectly represented. For many (most?) experiments this is an adequate representation. However, in cases where the deformation gradients are large or in highly curved objects, the linear assumption will result in errors. Examples where this is certainly the case include: deformation at grain boundaries, stress concentrations, and material discontinuities. For these situations, a higher order shape function would be more appropriate and more capable of capturing the actual deformation field within the subset. When the user cannot change the shape function (or it is unknown), the errors due to a shape function that does not accurately describe the deformation field

[a] Triggs, B., et al. Bundle Adjustment-A Modern Synthesis.

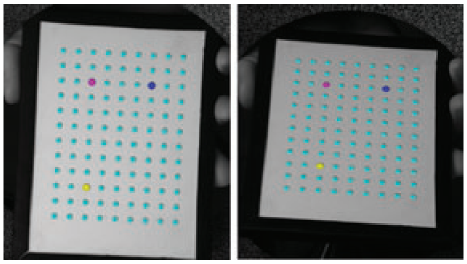

Left and right camera calibration images.

[‡] This is of course not true for the ‘‘global methods’’ which use the entire Region-of-Interest (ROI) to calculate the results: displacement, strain, etc. and will be discussed in a future article.

[*] In some cases the software will transform (rectify) the image from the left to the right camera using perspective projection. This is important when using the affine shape function which cannot adequately represent this distortion.

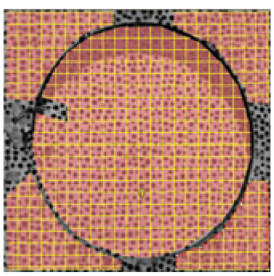

ROI broken into subsets with zero subset overlap.

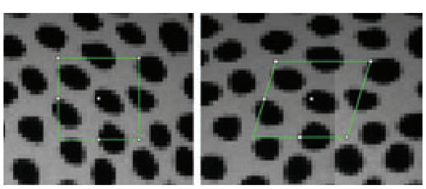

An affine shape function demonstrating its ability to represent translation, strain, and shear.

can be minimized by choosing small subsets. In order to maintain an acceptable noise level when using small subsets, it is important to maximize the information content in the image by applying a speckle pattern with high contrast and small speckles. To ensure an appropriately sized speckle, both the white and black areas need to be at least 3–5 pixels in size. When this is all done properly, a subset size of around 20 pixels can be used and often gives a good compromise between spatial resolution and matching noise. Subset size selection and speckling are complicated topics that will be covered in greater detail in future articles.

Next Time: A discussion of interpolation functions and subset matching.

Hidden Components of 3D-DIC: Interpolation and Matching—Part 2

by Phillip Reu

Introduction

Digital image correlation (DIC) requires three steps: calibrating the stereo-rig (in case of 3D-DIC), obtaining the deformed images, and running the analysis. Determining the best method for conducting these steps requires definition of the DIC ''hidden components.'' The last article covered calibration and subset shape functions. This article will discuss interpolation and subset matching. The purpose is to introduce the reader to DIC terminology and the function of each component. This will help with future articles as we discuss the practical aspects for conducting quality full-field measurements.

The ''hidden components'' of DIC

1. Calibration	(3D & Maybe 2D)
2. Subset shape function	(Both 3D & 2D)
3. **Grey level interpolation**	(Both 3D & 2D)
4. **Subset matching**	(Both 3D & 2D)
5. Triangulation	(3D only)
6. Post-processing	(Both 3D & 2D)

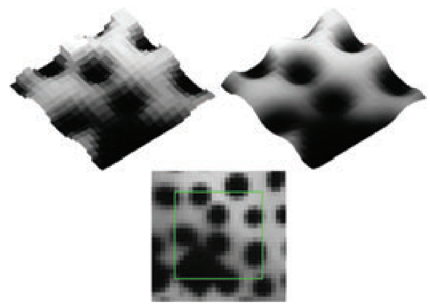

Raw image data (left) and interpolated data (right) for a subset (bottom). The grey levels on the right are the interpolated values ''between'' the pixels.

Interpolation

The ''magical''capabilities of DIC come from its ability to supply results that are *subpixel*. This ability to reliably resolve displacements of $1/100^{th}$ of a pixel enables high-resolution displacement measurements of relatively large fields-of-view. For example, a 5-Megapixel camera looking at a 100-mm sample would have a displacement resolution of approximately 1.5 microns[1]. Without this capability, DIC would be of limited use in experimental mechanics. As with all of the components of DIC, decisions regarding what method of interpolation to use must be made. The choice of interpolation method is important because it has been shown that not all methods are equal[2] and the decision usually reduces to computation time versus interpolation quality. Over the years a number of different interpolation schemes have been used including linear and cubic-polynomial. These two, because of their much larger errors, are no longer recommended. So what should you use? A quality interpolant is defined by its ability to suppress both the interpolation bias and the noise bias (illustrated on the right). Interpolation bias is seen at the subpixel level by a ''sinusoidal'' shaped error in the solution, with the lowest errors at the integer and ½-pixel locations. The noise bias is a linear error that results from the pixel level image noise. In order to improve the strain results, these two error sources must be suppressed. To accomplish this goal, a ''higher order'' more computationally expensive interpolant is required, with a cubic B-spline being the simplest interpolant that will supply acceptable[3] results. Unfortunately, with many software codes,

[1] For 1/30 pixel accuracy.

[2] Schreier, H.W., Braasch, J.R., and Sutton, M.A., Optical Engineering, 2000. **39**(11): p. 2915–2921. and Wang, Y.Q., et al., Strain, 2009. **45**(2): p. 160–178.

[3] ''Acceptable'' is an inexact measure of the uncertainty of the measurement. The reader will note that no answer to this is attempted at this point.

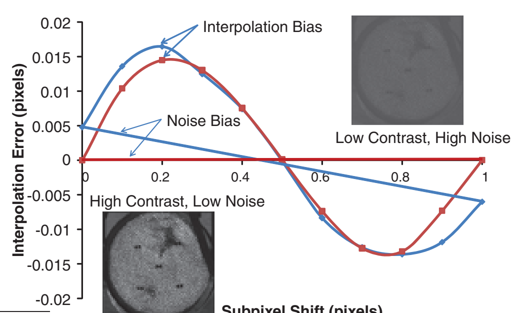

Popular Interpolation Methods

1. Linear (should not be used)
2. Cubic polynomial (marginal)
3. Cubic B-spline (adequate)
4. Optimized interpolation filter (best)
5. Quadratic quadrilateral (unknown)

The Art and Application of Digital Image Correlation is written by Phillip L. Reu. He received his PhD from the University of Wisconsin-Madison and is currently a Principal Member of Technical Staff at Sandia National Laboratory. He began working with digital image correlation in 2004 and is focused on understanding the influence of the unavoidable compromises made in field measurements on the final DIC uncertainty. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract No. DE-AC04-94AL85000. Email: phillip.reu.DIC@gmail.com

the user doesn't always know what interpolant is being used. A future article will discuss methods of assessing the matching quality of any software package via numerically shifted images. What is *always* true, regardless of the interpolant, is that a smoother speckle pattern with good contrast and low image noise will always yield better results.

Subset Matching

To determine the displacement of any subset, it is necessary to match a region of the reference image with a deformed image, or to match between the left and right cameras. As mentioned above, this will include subpixel information from the interpolation, as well as displacement and deformation from the shape function. The matching is done by minimizing a function that includes the pixel intensity information from both the reference image and the deformed images. The intensity values of the subset are correlated or matched with one another until a functional is minimized (or maximized). When a minimum is reached, you have found the displacement and deformation which represents the best match possible. Historically, the *correlation* in digital image correlation is derived from the name of the cross-correlation (CC) metric[4]. A related minimization function is the sum-squared difference (SSD); both of these should be used with care because they are unable to compensate for the nearly unavoidable changes in brightness and contrast[5] during a typical 3D experiment. Other methods, such as the normalized cross-correlation (NCC), normalized sum-squared difference (NSSD) and zero-normalized sum-squared difference (ZNSSD), are able to compensate for speckle and lighting changes and are better functions to use in nearly all cases.

Next Time: A look at Triangulation and Postprocessing.

Representative Minimization Functions

F is the reference subset grey values
G is the deformed subset grey values
χ is the minimization functional

Normalized Cross Correlation

$$\chi^2_{\mathrm{NCC}} = \frac{\sum_{\mathrm{pixels}} FG}{\sqrt{\sum_{\mathrm{pixels}} F^2 \sum_{\mathrm{pixels}} G^2}}$$

Normalized Sum Squared Difference

$$\chi^2_{\mathrm{NSSD}} = \sum \left(\frac{\sum_{\mathrm{pixels}} FG}{\sum_{\mathrm{pixels}} G^2} G - F \right)^2$$

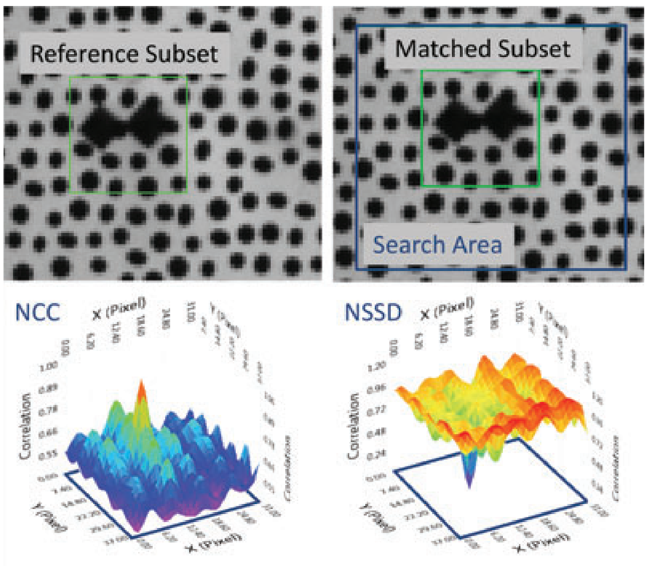

The peak (or valley) indicates the ideal match in the search area.

[4]This matching can be done either in the spatial domain or the frequency domain via the FFT; most implementations now use the spatial domain.

[5]*Brightness* is the overall intensity of the image in grey scale counts. *Contrast* is the difference in counts between the bright areas and the dark areas of the speckle pattern.

Hidden Components of 3D-DIC: Triangulation and Post-processing—Part 3

by Phillip Reu

Introduction

Digital image correlation (DIC) is a technique that uses digital camera images to measure shape and displacement. In order to do this, there are a number of ''hidden components'' that take place behind the scenes of the software interface. This article covers the final two: triangulation and post-processing.

The "hidden components" of DIC

1. Calibration	(3D & Maybe 2D)
2. Subset shape function	(Both 3D & 2D)
3. Grey scale interpolation	(Both 3D & 2D)
4. Subset matching	(Both 3D & 2D)
5. **Triangulation**	**(3D only)**
6. **Post-processing**	**(Both 3D & 2D)**

Triangulation

At this point in the analysis, each subset in the image has been matched relative to the reference image. This includes any translation and deformation of the subset relative to the undeformed (by definition) reference image, and the matching subset that is both translated and deformed in the corresponding stereo image. Now the calibration information can be used for triangulation. The intrinsic parameters allow a pixel in the camera image to be projected back out into space. With only one camera this does not tell you how far from the camera the object is. But with two cameras, using the extrinsic stereo parameters, the two projected lines will cross somewhere in space. The triangulation process unambiguously locates the lines projected from the centers of the corresponding subsets to a unique 3D point in space. The difficulty is that the lines do not intersect! Consequently, there have been a number of methods created to minimize this intersection error[1]. To find the optimum triangulation point, an assumption is usually made that the calibration is correct and the lack of intersection is caused by a subset mismatch caused by noise in the images[2]. So instead of simply taking the *true* point to be half-way between the nearest intersection of the two lines, a minimization process is used that allows the subset location in the images to be varied until the rays intersect. The minimized solution finds the optimum triangulated 3D point. An important feature of most codes is that the triangulation accuracy is being monitored for

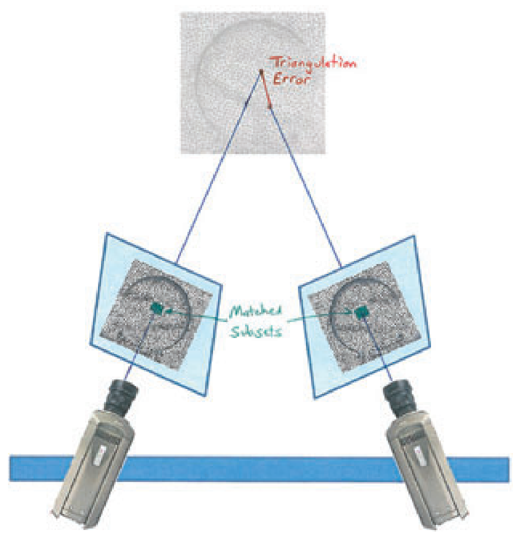

Matching subsets to be triangulated.

[1]Hartley, R.I. and P. Sturm, *Triangulation*. Computer Vision and Image Understanding, 1997. **68**(2): p. 146–157.

[2]There are a number of ways of doing this including letting the calibration parameters vary. However, I have described the most common approach.

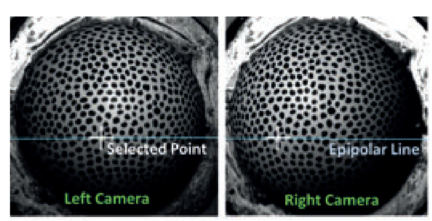

Epipolar line shown on both images.

The Art and Application of Digital Image Correlation is written by Phillip L. Reu. He received his PhD from the University of Wisconsin–Madison and is currently a Principal Member of Technical Staff at Sandia National Laboratory. He began working with digital image correlation in 2004 and is focused on understanding the influence of the unavoidable compromises made in field measurements on the final DIC uncertainty. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract No. DE-AC04-94AL85000. Email: phillip.reu.DIC@gmail.com

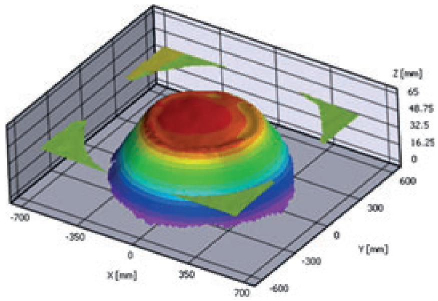

Dense 3D Field of Data. Displacement is found by subtracting the shape of two load steps.

$$\varepsilon = \frac{l_o - l}{l_o} = \frac{\Delta l}{l_o} = \frac{\text{(pixels)}}{\text{(pixels)}},$$

Where l_o is the original length
l is the current length, and
ε is the strain

[3] The 1D case illustrated is a simplification to visualize the process. However, this is usually done as a tensor that includes all the strain components at once using both data smoothing and curve fitting. Regardless, the conclusions drawn here are analogous for 2D-strain calculations.

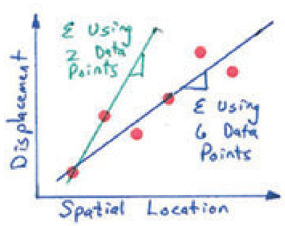

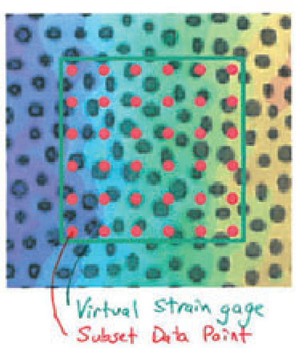

$$\text{Virtual Gage Size (mm)} = \left(\begin{array}{c}\text{\# of Data}\\ \text{Points used}\end{array}\right)\left(\text{stepsize (pixels)}\right)\left(\text{Pixel size }\left(\frac{\text{mm}}{\text{pixel}}\right)\right)$$

every subset using the epipolar constraint. The epipolar constraint results from the camera calibration, where for any pixel in either image, a line can be projected onto the other image. Now if the subset match and the calibration are good, the correlated pixel location should be extremely close to the epipolar line, much less than a pixel in fact! If this isn't the case something has invalidated the calibration, usually camera motion, or you have mismatched the subsets.

Post-Processing

We now have a dense set of 3D (or 2D) data points which defines the shape of the object. What now? Because this is experimental mechanics, strain is usually the next calculation completed. The commercial codes calculate the strain by taking the slope of the calculated displacements. Remember that the raw DIC results are always 3D points, with displacements being found by subtracting two 3D point clouds. Using ''the slope'' (or derivative) of the displacement[3] to calculate the strain follows naturally from the definition of strain: a change in length divided by the original length. Because we are measuring the slope, noise in the data is amplified making the strain results noisier than the displacement results. This is illustrated by a case using only two points. Any noise in the displacement data will be greatly amplified as illustrated in the figure. To minimize the noise, more points are used to help smooth the results. The number of data points selected for the slope calculation creates a ''virtual strain gage.'' A decision now needs to be made by the experimentalist: using more points to create a larger virtual gage provides better smoothing to reduce the strain noise, but smoothing removes strain gradients: often the most interesting result. A compromise between smoothing and strain averaging is ***always*** necessary. The physical size of the virtual gage is calculated by taking the number of points used in the strain calculation and multiplying that by the distance between points. It is surprising how quickly the gages become relatively large in the attempt to minimize the strain noise. Remember: the strain reported by the DIC software is the *average strain* in the region of the virtual gage! An important concept related to the virtual gage size is that the strain resolution does not greatly improve with increased camera resolution or smaller fields-of-view. This is a common misperception because the displacement resolution does improve. Why is the strain resolution fixed? Because both the numerator (Δl) and the denominator (l) change at the same rate and any increase in displacement resolution is cancelled out. This fact defines a fundamental limit of the strain resolution based solely on the noise of the images and errors in matching. Increased camera resolution will, however, for a given strain noise level, improve the spatial resolution of the strain, that is, decrease the size of the virtual gage.

Other post-processing parameters of interest often include velocity, curvature, strain rate, and displacement magnitude. Note that both velocity and strain rate are derivatives with respect to time and can also suffer from noise issues. The raw displacement data at this point can be manipulated to give other information: for example with axisymmetric measurements points at the same radius can be averaged together to reduce the noise and improve the measurement. Strain around a crack tip can be integrated to find stress intensity factors for fracture mechanics. You are limited only by your imagination. Remember that the primary measurement is always displacement and displacements should be used whenever possible.

Next Time: How to design the stereo-rig for a DIC experiment.

Stereo-Rig Design: Creating the Stereo-Rig Layout—Part 1

by Phillip Reu

Introduction

The first step in a successful digital image correlation (DIC) experiment is to design the stereo-rig; a two-camera unit that is used to take simultaneous images. The rig design involves taking into consideration the needed displacement accuracy, field-of-view, and preferred spatial resolution. Generally speaking you can count on a displacement accuracy of between 1/100th and 1/20th of a pixel with the out-of-plane resolution being the worst. The actual resolution is dictated by: camera noise, camera resolution, speckle quality, lens quality, camera motion, lighting conditions, *et cetera*. The list is long, and contributes to the difficulty of accurate uncertainty quantification (UQ) for DIC. To oversimplify things somewhat; to make accurate measurements one must obtain *high quality* images from a *rigid* stereo-rig.

Stereo-rig Design Decisions

1. **Stereo-rig Layout**
2. Camera Selection
3. Lens Selection
4. Stereo-angle Selection
5. Lighting

How to Estimate the Measurement Parameters Pre-Test

The first step in designing the stereo-rig is determining your desired field-of-view (FOV). FOV together with your camera resolution will define both the displacement accuracy and the spatial resolution of your results. Spatial resolution is defined as the distance between *independent* measurement points, either in pixels at the camera sensor or better in meters at the measurement surface. The spatial resolution is usually increased by overlapping the subset solutions. However, at some amount of subset overlap (around ½ the subset size), the measurements are no longer truly independent, and you are not adding any new measurement information to the data set. Maximizing the spatial resolution is often important for the strain calculation because the spatial separation of the data points will define the size of the virtual strain gage (see Article 4). With little guidance from the literature a good rule of thumb is to not make the step-size any smaller than ¼ the subset size[1]. Smaller step-sizes

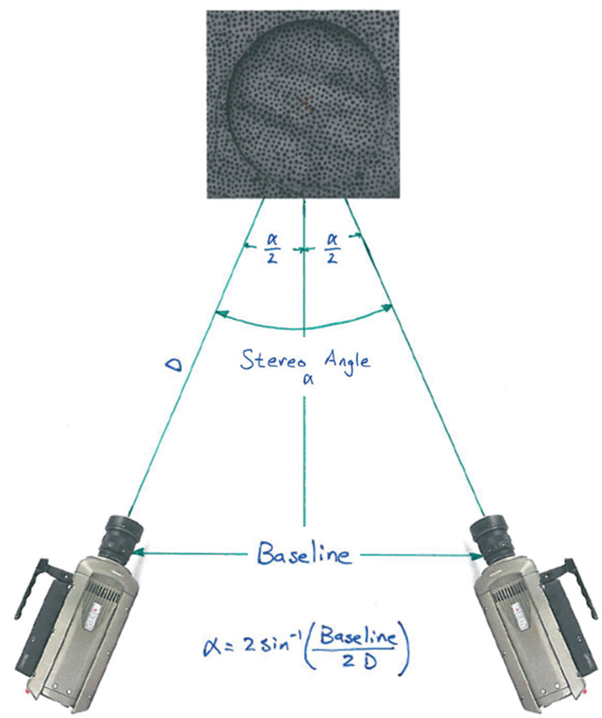

Calculating the stereo angle.

[1] Historically, a step-size of ½ was used to prevent the search algorithms from getting lost. Recent work* has recommended a step size of ≥ ½-subset size to minimize noise in the strain calculation.

*Ke, X.D. Experimental Mechanics, 2011, **51**(4): p. 423-441.

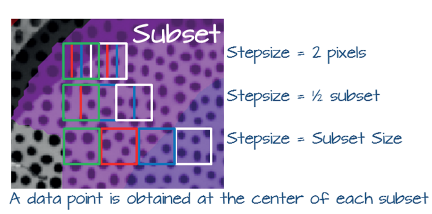

Subset and Step-size.

The Art and Application of Digital Image Correlation is written by Phillip L. Reu. He received his PhD from the University of Wisconsin–Madison and is currently a Principal Member of Technical Staff at Sandia National Laboratory. He began working with digital image correlation in 2004 and is focused on understanding the influence of the unavoidable compromises made in field measurements on the final DIC uncertainty. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract No. DE-AC04-94AL85000. Email: Phillip.Reu.DIC@gmail.com

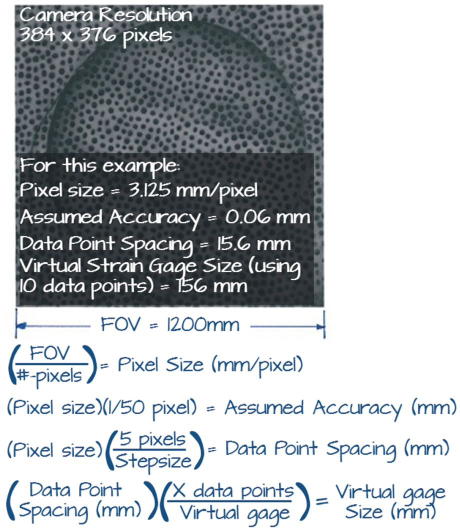

Experiment	Estimate (px)
Lab—good speckles	1/100 to 1/50
Lab—poor speckles	1/20 to 1/50
High-speed work	1/20 to 1/50
SEM	1/20

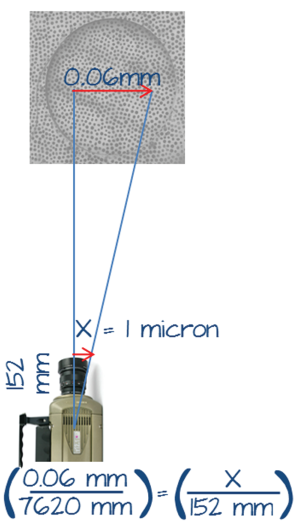

[2]Other codes call this: ''Refine projection parameter'' or ''self-calibration''

increase the solution time, because of the added subset correlations, without adding independent measurements.

To estimate the, spatial resolution, or the data-point spacing in millimeters, the FOV (mm) is divided by the number of available camera pixels, and then multiplied by the step-size (pixels). You can then use this information to calculate the virtual strain gage size by multiplying this by the number of points used to calculate the strain term. To estimate the measurement accuracy, the FOV is again divided by the number of pixels and then multiplied by the estimated correlation accuracy. The ''expected'' correlation accuracy presented in the table is a rule-of-thumb estimate, and the user should take care to adjust the predicted accuracy based on the anticipated experimental conditions. However, this estimate is useful in the experiment design phase to help determine camera resolution, lens choices, FOV, and frame rates of your stereo-rig.

Stereo-Rig Camera Mounting

Camera mounting is a critical element for making quality DIC measurements. To understand this, let's think about camera motion: If we are aiming for measurement accuracies in the range of $1/50^{th}$ to $1/100^{th}$ of a pixel, we need to ensure that the motion of the camera is less than this amount. Furthermore, any *relative* camera motion will compromise the extrinsic camera calibration used to triangulate the results in 3D space. This may seem like an easy target to meet, but it is not! This can be illustrated by calculating the assumed accuracy at the sample and then translating that to a camera motion at the end of the lens. For most situations, this motion is sub-micron for lab measurements, and micrometer scale for larger FOVs. What this practically means is that the cameras must be as rigidly mounted as possible. This is why commercial systems are sold with a single rigid bar on which the cameras are mounted. Tripod heads are not used; instead a simpler and more rigid mount that supplies only a pan rotation is used. The pan motion is required to set the overlapping camera views and then the stereo-rig can be aimed as a unit by mounting the stereo-rig on a tripod. Motion of the stereo-rig will add a rigid-body-motion to the results but not invalidate the camera calibration. Furthermore, rigid body motion is not usually a problem when deformation and strain are the desired result as it can easily be removed. If absolute position in space is desired, stereo-rig motion *will* corrupt these results. Practically speaking, for most experiments where displacement is desired, the errors caused by a well-designed and mounted stereo-rig are small compared to the motion of the measurement object. Because relative camera motion after calibration is often an issue, most software packages have included a camera re-orientation calibration option[2]. This allows the experimentalist to re-calculate the extrinsic camera parameters from a speckle pattern seen by both cameras.

In outdoor experiments the constraints on relative camera motion are difficult to meet and the calibration will always be compromised to some extent in these tests. A future article will discuss what can be done in these situations.

Next Time: Camera selection.

Experimental Techniques **36** (2012) 3–4

Stereo-rig Design: Camera Selection—Part 2

by Phillip Reu

Introduction

The selection of the field-of-view (FOV), the lens, and the digital camera are all intertwined. For example changing the digital camera detector size, will automatically change the FOV of a lens. Therefore, when designing a custom system the FOV, lens, and camera are chosen as a unit. More often, you already have the camera (when using a commercial system for example) and the easiest thing to change is the lens and standoff to control your FOV. This article discusses the choices involved in finding a good DIC camera and the effect of the camera choice on the ''hidden components'' of DIC. Of prime interest is the detector/pixel size, lens mounting, camera noise and camera rate.

Stereo-rig Design Decisions

1. Stereo-rig Layout
2. **Camera Selection**
3. Lens Selection
4. Stereo-angle Selection
5. Lighting

Interface	Speed	Synch.	Cable Length
Firewire	Good	Built In	Short
CoaXPress	6 Gbit	Built In	V. Long
USB	Good	Separate	Short
GigE	Good	Separate	V. Long
CameraLink	4 Gbit	Built In	Moderate

[1]This is obviously an evolving area. The camera technology has been improving at an incredible rate.

Choosing a Digital Camera

The first decision (other than cost) is the desired computer interface[1]. Each of the major options has different strengths. These are outlined in the side table. If speed is your most important criterion (up to 500 fps), CoaXPress is the best interface option. If your frame rate needs are more modest (5-15 fps) any of the methods will work. At this time the camera vendors usually offer the same detector size and resolution in a number of the communication formats. For all cameras there is always a tradeoff between resolution and frame rate. Higher resolution cameras will always operate more slowly than a lower resolution camera. Many cameras allow you to ''window'' the area you want to look at, which will improve frame rates by reducing the pixel count transferred to the computer.

Photos by E. Bystrom

C-mount Lens (left) and F-mount (Nikon) (right).

A primary assumption of Stereo-DIC is that the stereo-images are acquired simultaneously. Therefore the ability to synchronize the cameras is critically important. A nice feature of both CoaXPress and FirewireB (1394b) is the built in synchronization over the cable making it ideal for DIC. The other interfaces will require a hardware synchronization solution. An important difference between ''machine-vision'' cameras and web or point-and-shoot cameras is the ability to synchronize cameras via a trigger line. Without synchronization, a camera should only be used for 2D-DIC.

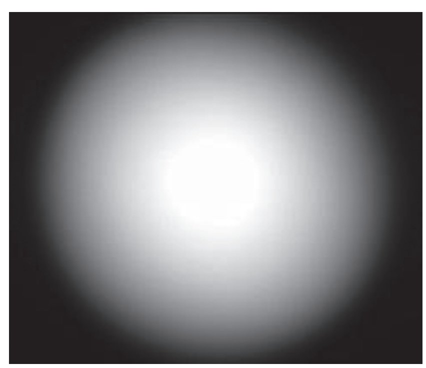

Vignetting.

The Art and Application of Digital Image Correlation is written by Phillip L. Reu. He received his PhD from the University of Wisconsin–Madison and is currently a Principal Member of Technical Staff at Sandia National Laboratory. He began working with digital image correlation in 2004 and is focused on understanding the influence of the unavoidable compromises made in field measurements on the final DIC uncertainty. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract No. DE-AC04-94AL85000. Email: Phillip.Reu.DIC@gmail.com

Your choice of camera will also affect your lens selection. Most machine vision cameras come with smaller detectors and threaded c-mount lenses that will be both cheaper and usually have lower distortions. The threaded c-mount lens also benefits from a more rigid camera mounting and is less likely to shift during the experiment and cause calibration and accuracy problems. As the camera detectors become larger, usually around the ⅔-inch to 1-inch size[2], c-mount lenses will cause image vignetting and a switch to the larger F-mount (Nikon) lenses is usually required. For high-speed work and very high resolution cameras, the mounting is always a Nikon mount[3] to account for the much larger sensor sizes in this type of camera.

A monochrome camera is always the best option. Adding color to the image does not improve the correlation, in fact due to the method of filtering to create a color image (Bayer for example), there is a net loss in resolution for DIC purposes. If the only camera available is a color camera, it may be used, but care should be exercised. Similarly, commercial point-and-shoot cameras can be used, but precautions must be taken to disable many of the built-in functions of the camera including: auto-focus, image stabilization, and any image compression. Monochrome machine vision cameras are always your best choice!

Noise and Pixel Size

The detector and pixel size are related to the total resolution of the system *and* the light sensitivity. Larger pixels will be more light-sensitive and have lower noise. However, with large pixels, as the resolution increases the sensor will become too large for most standard lenses. Typically, the camera sensitivity requiring large pixels is more important when selecting high-speed cameras, where obtaining enough light is always an issue. However for DIC camera selection, detector noise is essential because the measurement quality is directly related to the image noise and contrast (see Article 3 and the side figure here). Practically speaking for most modern machine vision cameras, both image noise and contrast are predominantly controlled by how well your sample is lighted rather than the camera itself. Better lighting will allow a lower camera gain, with a corresponding lower image noise and improved contrast. To quantify the image noise a number of static images may be averaged together pixel-by-pixel to find the noise distribution and level. After taking all of these factors into account, for most DIC applications a 2 to 5 Megapixel camera is a good compromise in resolution, sensitivity, and acceptable frame rates at an affordable price[4].

Next Time: Lens selection.

[2] Detector size is from the video tube days and is not reflective of the actual detector size.

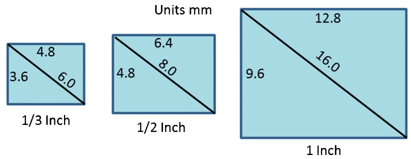

[3] There are other larger format lens mounts but they are very rare in digital imaging. Also, I have reservations about whether using a 14 megapixel camera at laboratory scale fields-of-view (inches or centimeters square) yields the resolution scaling one would expect.

Comparison of DIC results of two different images, one with low contrast and high noise (high gain) and the other with better contrast and lower noise. Note the difference in displacement noise in the static pre-test images shown on the bottom (identical scales).

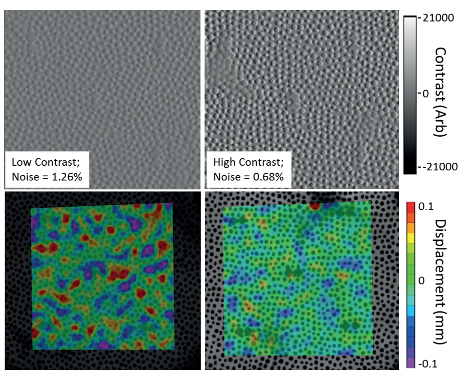

[4] This statement is a moving target.

Exercise for the Reader

The effect of noise on DIC can be investigated using a 2D DIC system. Acquire two static images and then create an exact copy of the reference image file and rename it. Run the correlation on the three images; the reference, the copied reference, and an independent static image. The DIC will show a perfect match for the copy while showing errors for the ''identical'' but independently obtained static images. Why? The camera image noise is different between the two independent images and creates both a bias error and a variance error in the match. The copied image has the exact same ''noise'' as the reference image and the ''correlated'' noise appears like a speckle and results in a perfect match.

Experimental Techniques **36** (2012) 3–4 © 2012, Society for Experimental Mechanics

THE ART AND APPLICATION OF DIC

Stereo-rig Design: Lens Selection—Part 3

by Phillip Reu

Introduction

Once the cameras have been specified and the field-of-view (FOV) is known, decisions regarding the lens can be made. For this discussion I will assume that you can control how far away the cameras will be from the experiment. This is not always the case! If a minimum camera standoff is required; your lens is now selected for you. Use that one, whether it is ideal or not. The rest of this article will discuss the considerations when the experimentalist has a choice in focal length. I will also point out that a good lens, if one can be afforded is better than a cheap one. Having said this, DIC has proven itself fairly robust at correcting for and working with lower quality lenses.

Stereo-rig Design Decisions
1. Stereo-rig Layout
2. Camera Selection
3. **Lens Selection**
4. Stereo-angle Selection
5. Lighting

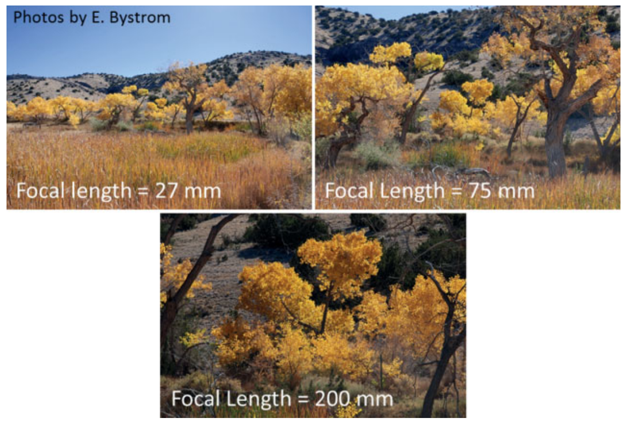

Lens Attributes

Lenses are defined by three primary attributes: the focal length, the aperture, and the resolution. Lenses come in either fixed-focal length, called prime lenses, or as zoom lenses, which will cover a range of focal lengths. The focal length is a measure of the magnifying power of the lens (see side figure). That is, the larger your focal length, the greater the magnification and the smaller your FOV, and vice versa. To estimate your FOV, the simple equation on the side is a good starting point; although with Macro lenses (used for close up work) it will be slightly off. The lens magnification can be increased simply and inexpensively by adding extension tubes between the lens and the camera. Care should be taken to not add *too* many extension tubes, as at some point, the resolution of the lens is exceeded and no effective magnification is added. This situation is the aptly named ''empty magnification'' and will have a negative impact on the DIC results. A lens is also defined by its largest aperture (also called the ''lens speed'') which specifies how much light can be gathered by the lens. Physically the aperture is an iris within the lens that controls the amount of light that will reach the detector. Frequently for DIC the aperture is reduced to increase the depth-of-field. Therefore it is not generally worth purchasing a large aperture or fast lens.

$$FOV \cong \frac{(Sensor\ size)\cdot(Distance\ to\ object)}{Focal\ Length}$$

$$Lens\ Mag. = \frac{Focal\ Length}{Focal\ Length - Distance\ to\ object}$$

$$Tube\ Mag. = \frac{Extension\ Tube}{Focal\ Length}$$

C-mount (left) and Nikon extension (right)

The Art and Application of Digital Image Correlation is written by Phillip L. Reu. He received his PhD from the University of Wisconsin–Madison and is currently a Principal Member of Technical Staff at Sandia National Laboratory. He began working with digital image correlation in 2004 and is focused on understanding the influence of the unavoidable compromises made in field measurements on the final DIC uncertainty. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract No. DE-AC04-94AL85000. Email: Phillip.Reu.DIC@gmail.com

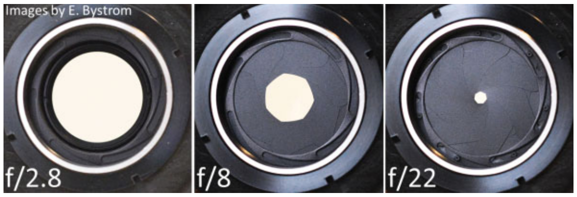

Illustration of the aperture

[1] $Resolution(lp/mm) = \frac{1}{2 \cdot Pixel\ Size\ (mm)}$

When purchasing a lens, the resolution is most often specified in line-pairs per millimeter (lp/mm). This measure indicates the resolving power of the lens based on how well closely spaced lines can be resolved when imaged through the lens. The required resolving power can be related to the pixel size of your camera using the sampling theorem[1]. It is prudent to select lenses that will out-resolve your detector. Remember that adding extension tubes will change the magnification and can quickly push the lens too far and create a situation with empty magnification.

Stereo-DIC Lens Selection

Commercial DIC systems usually come with prime lenses. Why is this? The short answer is that a prime lens is simpler in design and therefore has lower optical distortions. With a typical DIC system a variety of matched lenses are supplied with focal lengths ranging from 17-mm to 75-mm. This in coordination with some extension tubes will cover most of your laboratory situations. In outdoor setups, there is often less flexibility in positioning the stereo-rig, and zoom lenses are used to optimize the FOV. This is essential because the displacement and spatial resolution are direct functions of the number of pixels across the object as discussed in Article 5. As a result, you want to have as many pixels in the measurement area as possible.

[2] $\varepsilon_{error} = \frac{out\ of\ plane\ motion}{Effective\ Distance}$

[3] Sutton, M.A., et al., Optics and Lasers in Engineering, 2008. **46**(10): p. 746-757.

[4] For a quick telecentric tutorial see: http://www.opto-engineering.com/telecentric-lenses-tutorial.html and http://www.edmundoptics.com/technical-support/imaging/using-telecentric-lenses/

2D Lens Selection

For 2D DIC, particularly if calibration is not going to be done, a high quality *telecentric* lens is a great choice because of its ability to minimize the largest and most common error source in 2D-DIC: out-of-plane motion. Out-of-plane motion will create a strain error that is proportional to the motion divided by the effective distance[2,3]. For a standard lens this distance is simply the lens stand-off; but for a telecentric lens the effective distance is increased by a factor of 10× to 100×[4]. As telecentric lenses are expensive, a good quality middle focal length standard lens is the next best choice.

Example of a distortion field for a rigid borescope.

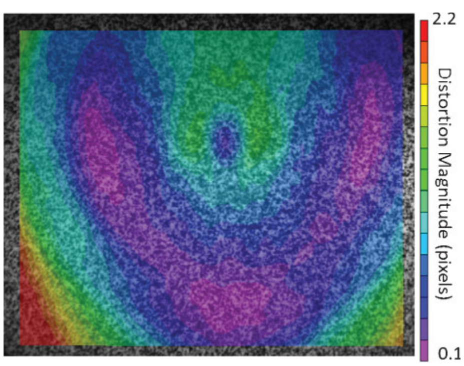

Special Circumstances

Everything in the optical path from the object to the detector is ‘‘the lens’’; including viewports, mirrors, and air currents. These non-traditional optics may create large and non-radial distortions in the images. Standard DIC implementations will *only* correct modest radial distortions and *will not* correct large or non-radial distortions. An example of this problem is shown in the side figure, where the distortion field of a rigid borescope is illustrated. These situations will require an advanced calibration and distortion correction scheme. Fortunately, with care, most optical systems can be made to work, including stereo-microscopes[5] and electron microscopes[6].

[5] Schreier, H.W., et al., Experimental Mechanics, 2004. **44**(3): p. 278-288.

[6] Sutton, M.A., et al., Experimental Mechanics, 2007. **47**(6): p. 775-787.

Experimental Techniques **37** (2013) 1–3 © 2013, Society for Experimental Mechanics

If in doubt, check your distortion level! To do this:

1. Translate a flat and rigid plate in-plane in the *x*- and *y*- directions over about 10% of the field-of-view taking 5 images in each direction.
2. Analyze the displacements with the DIC code.
3. Subtract the rigid-body motion using the average displacement of each in-plane step.
4. Any remaining displacements are a measure of your lens distortions (in real units for a calibrated system, or in pixels for an un-calibrated system).
5. The distortions can be emphasized by calculating the strain, which is the derivative of the displacement error.

The most important point for any experimentalist is to understand your entire measurement system.

Next Time: Stereo Angle Selection

Stereo-rig Design: Stereo-Angle Selection —Part 4

by Phillip Reu

Introduction

We have now selected a field-of-view (FOV) and corresponding cameras and lenses. The next question: What is the best stereo angle? First remember that the Stereo-DIC calibration yields both the intrinsic parameters, which tell you about your lens and camera system, and the extrinsic parameters which tell you about how the cameras relate to one another in space. The two extrinsic parameters most easily understood are the stereo-angle and the baseline. The stereo-angle is the angle between the two camera axes and the baseline is the distance between the cameras[1]. Some simple trigonometric relationships can be used to estimate a target angle and baseline length during the experiment design. The ''exact'' angle is not important as it will be calculated during the calibration.

Stereo-rig Design Decisions

1. Stereo-rig Layout
2. Camera Selection
3. Lens Selection
4. **Stereo-angle Selection**
5. Lighting

[1] More precisely the optical centers of the cameras.

Determining the Stereo-angle

So what is the optimum angle? As with most DIC decisions, that depends. In general you will want an angle between 15 and 35 degrees. But which extreme you choose depends on what you want to measure and the shape of your measured surface. In this decision you will be trading the in-plane resolution for the out-of-plane resolution. The narrower stereo-angles (shorter baseline) will improve the in-plane results at the cost of a higher out-of-plane uncertainty. For most experimental mechanics experiments, where strain is the desired measurement, a narrower stereo-angle is then generally preferred. Unfortunately, there are two added complications: the uncertainty also varies with the location within the FOV and the lens focal length.

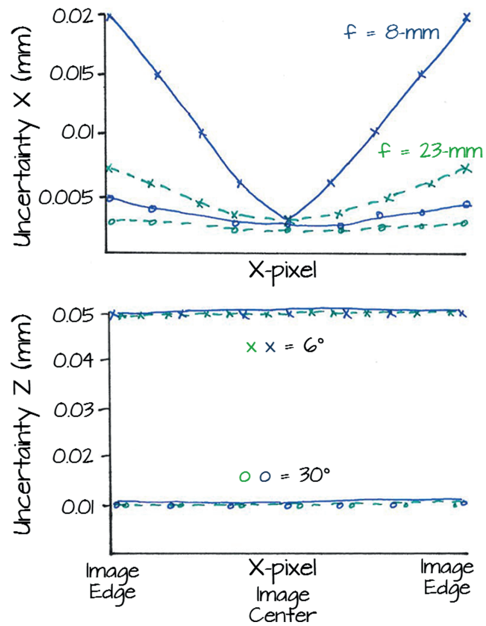

Uncertainty as a function of image location

The illustrations used in this paper are modified versions of figures from an important DIC uncertainty quantification paper by Ke[2]. First we will look at the most accurate location in the image. Not unexpectedly this is towards the center of the image and is illustrated in the side figure for an 8-mm and 23-mm lens. From this we can note that a longer focal length lens will yield a larger area with lower uncertainty for a corresponding stereo-angle.

Now we will look at the interaction of the stereo-angle and the lens focal length. The final figure shows two different lens focal lengths

[2] Ke, X.D Experimental Mechanics, 2011. 51(4): p. 423–441.

The Art and Application of Digital Image Correlation is written by Phillip L. Reu. He received his PhD from the University of Wisconsin–Madison and is currently a Principal Member of Technical Staff at Sandia National Laboratory. He began working with digital image correlation in 2004 and is focused on understanding the influence of the unavoidable compromises made in field measurements on the final DIC uncertainty. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract No. DE-AC04-94AL85000. Email: Phillip.Reu.DIC@gmail.com

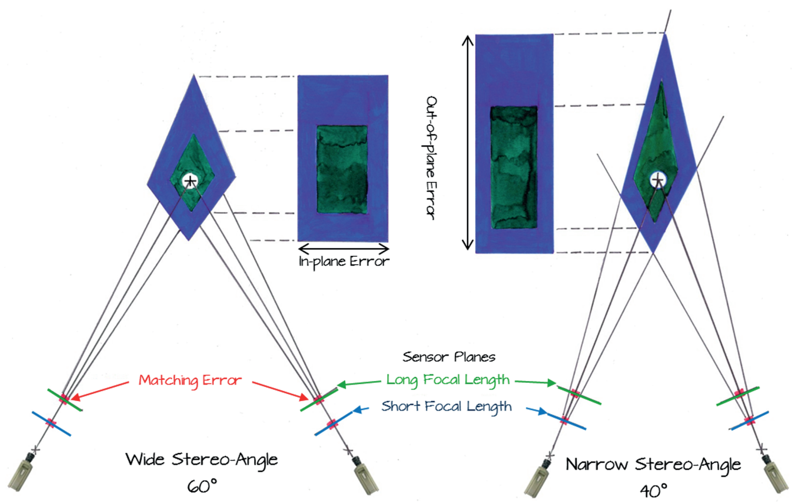

Uncertainty as a function of the stereo angle and focal length and matching accuracy.

and two different stereo angles (enlarged for illustrative purposes). The line labeled ''sensor plane'' in the figure represents the image sensor with the central line showing the ''true'' subset match and the red box indicating the uncertainty of the stereo-correlation. If you project this matching error through the sensor plane to the location of the object you can visualize the uncertainty as the shaded diamond shaped area. The green shaded region is for the 75-mm lens and the blue is for the 8-mm lens. The focal length has the effect of moving the sensor plane towards the object, which in turn reduces the influence of the matching error when projected to the object. Finally, the shaded rectangles are proportional to the in-plane and out-of-plan errors. So to reduce the uncertainty, we would prefer a longer focal length lens (within reason, lenses longer than 200-mm can cause issues with the calibration).

Conclusions

The two figures give a good idea of the compromises that you are making with the stereo-rig. Study them.

To attempt to sum this up in words:

1. Use longer focal length lenses.
2. Use a larger stereo-angle (longer baseline) to improve out-of-plane results.
3. If using a wide angle lens, use a stereo-angle of at least 25 degrees.
4. Keep your measurements towards the center of the FOV.

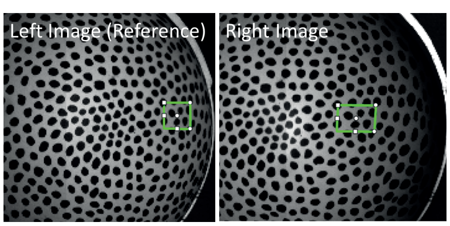

Notice the grossly deformed subset in the matching image for a spherical part.

This has one more complication: The shape of your object. If you are looking only at flat plates there are less issues, but when looking at curved surfaces, large stereo angles will cause difficulties for both the matching and the shape function (see side figure). This may require that a compromise be made on the stereo-angle to not have grossly distorted subsets (facets) at the edge of a curved object. Understand the compromises!

Special Circumstances

For large outdoor setups, there is likely a need to compromise between the stereo-angle and the baseline. Remember, the calibration requires that the cameras remain rigidly fixed relative to each other. This is most likely to be true if both cameras are mounted on the same bar. Therefore, it may be better to have a smaller angle in order to keep the cameras together on a single beam.

Next Time: Lighting

THE ART AND APPLICATION OF DIC

Stereo-rig Design: Lighting—Part 5

by Phillip Reu

Introduction

Lighting is a key component in creating great DIC measurements. Why? Measurement quality is *directly* related to the image quality. The image contrast and noise is determined by the interaction of the speckle pattern and the lighting. Often times the lighting can help salvage a situation where the speckle pattern is less than ideal; alternately, bad lighting can easily ruin a good speckle pattern. Take the time to get the lighting right!

Stereo-rig Design Decisions

1. Stereo-rig Layout
2. Camera Selection
3. Lens Selection
4. Stereo-angle Selection
5. **Lighting**

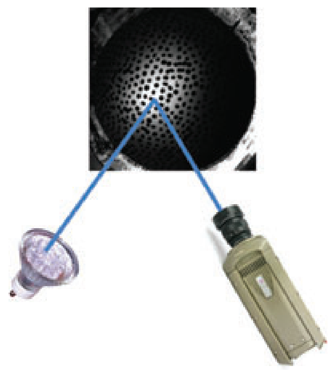

A Case of Bad Highlighting.

What is good DIC lighting?

The qualities in lighting you will be striving for are: flat, uniform, and of adequate intensity. Flat lighting means that the light seems to have no source: It is diffuse. This creates light that does not generate highlights or bright areas in the image. The elimination of highlights is more easily accomplished when looking at objects that are flat and remain flat, because the only requirement is to ensure that the light source has no direct path from the object to the camera. Good lighting is greatly aided by using the right paint for the surface. *Only flat or matte paints should be used*. A gloss or semi-gloss paint will be difficult to light without creating highlights. Even with flat paint, when looking at curved objects, or objects which deform toward or away from the camera, there will be problematic areas that will align the light with the camera (see side figure). It is in these situations where using a method to diffuse the light will be most beneficial. A common approach is to use a diffusion material of translucent plastic available at many photography stores. The use of this material is illustrated on the right for an experiment designed to simultaneously measure all sides of a sphere. Another method is to reflect the light source off of a matte surface such as paper or Styrofoam or even the walls of the room. All of these methods require a significant amount of light because you will be ''throwing away'' a lot of the light to create the flat illumination.

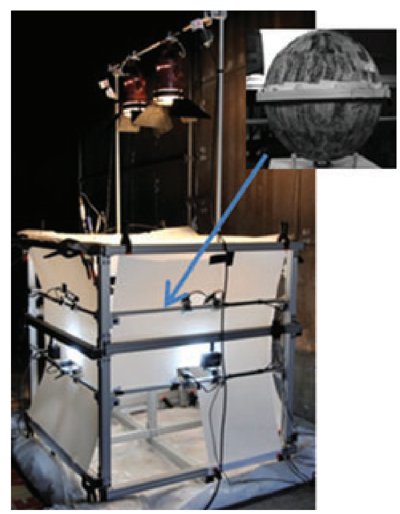

Extraodinary efforts to create flat lighting.

The Art and Application of Digital Image Correlation is written by Phillip L. Reu (Phillip.Reu.DIC@gmail.com). He received his PhD from the University of Wisconsin–Madison and is currently a Principal Member of Technical Staff at Sandia National Laboratory. He began working with digital image correlation in 2004 and is focused on understanding the influence of the unavoidable compromises made in field measurements on the final DIC uncertainty. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract No. DE-AC04-94AL85000.

Adequate illumination

DIC often requires a lot of light for two reasons: First, creating flat illumination (as already discussed) requires extra ''unused'' light, and second, the lens aperture is often stopped-down to increase the depth-of-field[1]. This can lead to two problems during the experiment; specimen heating and heat waves. Heating of the specimen can be quickly determined by measuring the temperature, often simply by holding your hand in front of the lights. Depending on the specimen the unwanted heating may ruin the desired measurement; particularly if the lights will be on for extended periods of time. The production of heat waves is a more insidious problem. Heat waves are tricky because they often cannot be seen directly in the images, and will appear in the results as displacements that often seem reasonable. The side figure shows false displacements caused by heating of the air (left) and the actual deformation at the same load step (right). Heat waves may not only be caused by the lighting, but also emanate from a heating part. Fortunately for many experiments a simple fan can be used to remove the heat waves and restore the measurement.

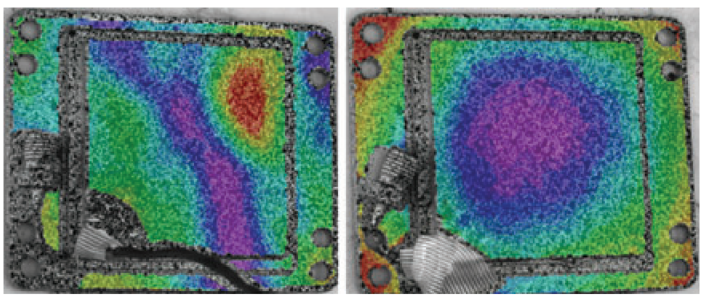

(Left) Illustration of heat waves masking the out-of-plane displacement. (Right) Correct out-of-plane measurement.

The final test for the lighting is to check the number of counts or grey levels! Remember that your matching quality is based on the ratio of the image noise to the image contrast. Contrast is proportional to the difference in grey levels between the bright speckles and the dark speckles. This difference can be found with the live cameras by hovering the cursor over different portions of the image and checking the counts in both a white and a black region. You would like as many counts as possible – within reason. As a best practice, the lighting should be adjusted until there are at least 50 counts between black and white for an 8-bit camera. 12-bit cameras will typically have more counts; however, with low noise, contrast greater than 50-counts does not improve the matching.

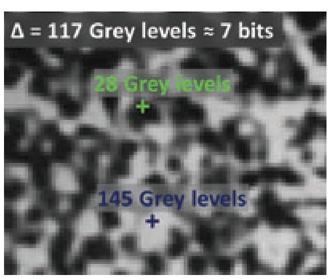

The optimum light source

The best lighting for DIC is an LED source. This is because it can supply a lot of light with less heating than a comparable halogen or incandescent light (such as a painter's light). White light is not required. Remember DIC uses monochrome cameras, so a green or red LED will work just fine. Typically green is chosen because they are readily available and the camera detector is more sensitive at this color. You may be tempted to use lasers for DIC lighting because they can supply abundant amounts of light very cheaply. Do not do this! The coherent nature of laser illumination will create laser speckling that will mask the painted speckle pattern. And furthermore, the beautiful speckles created by the laser cannot be used for DIC because of the rapid evolution of the speckle pattern during sample motion and strain, which causes the speckles to change and become ''uncorrelated''.

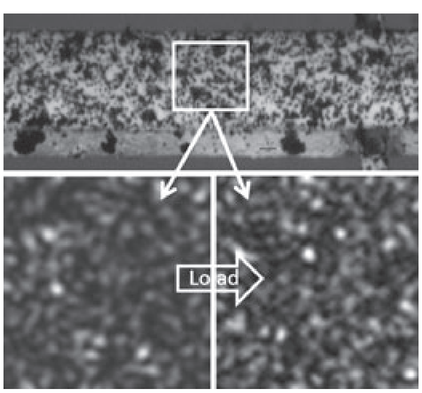

Laser speckle masking a painted speckle pattern and evolving under loading.

Next Time: Pre-calibration routines

[1] Reu, P., *Experimental Techniques*, 2013. **37**(1): p. 1-3.

Calibration: Pre-Calibration Routines

by Phillip Reu

Introduction

The camera, lens, stereo-angle, field-of-view (FOV), and lighting are now selected for the DIC experiment. The next step is to calibrate the DIC system. Calibration is a breakpoint in the DIC process. Once calibrated, *nothing* can be changed in the camera system. The stereo-rig as a unit may be very carefully moved, but anything that violates any calibration parameter cannot be done, including changing the focus, the aperture, the field-of-view, or the camera pointing. This article examines a number of important pre-calibration procedures, listed on the side, which should be followed to prevent having to needlessly recalibrate the stereo-rig. While not difficult, calibration can be time consuming, particularly when having to be done multiple times before starting the experiment.

Conducting a Calibration

1. **Pre-calibration routines**
2. 2D Calibration
3. A good calibration image
4. Stereo Calibration
5. Sanity Checks on the Calibration
6. Care and Feeding of a Stereo-rig

Pre-Calibration Checklist

1. Check for contamination on lens/CCD
2. Align the stereo-rig
3. Set the aperture, exposure and focus
4. Calibrate the stereo-rig

Check for Dirt

The most tempting step to skip is to check for dirt. Circumventing this easy step will be regretted! To look for contamination, find a uniform colored surface, usually white or grey. Defocus the lens and look for any blobs. This *cannot* be done while looking at a pattern, as the dirt spots will be completely hidden by the speckles. To determine if the dirt is on the lens or the detector, a nice trick is to rotate the lens with a live image. If the dirt moves, it is on the lens. If it remains stationary it is on the detector. Carefully clean the dirty components (according to the manufacturer's recommendations) and re-check as dirt can be difficult to remove. Sometimes camera detectors, particularly high-resolution cameras, have dead pixels that will look like dirt and may also be found with this technique. Dead pixels are dealt with in different ways by the camera vendors. Sometimes they are mapped out in the camera hardware via pixel mapping, or corrected for in the software. Regardless, care should be taken when dead pixels exist. The side figure shows a speckle pattern with dirt on the lens, and the resulting poor correlation in the region of the dirt. The poor matching is caused by the dirt being a stationary ''speckle'', which causes problems when trying to match the intensity in a subset between the reference and deformed image.

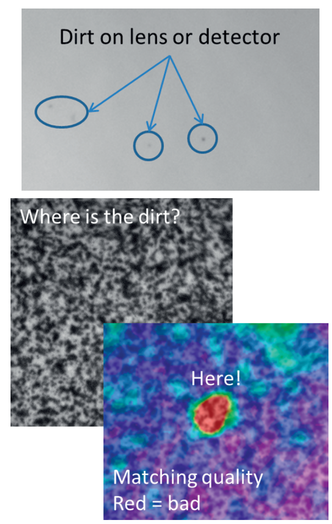

The Art and Application of Digital Image Correlation is written by Phillip L. Reu (Phillip.Reu.DIC@gmail.com). He received his PhD from the University of Wisconsin–Madison and is currently a Principal Member of Technical Staff at Sandia National Laboratory. He began working with digital image correlation in 2004 and is focused on understanding the influence of the unavoidable compromises made in field measurements on the final DIC uncertainty. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract No. DE-AC04-94AL85000.

[1] Reu, P., *Exp Tech*, 2013. 37(2): p. 1-2.

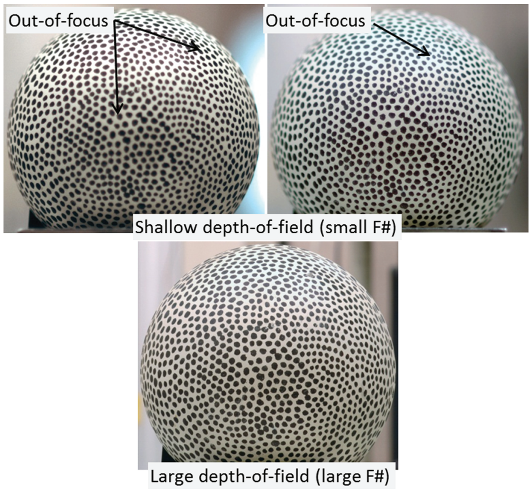

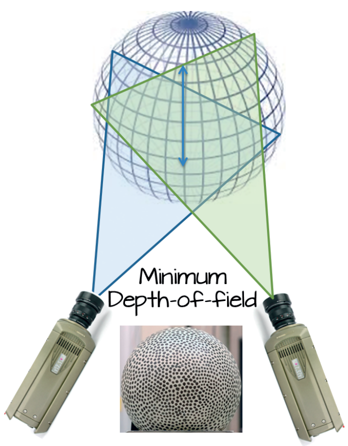

The object must be in focus in *both* cameras anywhere correlation is to be done.

Align the Stereo-rig

With your test object in place, carefully set the FOV and camera pointing angles to get the optimum image and stereo-angle[1]. The FOV must include all parts of the object to be measured in *both* cameras. Arrange this carefully; once the system is calibrated, the cameras cannot be moved without re-calibration.

Set the focus, aperture and exposure

The focus, aperture, and exposure are all set at the same time because they are all interrelated—the depth-of-field is controlled by the aperture which in turn will change the exposure or lighting. To start: Set the lighting, aperture and exposure to have at least 50 counts on the speckle pattern. Once the lighting is good, the focus can be optimized. The best approach is to open the aperture wide open (smaller F/#) in order to have the smallest depth-of-field while simultaneously changing the exposure so the image is not over-exposed. The location of best-focus can now more easily be found. This process is illustrated on the side figure with three images: Each with a different focus plane and depth of field. As the object in this test is going to move towards the camera, it is best to have the best-focus at the front. The bottom section of the figure is with a reduced aperture (larger F/#) that yields greater depth-of-focus resulting in the entire object being sharp. Typically a DIC experiment is run with a fairly small aperture to improve the depth-of-field. Even with flat objects, a reasonably large depth-of-field is required because of the perspective difference between the two camera views. The need for good depth-of-field, and exposures short enough to freeze the object motion, often require a large amount of light.

Check the calibration volume

The final check is to move the calibration target around in the desired calibration volume checking the focus. The calibration volume is the 3D region in which the object will move and be measured. Sometimes the calibration process may require more depth of field than the actual experiment, particularly for small fields-of-view with high magnification.

Now that the pre-calibration process is complete: The calibration may be done.

Next Time: What is a good calibration image?

THE ART AND APPLICATION OF DIC

Calibration: 2D Calibration

by Phillip Reu

Introduction

Calibration is not only for Stereo-DIC; depending on your experiment, it may be important for 2D-DIC as well. The reasons to calibrate in 2D include correcting for lens distortions, obtaining the correct image scaling for units in millimeters[1], and correcting for non-planar surfaces. For many 2D experiments skipping the calibration may be justified; however, make sure you understand the repercussions.

Conducting a Calibration.

1. Pre-calibration routines
2. **2D Calibration**
3. A good calibration image
4. Stereo Calibration
5. Sanity Checks on the Calibration
6. Care and Feeding of a Stereo-rig

[1]Or other engineering units.

Single Camera Calibration

To begin to understand this topic, it is useful to look at what a camera calibration is attempting to do. At its simplest, it is relating the physical 3D world to a 2D camera image via a mathematical model; usually this is the ''Pinhole Camera'' model. The schematic on the right shows the model parameters typically used and represents their physical meaning. These parameters are called ''Intrinsic'' parameters because they are related to the physical system of the camera and lens[2,3]. They include the image center (C_x, C_y), skew (Θ), and the focal length of the lens (f_x, f_y), which is often expressed in pixels, and should approximately match the focal length of the physical lens when multiplied by the pixel dimensions[4]. Another group of calibration parameters are those to correct for lens distortions. Most often with medium and long focal length lenses the imperfections are not visible, but can often be seen in images from ''wide-angle'' lenses (see side figure). It is important to note that just because a distortion cannot be seen, does not mean it does not exist. Remember, we would like to have displacements that are subpixel. By definition, the human eye looking at an image cannot ''see'' subpixel.

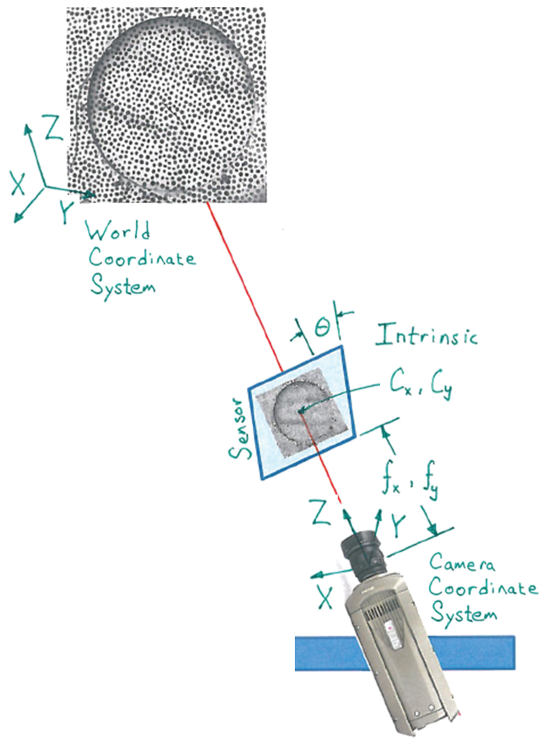

2D Camera Model (Typical)

[2]Sutton, M.A., J.J. Orteu, and H.W. Schreier, *Image Correlation for Shape, Motion and Deformation Measurements*. 2009, New York, NY: Springer.

[3]Reu, P., *Experimental Techniques*, 2012. **36**(2): p. 3–5.

[4]If used without extension tubes

Checking and Correcting for Lens Distortions

There are some real risks when 2D-calibration is going to be skipped. The primary one is that the measured displacement or strain results

The Art and Application of Digital Image Correlation is written by Phillip L. Reu (Phillip.Reu.DIC@gmail.com). He received his PhD from the University of Wisconsin–Madison and is currently a Principal Member of Technical Staff at Sandia National Laboratory. He began working with digital image correlation in 2004 and is focused on understanding the influence of the unavoidable compromises made in field measurements on the final DIC uncertainty. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract No. DE-AC04-94AL85000.

Distortions caused by a wide angle lens.

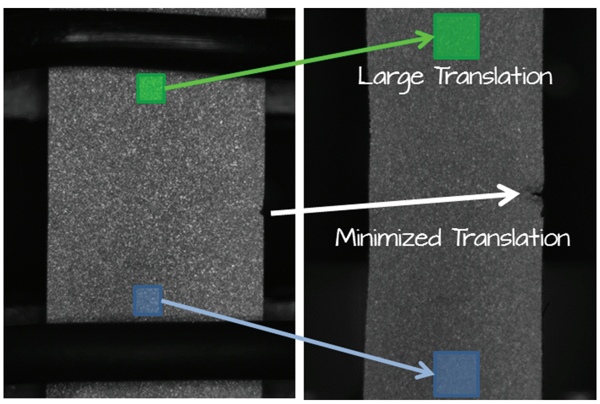

Note that the top and bottom of the sample travel a large distance. However, the region of interest was maintained near the center of the image.

[5]Note this is not an exact measure of the distortions, just a quick check. And out-of-plane motion could cause problems as well.

[6] $Strain\ Error = \frac{Out-of-planeMotion}{\text{Distance to the Object}}$

Important Reasons for 2D Calibration

1. Corrects lens distortions
2. Corrects for non-planar surfaces
3. Scales the image to millimeters rather than pixels

may be corrupted by the lens distortions. This is more important when the object is going to be loaded and the loading then results in a translation of the test item across the image, as this will exacerbate the problem. Unfortunately, the motion condition is more common than not as illustrated in a typical tensile testing experiment (see the side figure) where the sample translates across the field-of-view as it is stretched.

So it seems prudent to take a quick look at the lens distortions. Fortunately this is fairly simple. Take your sample and without loading, translate it across the FOV. For example, the upper grip of the tensile test machine can be released, and the lower grip lowered while imaging[5]. The images can then be analyzed as usual. Be sure to calculate both the displacement, which is rigid-body only at this point, and the strain for all of the images. Because there is no loading, the strain is a calculation of the displacement gradient, which is very sensitive to the lens distortions. Any pattern in the displacement or strain field indicates that there are lens distortions. If the displacement and strain fields appear as simple ''noise'' lens distortions are not likely to be a problem. This procedure, beyond checking the lens, will also determine the displacement and strain noise floor of your current setup and may suggest areas of improvement in speckling and lighting.

Possible Problems with 2D

There are a number of things the single camera model cannot tell you: Most importantly is the distance from the camera to the object. (That is without a number of usually impractical assumptions.) Because of this, the 2D assumption that the camera remains exactly the same distance from the test object is critically important. Any motion of the object or camera can easily corrupt the results by causing the object to appear to grow or shrink. This will introduce a strain error that can be estimated with the following equation[6]. If you cannot ensure in-plane motion, or the errors introduced by the motion are too large, use stereo-DIC, which removes the in-plane assumptions.

Next Time: Stereo Calibration.

Experimental Techniques **37** (2013) 1–2 © 2013, Society for Experimental Mechanics

Calibration: A good calibration image

by Phillip Reu

Introduction

For Stereo-DIC to work, the rig must be calibrated. Calibration is also advised for many 2D experiments as well. Article 2[1] discussed the parameters that are found during the calibration. This article will discuss what the calibration images should look like. Different commercial systems handle this differently; some systems require a prescribed set of calibration motions while others allow the user to choose the calibration images to be used. Irrespective of the process, all calibrations are best done by providing a series of quality images that ''fill'' the region where the measurement will be made, referred to as the ''calibration volume.'' A future article will discuss how many images are required, this article discusses: What does a good calibration image *look* like?

A Good Image

A short list of the attributes of a good calibration image is given on the side. The key is to understand that the software is going to attempt to extract a large number of points from an image of a calibration object and mathematically solve for the calibration parameters. Consequently, anything that corrupts the point extraction process will degrade the calibration quality, including: image noise, poor contrast, blur, and small features.

Calibration issues and solutions

Lighting is critical for a good calibration image. Hopefully, you have already achieved ''flat'' lighting to improve your DIC speckle images[2], and if so, calibration will be simpler. In practice this is difficult. The need to change the angle of the target will change the illumination and contrast in the different orientations. To compensate for this, monitor the live view and adjust the exposure to ensure that the images are neither under nor over exposed. Unlike the aperture (set on the lens with the f-stop), the exposure can be varied during the calibration process to compensate for any change in lighting without changing any calibration parameters. Furthermore, if needed, the exposure can be controlled independently on the two cameras as long as the cameras remain synchronized. Changing the exposure is

Conducting a Calibration

1. Pre-calibration routines
2. 2D Calibration
3. **A good calibration image**
4. Stereo Calibration
5. Sanity Checks on the Calibration
6. Care and Feeding of a Stereo-rig

[1]Reu, P., *Exp. Tech.*, 2012. **36**(2): p. 3-5.

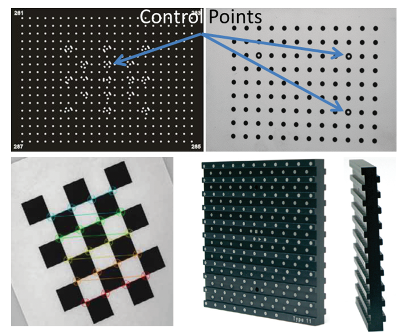

Typical Calibration Targets

Attributes of a good calibration image

1. Adequate and even illumination
2. In focus
3. No board motion (blur)
4. Cameras synchronized
5. Control points in FOV
6. Fill the calibration volume
7. Adequate dot (feature) size
8. Take care of your calibration targets

[2]Reu, P., *Exp. Tech.*, 2013. **37**(1): p. 1-3.

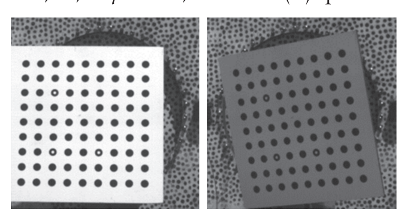

The left image, which is reflecting the sun, is over exposed. The right image is under-exposed. Adjust exposure to optimize contrast throughout calibration.

The Art and Application of Digital Image Correlation is written by Phillip L. Reu (Phillip.Reu.DIC@gmail.com). He received his PhD from the University of Wisconsin–Madison and is currently a Principal Member of Technical Staff at Sandia National Laboratory. He began working with digital image correlation in 2004 and is focused on understanding the influence of the unavoidable compromises made in field measurements on the final DIC uncertainty. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract No. DE-AC04-94AL85000.

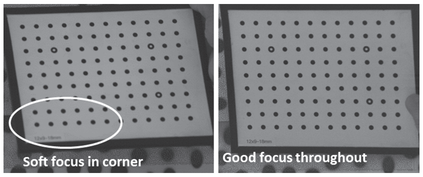

Focus problems

preferable to changing the gain, as this will add noise to the image, which will be bad for both the calibration and the DIC matching. Over-exposed images are created when there is a straight path from the light source to one of the cameras and will occur only over a small range of angles. This unfortunate alignment can often be avoided by slightly adjusting the calibration target angle. The small angle change required creates no calibration problems because there are no ''exact'' target positions to obtain, only that a variety of angles and positions be acquired.

It is important that the calibration grid remain in focus. Poor focus can lead to extraction errors and will decrease the quality of the calibration. Maintaining focus can be particularly difficult when using high magnification as the depth-of-field is much smaller. To avoid problems, monitor the image focus as you acquire the calibration images. Image blur due to target motion during exposure will also cause problems. Image blur is controlled by the exposure time, in this case, shorter is better and will improve calibrations. For a small FOV, a jig to hold the calibration target in place rather than hand-held is preferable to minimize grid motion. For any type of calibration it is best to steady your hands on a solid base while holding the target (See side figure). As a general rule, an exposure of 25 milliseconds or less will help prevent motion blur. A common calibration error can occur if there is a lack of synchronization of the cameras. This can be diagnosed if the individual cameras are solving correctly for the intrinsic parameters, but fail to find the extrinsic parameters[3]. Every calibration target will contain points (dots or corners) that are special locations on the target that give the calibration routine the scale of the view and are used by the software to move from pixels to engineering measurement units, such as millimeters. Some software will also use control points to define a global coordinate system. The control points must be in the FOV or the calibration image cannot be used. Usually the calibration routine will remove that image from the calibration, or the software will not accept this as a valid image. It is best that the calibration image fill the entire FOV. This is not required, and indeed is not possible with a large FOV. But for lab experiments if the calibration grid fills the FOV, fewer images will be required to calibrate the system. To ensure a successful calibration the dots to be extracted must be 10-30 pixels in diameter, for the square features they should be 15-20 pixels. Finally, the calibration target should be rotated, tilted and translated throughout the entire calibration volume. It is best to follow a set process of tilting the target, rotating the target, and then translating/plunging the target, taking images throughout the process. For the DIC packages that require prescribed motions, these will be defined by the software. Taking care of the calibration targets is important. Soiling, bending, scratching or anything that mars the surface may cause feature extraction problems. Inaccurately extracted features will in turn decrease the accuracy of the calibration. The calibration routine is robust as it is an ''averaging'' process using many extracted features to find the calibration parameters. However, a number of poorly extracted points will cause calibration errors that will result in a measurement bias.

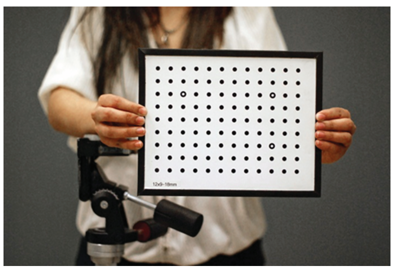

Hand holding a calibration grid on a tripod

[3]Other indicators of a bad calibration are poor calibration scores, and large uncertainties in the reported parameters.

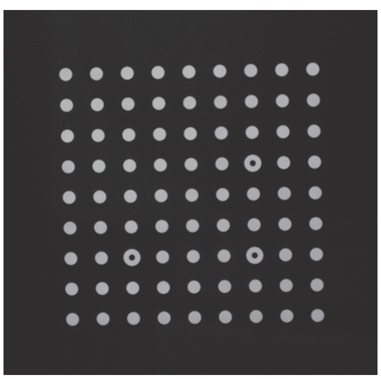

Back-lit calibration target

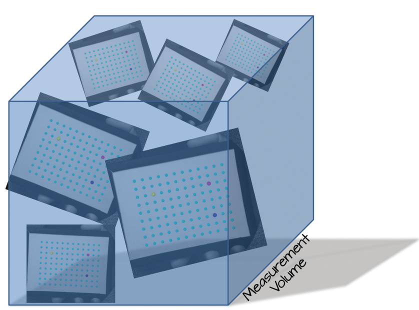

Conclusions

Finally, monitoring the live image for contrast, focus and position is extremely useful. Whenever possible position the DIC computer so it can be seen from where the calibration target is being held. If that is not possible, an assistant is important to track and ensure good image quality. Some DIC codes automatically capture an image when it determines that the live image is good. As with many things in DIC, the quality comes down to having a ''good'' image. For calibration this means adequate contrast, appropriate sized calibration features, low image noise, good lighting, and a stable synchronized stereo-rig. When the calibration scores are indicating problems, you have violated one of these precepts (guidelines).

Next Time: 2D Calibration.

Calibration: Stereo Calibration

by Phillip Reu

Introduction

Stereo-calibration is attempting to find all of the intrinsic and extrinsic camera parameters that are needed to triangulate in the measurement volume. An illustration of the camera model is shown on the right. These parameters are found by extracting calibration points from a series of images. Usually centers of dots or corners of squares. A previous article covered what a good calibration image should look like; this article will look at how best to accomplish obtaining a good calibration with the next article discussing checking the stereo calibration. A final note for brevity: When I refer to a calibration image in this article, it refers to a matched and synchronized image-pair from the stereo cameras.

How many calibration images are needed?

Because taking calibration images can be tedious, a very common question is: How many calibration images do I need? This is a difficult question to answer, but I will attempt to provide some guidelines assuming that you have already acquired *quality* calibration images. If not reread the previous article! First, it is helpful to remember that the calibration is an optimization process to find the camera model parameters using the extracted target points. So generally speaking the larger the number of points extracted the better the minimization. A large number of points can be acquired by having as many points extracted in each image as possible and/or having a large number of images. Increasing the number of extracted points per image is why it is preferable to have a calibration target that fills the field-of-view (FOV). Filling the FOV has a second benefit of ensuring that there are extracted points in the corners of the image, which will assist with the minimization of the lens distortion parameters. The only other way to add points is by taking more images. Most calibration algorithms require at least 4 images to even attempt a calibration[1]. This is clearly not enough images! Recent work using real calibration images and running 1000's of repeated calibrations yields the result that 25 images is a good *minimum*, with 100 calibration image pairs yielding the best calibration[2]. Beyond 500 images, the calibration minimization can become mathematically time consuming with only

The Art and Application of Digital Image Correlation is written by Phillip L. Reu (Phillip.Reu.DIC@gmail.com). He received his PhD from the University of Wisconsin–Madison and is currently a Principal Member of Technical Staff at Sandia National Laboratory. He began working with digital image correlation in 2004 and is focused on understanding the influence of the unavoidable compromises made in field measurements on the final DIC uncertainty. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract No. DE-AC04-94AL85000.

Conducting a Calibration

1. Pre-calibration routines
2. 2D Calibration
3. A good calibration image
4. **Stereo Calibration**
5. Sanity Checks on the Calibration
6. Care and Feeding of a Stereo-rig

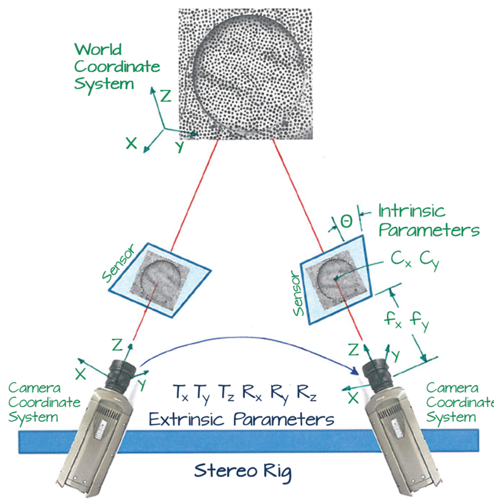

Attributes of a good calibration image

1. Adequate and even illumination
2. In focus
3. No board motion
4. Cameras synchronized
5. Control points in FOV
6. Fill the calibration volume
7. Adequate dot size
8. Take care of your calibration targets

[1] One image may yield a successful calibration if you have a 3D object with known extracted 3D coordinate points.

[2] Reu, P.L. *Experimental Mechanics*, 2013: p. 1-20.

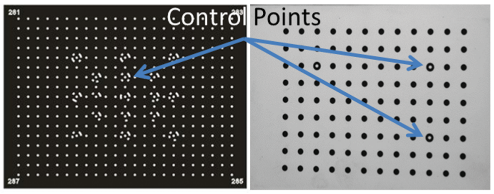

Calibration Volume

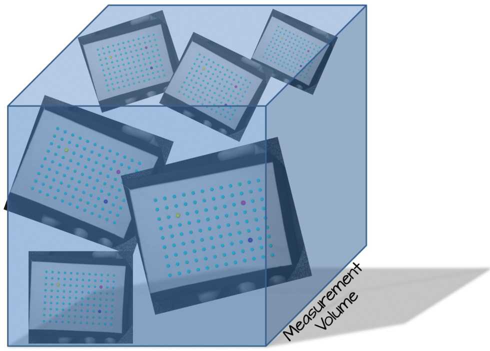

minimal improvements in the calibration results. Therefore, as a best practice, 50 to 100 image pairs are recommended. There are some situations, particularly at high-magnification, where it can be difficult to acquire good calibration images, and so in these situations, I typically cheat on the lower end with 25 images.

The final point regarding the number of calibration images is that the target motions should ''fill'' the desired measurement volume, with a good representation of angles and motions. That is, even 100 images centered and unmoving in the middle of the image will not give a good calibration. (In fact the minimization is likely to completely fail in this situation!) This indicates that without a good distribution of extracted points throughout the volume, the minimization will not work very well, and will yield poor calibration results.

What is the residual? Do I have a good calibration?

During the calibration process, the calibration quality is most often measured by calculating a ''calibration residual.'' This is illustrated in the side figure where the green cross-hairs represent the extracted pixel locations that are used to calculate the calibration parameters. Then, using the calculated calibration parameters, the extracted target locations are reprojected onto the sensor. The difference between the calculated point and the extracted point is the residual. These are summed together[3] for each image, and for all the images together to yield a calibration score. This number is a general measure of the quality of the calibration. High residuals will often indicate problems with extraction of some (or all) points, which happens when the calibration images are poor. Another common scenario is when the individual cameras will calibrate correctly but fail as a system and usually indicates a problem with the camera synchronization.

Before testing it is important that you have a good calibration. The calibration residual is the first measure of this, but probably more useful is to acquire some static images of the test object and run a quick analysis. This important step is the topic of the next article.

Next Time: Calibration sanity checks.

[3]There are a number of ways to ''sum'' the residuals together, including RMSS, averaging, and scaling by number of points.

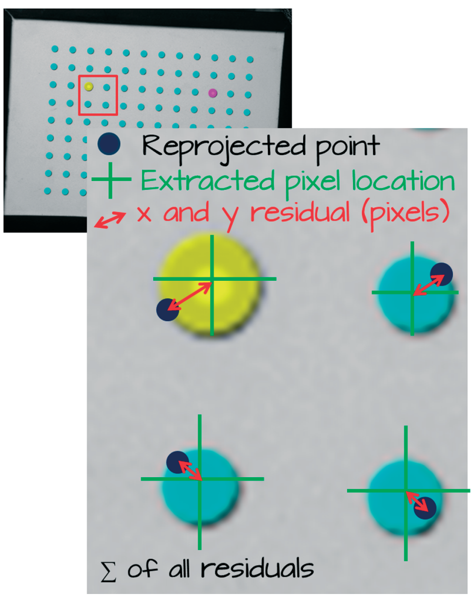

Experimental Techniques **38** (2014) 1–2 © 2014, Society for Experimental Mechanics

THE ART AND APPLICATION OF DIC

Calibration: Sanity Checks

by Phillip Reu

Introduction

The final step before initiating a costly and time consuming experiment is to examine your calibration, and to answer the question: Is this calibration good? There are a number of things that can be checked in order to have confidence in the calibration. Some of these checks will vary from software to software, with some making more or less sense. One challenge of this series is to provide best practices that cover all the commercial codes and the self-written University codes. Regardless of the code, some version of these checks should be possible because they all share a common base in the hidden components of DIC. Some software more carefully constricts what the user is allowed to do to prevent some of these issues, with some loss of generality and flexibility.

Calibration parameter check

There are a number of intrinsic and extrinsic calibration parameters that have a physical meaning that should be checked. The simplest of these is the stereo-angle. This is defined as the angle between the two cameras and should be approximately equal to the physical setup. Directly related to this is the baseline, defined as the distance between the optical centers[1] of both cameras. It is difficult to measure either of these values accurately, so they only provide a general check on the calibration.

The intrinsic parameter with the most physical meaning is the focal length. If the software displays this parameter in pixels, it may be converted to the actual lens focal length by multiplying by the pixel size in millimeters.[2] This value should be close to the specified focal length of the lens, or the focal length at the current zoom setting read from the lens. The final intrinsic parameter to check is the center pixel location. Almost always, this should be near the physical center of the detector. For instance if the detector is 1280×1024 pixels, the center pixel should be approximately 640 in the horizontal direction and 512 in the vertical direction. Remember, none of these parameters are exact. All of them are optimized and may vary from the expected exact value. However, if things are too far off, care should be taken to ascertain whether a good calibration has been obtained using other means.

Conducting a Calibration

1. **Pre-calibration routines**
2. 2D Calibration
3. A good calibration image
4. Stereo Calibration
5. **Sanity Checks on the Calibration**
6. Care and Feeding of a Stereo-rig

Sanity Checklist

1. Stereo Angle
2. Baseline
3. Focal Length
4. Epipolar Constraint
5. Point distance/object shape

[1] For measurement purposes, this is nominally the aperture or f-stop on the lens.

[2] Be careful with non-square pixels.

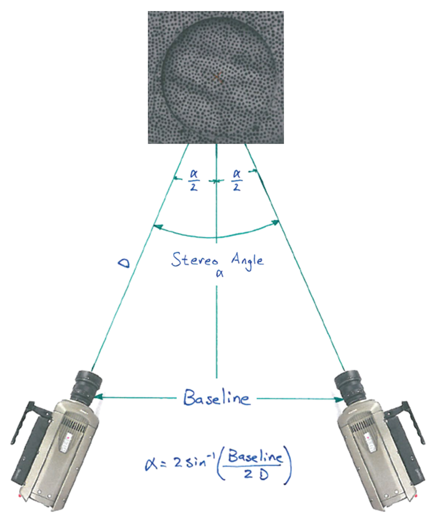

The Art and Application of Digital Image Correlation is written by Phillip L. Reu (Phillip.Reu.DIC@gmail.com). He received his PhD from the University of Wisconsin–Madison and is currently a Principal Member of Technical Staff at Sandia National Laboratory. He began working with digital image correlation in 2004 and is focused on understanding the influence of the unavoidable compromises made in field measurements on the final DIC uncertainty. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract No. DE-AC04-94AL85000.

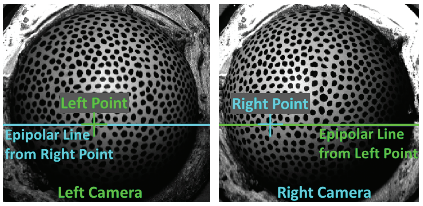

Epipolar Line

[3]A calibrated stereo-rig has the nice property that for every pixel in the left camera, all matching points will lay on a line in the right camera (and vice versa). This is illustrated above showing both epipolar lines for both points.

[4]A related parameter used by some software is the "residuum."

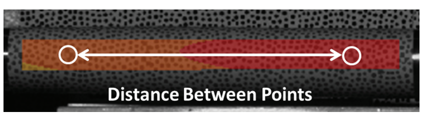

[5]In fact, you can have a great calibration score, and have a completely wrong scale factor!

Correlation check

The calibration is ***not*** complete unless a correlation is done on static images of the test object! The importance of this step cannot be over emphasized. A calibration is not known to be good unless a successful triangulation on an actual image set is completed. This step also provides a good chance to check the correlation accuracy and will indicate any lighting and speckle pattern issues before the test is conducted. The process for doing this is to acquire at least 5 static images of the test object with no loading or motion of the object. The images can then be analyzed with the DIC software to calculate the position and shape of the object. The strain can also be calculated at this point, and the noise floor will provide a feel for the accuracy of the strain calculation. Note: this is not the actual uncertainty, because it does not include some of the bias errors that may be introduced, but will indicate the lowest strain error possible with the current configuration. At the time of correlation, if there are gross problems with the calibration, the epipolar constraint[3] will be violated and the software will provide a warning. If this occurs, it may be as simple a problem as the left and right camera numbers were switched between the calibration and experiment. Note that the epipolar constraint is an important indicator of the calibration quality, as well as how well the calibration is holding up during a test.[4]

It is often helpful to check the location of the object relative to the camera system. This depends on where the coordinate system has been defined during the calibration and the correlation. It is preferable to leave the coordinate system at the cameras and then transform the results to a known coordinate system on the object, but this is a topic for a future article. If this has been done, the distance between the object and the cameras can be checked against the results.

The final check is to look at some known dimension of the object. This can be many things: radius of a cylinder, flatness of the plate, or distance between two known locations. These checks will test to make sure that the scaling and triangulation are correct.[5] There can be some difficulty in measuring the distance between points because the data point selection is only as good as a single pixel and the DIC results for distance are typically subpixel in accuracy. Therefore the comparison is only accurate to the pixel size on the object.

The stereo-rig is now setup and calibrated. The calibration and image quality have been checked using static images to find the noise floor and test the calibration. It is time to do the experiment!

Next Time: Care and feeding of a stereo-rig.

Calibration: Care and Feeding of a Stereo-rig

by Phillip Reu

Introduction

Great effort has been invested in setting-up and calibrating the stereo-rig. It should be treated with care. *Any* change to the cameras including bumping of the camera or lens could invalidate the calibration and may compromise your experiment. More than anything, the goal with a stereo-rig is to hold the cameras stationary *relative* to each other.

Maintaining the calibration

User error is often the biggest threat to the calibration. This happens when the cameras are not optimally setup; that is the cameras are not pointed correctly, the focus is off, or there is dirt on the lens or detector. A shortcut is then taken to fix the problem with hopes that the calibration will be OK. It is not. You need to recalibrate[2]. In some situations, if the lenses are left alone, a re-orientation calibration can restore the extrinsic parameters; however it would be better to re-calibrate.

Mounting the cameras together on a rigid beam, as a stereo-rig, is highly recommended and allows the unit to be carefully re-pointed, and is helpful when mounting test samples that often do not end up in exactly the same place (in a tensile testing machine for example). Re-pointing is necessary because a tight field-of-view is desirable for DIC, with the sample filling the entire image. Remember that the displacement resolution and the virtual gage size are both directly related to the number of pixels on the sample. Therefore, small movements of the stereo-rig are often helpful to center the sample in the image. With rigid camera mounting this is not a problem. In fact, with the right mounting hardware, the entire frame can be continuously moved during the experiment to track the motion of an object while keeping a tight FOV[3].

How long can one go between calibrations? How lucky do you feel? If a large number of non-violent tests (i.e. tests that are not going to shake or move the cameras) are going to be run, you can typically calibrate once a day. Better is to calibrate before the first test and then again after the final test. While it is difficult to directly compare all of the camera parameters between the two calibrations[4], the data from

Conducting a Calibration

1. Pre-calibration routines
2. 2D Calibration
3. A good calibration image
4. Stereo Calibration
5. Sanity Checks on the Calibration
6. **Care and Feeding of a Stereo-rig[1]**

[1]Tim Schmidt (Trillion Systems) is credited with this term and concept.

[2]Hopefully you have followed my earlier checklist and are not in this position.

How to care for your stereo-rig

1. Be careful with the stereo-rig
2. Don't refocus the cameras
3. Watch the epipolar constraint
4. Beware of environmental effects (heat, vibration and shock)
5. Calibrate regularly

Acceptable Stereo-rig Changes

1. Camera gain and exposure
2. Carefully moving the entire stereo-rig
3. Lighting changes

[3]Littell, J., Schmidt, T., et al., *Experimental and Applied Mechanics*, 2013, **4**: p. 111–123.

[4]This is because of the covariance of the camera parameters. The calibration process creates a situation where one parameter may increase and another decrease to compensate while still yielding the same triangulation results[5]. Generally camera parameter comparisons can only be to the level of ''nearly'' the same.

[5]Reu, P.L., Experimental Mechanics, 2013. **53**(9): p. 1661–1680.

The Art and Application of Digital Image Correlation is written by Phillip L. Reu (Phillip.Reu.DIC@gmail.com). He received his PhD from the University of Wisconsin–Madison and is currently a Principal Member of Technical Staff at Sandia National Laboratory. He began working with digital image correlation in 2004 and is focused on understanding the influence of the unavoidable compromises made in field measurements on the final DIC uncertainty. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract No. DE-AC04-94AL85000.

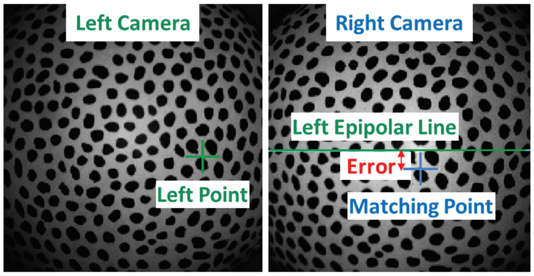

Matching point is off the epipolar line indicating a compromised calibration

the different tests can be analyzed with both calibrations and the *results* can be compared. Any drift during the day can likely be corrected using a reorientation calibration. However, as a calibration only takes 15 minutes, it is safest to calibrate before each test. For important tests a calibration both before and after the test is recommended.

Many common DIC measurements utilize material testing machines or other actuation systems where vibrations and impacts may be present. To mitigate this, taking time to consider how best to mount the camera is important. The goal is to minimize vibrations that reach the stereo-rig. Usually positioning the cameras on a tripod separate from the machine is adequate. However, for more dynamic tests, high strain-rate or impact tests for example, other mounting means may need to be investigated.

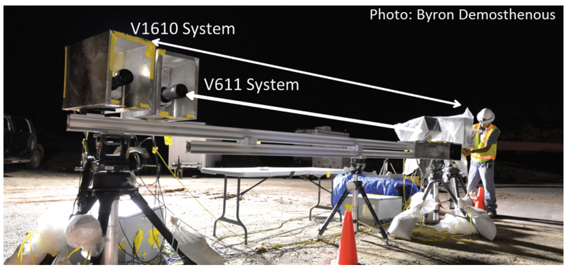

Attempt at a large scale mounting using rigid beams and mounting plates.

The epipolar constraint

The stereo-rig calibration results mathematically in what is called the ''Epipolar Constraint'' (see side figure). The result is that for any point in one camera, the corresponding point in the other camera *must* lay along a line. This is valuable because if either camera moves[6], this constraint will quickly be violated. Most software reports some version of this as part of the triangulation process and should be subpixel. A growing error during a test or between two tests indicates that the calibration is drifting and should be redone, or possibly a re-orientation should be done as provided for in many software codes.

[6]Except in the unlikely situation of translation of the image along the epipolar line.

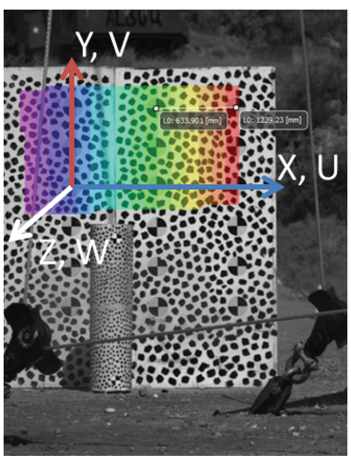

Backboard used to check if cameras have moved and provide scaling information.

Outdoor setups and continuous calibration

DIC has shown itself to be robust and flexible for camera setup. However, care should still be taken when mounting the cameras separately or on a very long beam. The best way to keep track of this is to monitor the epipolar constraint as already discussed. Indeed, it is in the large outdoor setups, where camera re-orientation is most valuable (or possibly continuous calibration). In these situations, it is helpful to have a large speckle pattern somewhere in the background that can be used for orientation calibration, as well as checking the scale. Putting fiducials on the target is very helpful for this, as you can easily check the distance between two known points to confirm your calibration. Furthermore, camera motion can be investigated by simply tracking the stationary speckle pattern in 2D: If the pattern appears to move, the camera calibration has been compromised.

Conclusions

The camera rig must be treated with care. Any motion of the cameras relative to each other will violate the calibration and produce errors in the 3D data. Over longer term testing make sure to monitor the calibration via the epipolar constraint (or other metric) and calibrate as needed. For camera mounting: Rigid is best.

Next Time:
Speckles and their relationship to the digital camera.

 Experimental Techniques **38** (2014) 1–2

THE ART AND APPLICATION OF DIC

Speckles and their relationship to the digital camera

by Phillip Reu

Introduction

Digital image correlation (DIC) was long known as ‘‘speckle correlation’’ – and that for a very good reason; the surface used for the correlation mathematics was a contrasting ‘‘speckled pattern’’. This defining characteristic continues and is a popular subject when talking with other DIC practitioners. A favorite conversation is almost always ‘‘How did you produce your speckle pattern?’’ This is understandable if one has attempted to create a high-quality pattern; particularly if their work spans various scales as each speckle size has different tricks. In the next 5 articles, I will introduce the reader to the complex and important topic of speckling: The how and the why?

Properties of a speckle

1. Size (pixels)
2. Contrast (grey levels)
3. Speckle edge sharpness (contrast gradient)
4. Speckle density and distribution (coverage)

Upcoming Articles

We will discuss the four important attributes of a speckle pattern: Size, contrast, speckle edge, and speckle density. These four features describe an ideal (or not so ideal) speckle as it interacts with the hidden components of DIC. Once these are understood, the following articles will attempt to describe how to obtain good patterns at various sizes. Many of the techniques that will be presented will work for one image resolution but not another, because the speckles will be either too large or too small. Be careful, speckles that are too small are aliased and should be avoided! Before discussing the attributes of speckles, let’s review how a digital camera interacts with the image of the speckle pattern.

Digital Cameras and Speckles

To start the conversation, it is a good idea to review a number of terms that will be used constantly in the following articles and investigate their relationship with the digital camera. The most important are contrast, grey level (counts), and noise. A digital camera pixel converts light into electrons, with a certain light intensity converted into a given number of electrons. The electrons are then ‘‘counted’’ using an analog-to-digital (A/D) converter that has a specified resolution expressed in bits. All digital cameras will have a specified number of bits of resolution, 8-, 12-, or 14-bit are the most common. This correspondingly converts the analog electrical signal into a digital signal with a given resolution, 256 counts (8-bit), 4096 counts (12-bit) or 16384 counts (14-bit)[1]. Digital cameras also have ‘‘gain’’

[1] Dynamic range is the ratio of light intensity between the pixel full-well capacity and the readout noise. This range is then digitized by the A/D converter with the number of available bits. However, the dynamic range is usually limited by the pixel rather than the A/D properties. (http://www.cambridgeincolour.com)

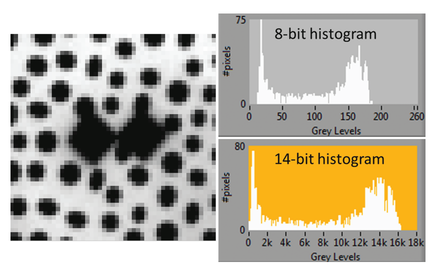

Speckle image and the corresponding histogram for both 14- and 8-bit image types.

The Art and Application of Digital Image Correlation is written by Phillip L. Reu (Phillip.Reu.DIC@gmail.com). He received his PhD from the University of Wisconsin–Madison and is currently a Principal Member of Technical Staff at Sandia National Laboratory. He began working with digital image correlation in 2004 and is focused on understanding the influence of the unavoidable compromises made in field measurements on the final DIC uncertainty. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy’s National Nuclear Security Administration under contract No. DE-AC04-94AL85000.

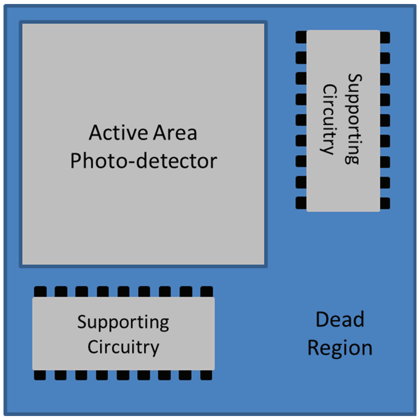

Schematic of a pixel (CMOS)

Typical Fill-Factors

Camera Model	Fill-Factor
Typical Machine Vision	>90%
Phantom V1610	64%
Photron SA-X2	58%
Shimadzu HPV-2	13% & 76%

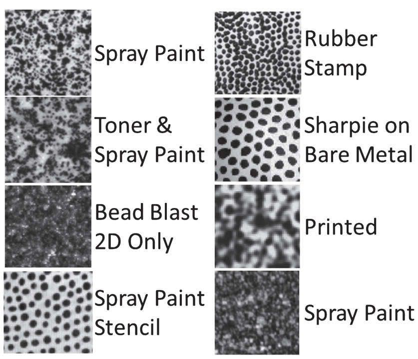

Examples of speckle pattern applications

which is supplied by an amplifier, much like the volume control on a music player. This gain (ISO selector on a point-and-shoot camera) can be used to increase the number of counts registered by the camera by amplifying the analog electron signal. Higher gains will not only amplify the signal, but unfortunately will also amplify the noise. So gain should be used with care.

Another camera attribute, particularly important for ultra-high speed imaging, is the camera fill-factor. Fill-factor is defined as the amount of the CCD or CMOS pixel that is light sensitive, in other words how much of the pixel area will convert photons into electrons. Typically the machine vision cameras used for DIC have a high fill-factor and this is not an issue.

The purpose of the speckle pattern, in cooperation with the lighting, is to create a high-contrast low noise speckle pattern in the digital image. It cannot be emphasized enough that lighting is as important as the speckle pattern in creating a great image! During the experiment, a general feel for the contrast can be obtained by observing the number of counts in your image between the white and black regions. The image histogram is also an indication of the contrast that quickly shows the range of counts that are contained in an image (figure on previous page). Be careful with this in that only the contrast in the speckle region to be correlated is important. Many times there will be bright and dark areas outside of the speckle region that may indicate a higher contrast range in the image than is contained in the speckles themselves. A few sample speckle patterns are illustrated in the side figure. How these were created and why they are good or bad is the topic of the upcoming articles.

Speckling Technique

The later speckling articles will share many of the techniques that have been discovered over the years that yield good speckle patterns. Any and all suggestions in this realm are welcome; please email me at Phillip.Reu.DIC@gmail.com with ideas. If I can use your ideas, full credit will be given. Remember that when generating speckles, the size, contrast, spatial distribution, and edge profile are all important.

Next Time: All about speckles: aliasing.

THE ART AND APPLICATION OF DIC

All about Speckles: Aliasing

by Phillip Reu

Introduction

Optimizing the speckle size is critical. The challenge in doing this is to make them as small as possible, without making them too small. Practically, this is very difficult and is why at every size scale, a new speckling technique is likely to be required. So how do we determine the optimum speckle size?

Properties of a speckle

1. **Size (pixels)**
2. Contrast (grey levels)
3. Speckle edge sharpness (contrast gradient)
4. Speckle density and distribution (coverage)

Estimating Speckle Size and Aliasing

A speckle must be at least 3 pixels in size, but this doesn't tell you how big it should physically be on the sample in millimeters. The physical size of the speckle can be found by first measuring the field-of-view (FOV) in millimeters and then knowing how many pixels you will have across the FOV to determine pixel size at the object (mm/pixel). From this the physical size of the speckle is calculated by multiplying the pixel size by a factor of 3 to calculate the minimum speckle size. However, *if* the minimum speckle size is used, the FOV cannot be increased, even a little, as this will make the speckles too small. FOV can be decreased, but will negatively impact the optimum speckle size by making them too large and wasting spatial resolution. Of the two scenarios, the second is the safest.

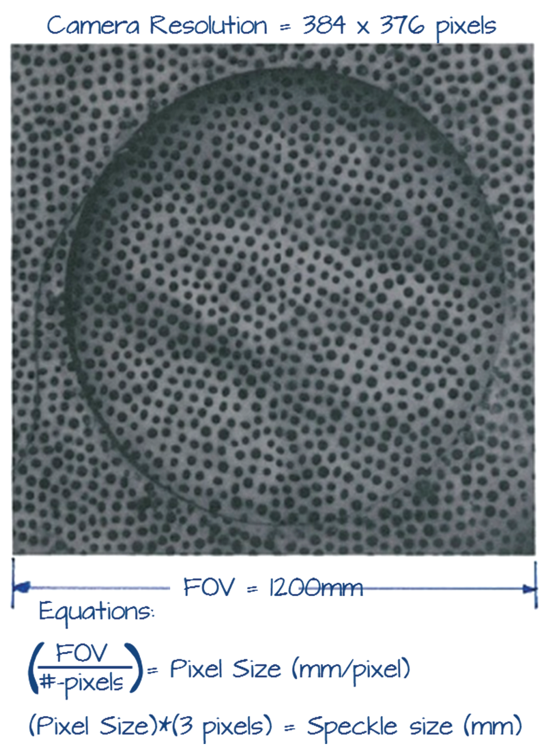

The 3-pixel requirement is a clear result of the sampling theorem. The theorem says that in order to unambiguously represent a signal, it must have at least 2 pixels (or samples) per period. For the case of a speckle; that is 2 pixels across the white and 2 pixels across the black. An astute reader may then argue that ≥ 2 is the optimum, however, there is also a practical reason for a speckle size of 3-pixels and greater; and that is because it is impossible to determine if a speckle is 2-pixels or smaller when viewing a live image on the computer. Why? Because of aliasing! Once a speckle, or any signal for that matter, is digitized, it is not possible to determine if it *is* aliased. The true information is lost and cannot be recovered in software post-processing. This can practically be illustrated by visualizing how

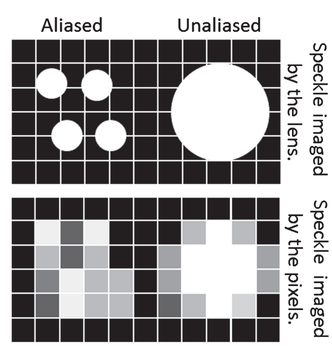

Illustration of (left) aliased and (right) unaliased speckle.

The Art and Application of Digital Image Correlation is written by Phillip L. Reu (Phillip.Reu.DIC@gmail.com). He received his PhD from the University of Wisconsin–Madison and is currently a Principal Member of Technical Staff at Sandia National Laboratory. He began working with digital image correlation in 2004 and is focused on understanding the influence of the unavoidable compromises made in field measurements on the final DIC uncertainty. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract No. DE-AC04-94AL85000.

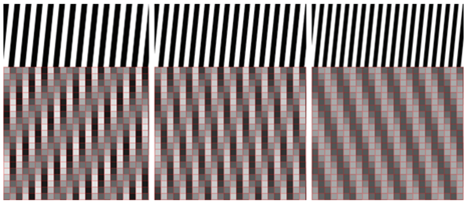

Illustration of aliasing using tilted lines of 3 frequencies as sampled by a CCD. (http://photoacute.com/)

[1]In fact most commercial cameras come with an antialiasing filter which adds just a little blur at the pixel level to filter out any frequencies that are too high. Machine vision cameras don't usually have anti-aliasing filters.

[2]Reu, Experimental Mechanics not yet published.

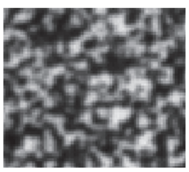

Good Pattern:
Actual Speckle Size = 4 pixels
Calculated Size
= 4.3 (autocorrelation)
= 2.9 (autothreshold)
= 3-4(Eyeball)

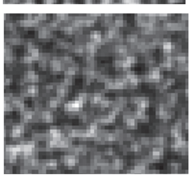

Aliased:
Actual Speckle Size = 1.5 pixels
Calculated Size
= 4.5 (autocorrelation)
= 2.2 (autothreshold)
= 2 to 3 (Eyeball)

Obtained from printing two different pattern sizes and acquiring both at the same time (Courtesy Pascal Lava – KU Leuven).

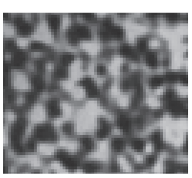

Good Pattern:
Actual Speckle Size = 4-5 pixels
Calculated Size
= 4.8 (autocorrelation)
= 3.6 (autothreshold)
= 3-5(Eyeball)

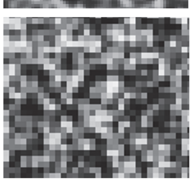

Aliased:
Actual Speckle Size = 1.5 pixels
Calculated Size
= 2.1 (autocorrelation)
= 1.6 (autothreshold)
= 2 to 3 (Eyeball)

Obtained via decimation of a higher resolution image.

[3]Reu, P., Experimental Mechanics, 2011. **51**(4): p. 443–452.

a speckle appears to the pixel. In the side figure the white blobs illustrate a one-, and three-pixel speckle as imaged by the lens onto the pixel array. For the one-pixel speckles, it is difficult to unambiguously locate the center of the speckles to a single pixel (let alone subpixel!). However, once a speckle reaches 3–4 pixels, the outline can be clearly seen in the pixel array. Also note that for the aliased speckle you still see a ''speckle'' like pattern in the digitized image. The reason for this is illustrated in the diagonal-line aliasing example (side figure). The signal to be sampled at the top clearly shows diagonal lines, but once sampled and aliased, the directionality and even the frequency are modified by the under-sampling of the detector grid. So the danger is not that the aliased speckle pattern will remove the speckles, they still appear as ≈2-pixels, but they are not correctly sampled by the detector and will lead to errors in the DIC measurement.

A speckle size of 3 is also recommended because it is very difficult when applying speckles via many techniques to *exactly* size the speckles. Because of the sizing issues it is safer to slightly oversize the speckles so any slightly smaller speckles are still greater than 2 pixels. An important point to remember is that the ''speckle'' includes both the round blob, and the contrasting region around it. I suppose this may be grouped under the speckle density topic, but it needs to be noted that neither the gap between speckles, nor the speckles themselves should be less than 3 pixels.

Unlike in temporal sampling, there are not many good ways to prevent aliasing with machine vision cameras[1]. In 2D DIC, you could theoretically use the lens as an antialiasing filter, by either the focus or the fundamental lens resolution (modulation transfer function[2]). However in stereo-DIC this is nearly impossible to do due to the inevitable uneven defocus.

Illustration of Aliased Speckle Patterns

To the left are two sets of speckle patterns that illustrate both an aliased and unaliased representation of the same pattern. The first one was created by printing the same pattern at two different sizes and imaging them simultaneously. The unaliased speckles are approximately 4 pixels and the various size measurement techniques (discussed in the next article) are able to correctly determine the size. Once aliased however, the speckle size measurement does not give reliable results. The second illustration was created by taking a high resolution image and decimating it to create an aliased speckle condition[3]. This allows exact control of the final speckle size. Again, the unaliased image has clearly defined speckles that are easy to see. The aliased pattern also has ''speckles,'' but both the eye and the automated methods do not correctly determine the speckle size. It is important to train your eye to recognize the aliased speckle condition. Go into the lab and create some aliased patterns and look at them at different magnifications. You will see them ''twinkle'': They just don't look right . . .

Conclusions

Aliased speckles should be avoided at all costs. Any aliasing in the pattern will add noise to the image, of one sort or another, and will compromise the accuracy of the measurement. This not only applies when all the speckles are aliased as has been illustrated here, but when there may be large number of undersized speckles in amongst the many correctly sized speckles. These small speckles will cause issues and may influence the minimum subset size. This is often the greatest difficulty in using spray painted speckles. It is very difficult to optimize the spray to have a tight distribution of speckle sizes without any that are too big or too small. Remember, speckles must be greater than 3 pixels in size.

Next Time: Speckle Size Measurement.

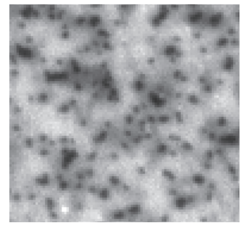

Questionable Pattern:
Calculated Size
= 8.6 (autocorrelation)
= 3.7 (autothreshold)
= 1-9(Eyeball)

Questionable speckle pattern with small speckles amongst nicely sized speckles.

Exercise for Reader
Take an image with speckles that are approximately 3 pixels in size. Translate the speckle pattern in 2D over a couple of pixels while imaging. For each image take a full resolution image and one with the Binning operation on the camera to create aliased speckles. If binning is not available, this can be done by printing two different speckle patterns, or by numerical binning of the full resolution image[3].Analyze the results.

All about speckles: Speckle Size Measurement

by Phillip Reu

Introduction

The previous article introduced the topic of speckle size and covered the very important issue of aliased speckles. It is difficult to tell once imaged, whether your speckles are aliased, but if you have *guaranteed* that they are not, it is helpful to measure the average speckle size and distribution. This article discusses methods of doing just that.

Properties of a speckle

1. **Size (pixels)**
2. Contrast (grey levels)
3. Speckle edge sharpness (contrast gradient)
4. Speckle density and distribution (coverage)

Speckle Size Measurement

To measure the speckle size in the image there are three common approaches of varying mathematical complexity: Autocorrelation, segmentation/blob analysis, and ''by eye.'' The standard method is the autocorrelation approach[2]. For nearly uniform speckle patterns this is a good approach and will quickly give you an estimation of the average speckle size. A more powerful method, however, may be ''blob'' detection, which is a standard image processing methodology available in LabVIEW, MATLAB[3], and many other image processing toolboxes. An example of this is shown in the side figure with a speckle pattern that has a large variation between a big speckle (a fiducial) and the rest of the pattern. As can be seen, the blob analysis is able to pick up these variations by reporting the distribution of speckle sizes and the percent coverage of the sample and will be important when we discuss speckle distribution. If interested, Lecompte applies this method using MATLAB[3]. The simplest and quickest method, as well as the one most likely to be used during testing, is ''by eye''. With some training of the eye, you can easily measure the speckle size and distributions by observing the pattern. To add a more quantitative approach, a representative speckle or two are found and measured in the image acquisition program. The acquisition program will display[4] both the current cursor location and the grey level at that position and allows one to easily determine both the speckle size and the contrast of the pattern simultaneously

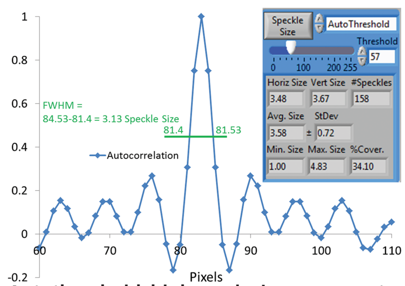

Autothreshold blob analysis versus autocorrelation analysis.

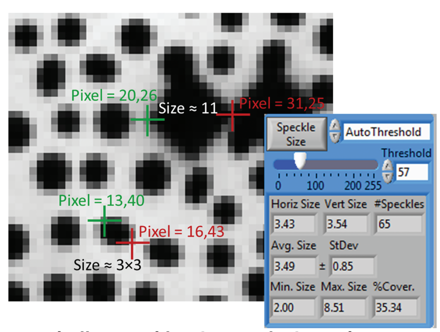

''Eye-ball'' Speckle size analysis and comparisons to autocorrelation and particle analysis.

The Art and Application of Digital Image Correlation is written by Phillip L. Reu (Phillip.Reu.DIC@gmail.com). He received his PhD from the University of Wisconsin–Madison and is currently a Principal Member of Technical Staff at Sandia National Laboratory. He began working with digital image correlation in 2004 and is focused on understanding the influence of the unavoidable compromises made in field measurements on the final DIC uncertainty. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract No. DE-AC04-94AL85000.

[2]Sutton, D.A., J.J. Orteu, and H.W. Schreier, *Image Correlation for Shape, Motion and Deformation Measurements*. 2009, New York, NY: Springer.

[3]Lecompte, D., et al., *Optics and Lasers in Engineering*, 2006. **44**(11): p. 1132–1145.

[4]Or it should!

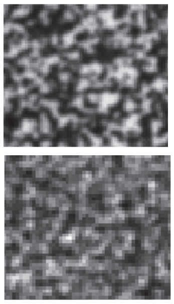

Good Pattern:
Actual Speckle Size = 4 pixels
Calculated Size
= 4.3 (autocorrelation)
= 2.9 (autothreshold)
= 3-4(Eyeball)

Aliased:
Actual Speckle Size = 1.5 pixels
Calculated Size
= 4.5 (autocorrelation)
= 2.2 (autothreshold)
= 2 to 3 (Eyeball)

Illustration of an aliased pattern and errors in extracting the correct speckle size.

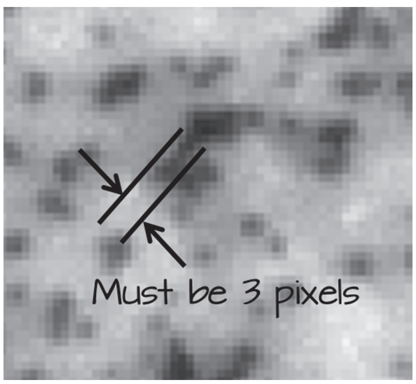

''Line'' type feature in a speckle pattern. Note: This applies to both the white and black regions.

[5]See the ''Hidden Components of DIC'' in the previous articles.

(see the annotated figure). There are two cautions with this approach; first, be careful that your viewer software does not filter the image before displaying. Possible display filtering can be determined by ensuring that when zoomed in, individual pixels can be seen. The only other time measurement by eye is difficult is when the speckle pattern is exactly at the threshold of 2 pixels. In this case, it is difficult to find the speckle edge and the other methods are typically more accurate.

As a reminder, once a pattern is aliased, measuring the speckle size cannot be done accurately by any method! This is demonstrated in the side figure with an aliased speckle pattern and the corresponding measurements of speckle size. The best that can be done is to train the eye to see this situation. The speckles do look different; they ''twinkle'' and seem to be poorly defined.

Speckle Shapes and Aliasing

There are only two real requirements for the speckle shape: the smallest feature must still be unaliased and there should be no directionality in the pattern. An example of this is the small line like feature illustrated in the figure. Directionality in speckling will be discussed thoroughly in a following article.

Speckle Size Defines Spatial Resolution

Why is there always an attempt to make the speckles as small as possible? Why not just make them relatively large, say 10 to 15 pixels? The reason is spatial resolution. Remember that the subset/facet size fundamentally limits the ability of DIC to measure shape or displacement or strain gradients. And, unfortunately, the subset limitation cannot be escaped, because to measure subpixel motion, a region of the pattern must be used for correlation[5]. To improve subset matching, there must be enough contrast in the matching region to yield a good signal for the minimization process. To do this their needs to be 2–3 speckles within a subset. I think you see where this is going:

1. To capture gradients – Small subsets are desired,
2. For small subsets – Small speckles are needed!
3. **But**, speckles must not be too small.

Conclusions

Remember this is your *smallest or largest* speckle. Not the average speckle size. If your subset/facet contains enough appropriately sized speckles, a few undersized speckles can probably be tolerated, but will certainly add noise to your measurements! It is best to get the speckle size correct for your field-of-view and magnification early in the experiment. Spray paint, probably the most popular method of speckling is hard to control and small speckles will often result. Be careful. The speckles must be 3 pixels as a best practice including both the white and black regions.

Next Time: Speckle contrast.