

# Development of an Analysis Engine for High Precision Slope Measurement

Ben Sheff, Jun Qian, Joseph Sullivan, Lahsen Assoufid

August 4, 2015

## Abstract

The Advanced Photon Source (APS) is a 7 GeV storage ring based hard x-ray synchrotron radiation source funded by the US Department of Energy Office of Science-Basic Energy Sciences scientific user facility, and located at Argonne National Laboratory, Lemont, Illinois, USA. The APS is currently developing plans for a major upgrade to improve x-ray beam brightness and coherent flux by 100 to 1,000 times over current values. In order to preserve these qualities, improved optics are necessary to carry out cutting edge experiments. The Optics group mission in the X-ray Science Division at Argonne focuses on designing and fabricating optics in support of the APS operations and upgrade.

Simulations have shown that the APS beamlines upgrade will require mirrors with slope error as low as 50 nrad. The optics group has commissioned an autocollimator-based slope measurement device (AC-LTP) to characterize mirror quality to meet future requirements. However, the autocollimator has a very limited angular measurement range, which presents a challenge when measuring highly curved mirrors. This project focuses on developing analysis tools and exploring advanced techniques for measuring these mirrors with a high degree of precision.

In developing these tools and techniques, new software was created to handle the data analysis and processing. OMEN (Optical Metrology ENgine) was developed as an intuitive interface to easily view plots of any data collected, and provides for a variety of analysis options that dynamically manipulate and display data. To accommodate non-linearities presented by measurements of highly curved mirrors, it incorporates options both for stitching together subaperatures, and for the use of calibrations curves. In this report, the software will be described in detail along with its application to real data obtained from flat, spherical, and elliptical mirrors.

## Introduction

The APS autocollimator-based long trace profiler (AC-LTP) is an advanced slope measurement device used to make surface slope measurements of high precision x-ray synchrotron radiation mirrors [1, 2]. The system was designed to accommodate multiple sensors and measurement probes. It is currently configured in the NOM concept [3], using an Elcomat 3000 autocollimator [4] and a mirror-based pentaprism combined with an iris providing a 2.4 mm diameter probe beam to scan a mirror surface and measure a line slope profile along the scan direction. The system is designed to measure mirrors up 1.5 m long with about 50 nrad resolution on a flat shape. To reduce errors from noise, several scans are typically acquired and analyzed and averaged. As a result, it outputs a large amount of data over which several routine processes will usually be done. Autocollimator sensors are known to have a limited angular range with a non-linear response to large angle measurements. This makes measurements of highly curved surfaces challenging to measure.

The goal of this software is to alleviate both the above challenges. This software is designed to allow a user to access his or her data easily and to complete a variety of standard analysis actions while dynamically seeing the effect, all in an intuitive format (figure 1). It is configured keeping in mind other analysis programs used within the optics group, and has the capability to interface using a few standard formats. This software also introduces analysis tools to allow investigation into curved mirror measurements. It has methods to add calibration curves and to stitch partial scans together to allow for rapid, accurate handling of highly curved mirror data.

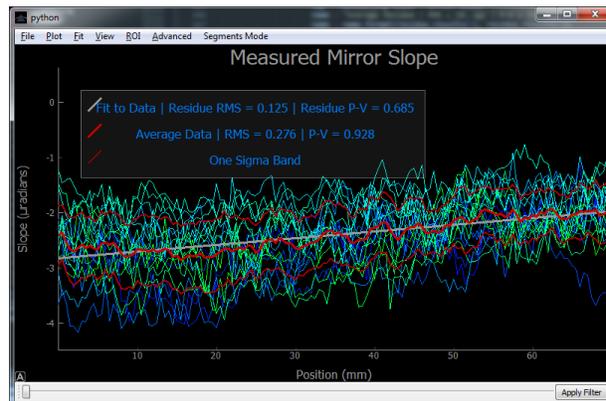


Figure 1: OMEN basic interface with twenty scans of a flat, 70 mm mirror loaded

## OMEN Fundamentals

### Basic Interface and Initialization

The interface, named the Optical Metrology Engine (OMEN) was designed using the package PyQt4 for Python 3.4. As a PyQt4 based Graphical User Interface (GUI), it has three primary components: The “central widget” of the window is an object that handles all the plotting, which allows the bulk of the window to be filled with the plot of interest for a user’s convenience; this is handled within the package PyQtGraph. The “Toolbar” of the window is an object containing the controls for the filtering options (see Filtering), set to sit at the bottom of the plot. Finally, the “Menubar” of the window is a set of drop down menus to allow selection among a variety of analysis options and tools.

The program as a whole follows a particular logic in its design. Once the desired dataset is loaded, the user is expected to select some analysis option, tune that, apply it, and then proceed to the next desired option. After tuning the choice, OMEN sends a command to Raven to apply the selected option to the data, and then the plot is updated to reflect this change.

Once a user has selected a file or set of files of

interest, a Raven object is created. Raven is the primary backend class, handling most of the data analysis that can be done on a single scan or collection of repeated scans. Raven figures out what type of file was selected and, if it is a collection of raw files from the AC-LTP, it collects all files of the same name up to the scan number appended in the mda2ascii code. Raven saves the data to the relevant parameters internally, and if not specified in the file, calculates the remaining ones. These include:

- The mirror slope at each point, almost always specified in the given data file.
- The mirror height at each point, found by subtracting the average slope and integrating the slope data by simple left-handed Riemann sum
- The average value of the slope and height at each point, calculated by finding the average at each position of all the scans read
- The standard deviation of the slope and height at each point, calculated by finding the standard deviation at each position of all the scans read
- The fit function, described later. It is initially set to just be the data itself, but in OMEN is re-initialized to a circle

With the backend set up, OMEN reads out the data and provides options in the plot menu for what the operator would like to plot. There are also a variety of options in the view menu for adjusting the area of the graph shown, and more in the right-click menu of the plot, or through the right click menu available in the plot itself.

## Object Oriented

OMEN itself is a class. Each OMEN object creates and manages an analysis engine window. It also handles objects of the companion classes developed for this software. Raven is the primary

backend class. It handles data collection along with the bulk of analysis functions, described below. Lightning is used primarily for stitching segments together. It collects overlapping sets of data and fits them to each other, collecting an average plot. Finally, Medium is a minor class developed to handle input and output into a few standard forms. It is used mainly in Raven. This structure makes it fairly modular, allowing fairly easy modification of the code for possible future needs by adding new classes or methods. Many of the methods are also designed with this in mind, for example, by just adding a method in Raven to do fitting, and making an appropriate method and button in OMEN, it is fairly easy to add in other functions to the fit options.

## Main Functions

### Filtering

Often when running the long trace profiler, there will be some disturbance during a given run. Be it a large cart passing by causing unwanted vibrations, a pocket of air passing by, or a temperature fluctuation, these will usually ruin a given scan, creating a scan that deviates significantly from the shape of the mirror, as shown in some of the traces visible in figure 2. The purpose of the filter is to remove any such scans.

The filter operates by integrating the standard deviation of each profile. This is done by taking the active data (slope or height) of each individual scan, finding the number of standard deviations it is away from the average at each point, and if this is greater than one, adding it to a running total. This total is then compared to a number proportional to the number of data points per scan, and inverse to a user-specified value set by moving a slider at the bottom of the window. Any scans with a total deviation larger than this number are tagged to be removed from the data by removing them from an internal list of approved scans. For speed, the deviation total for each scan is pre-calculated and only updated when necessary, allowing an operator to dynamically adjust the filter by moving the slider and seeing what scans are removed before applying the filter. When an operator chooses to apply their filter, Raven will actually remove any tagged scans from its available data, and update the average and standard deviation data appropriately. This allows for iterative filtering as well.

### Fitting

Often the data collected will have some trend, indicating curvature in the mirror, problems with the data collection, or deflections in the mirror. It is useful to be able to find these curves, and

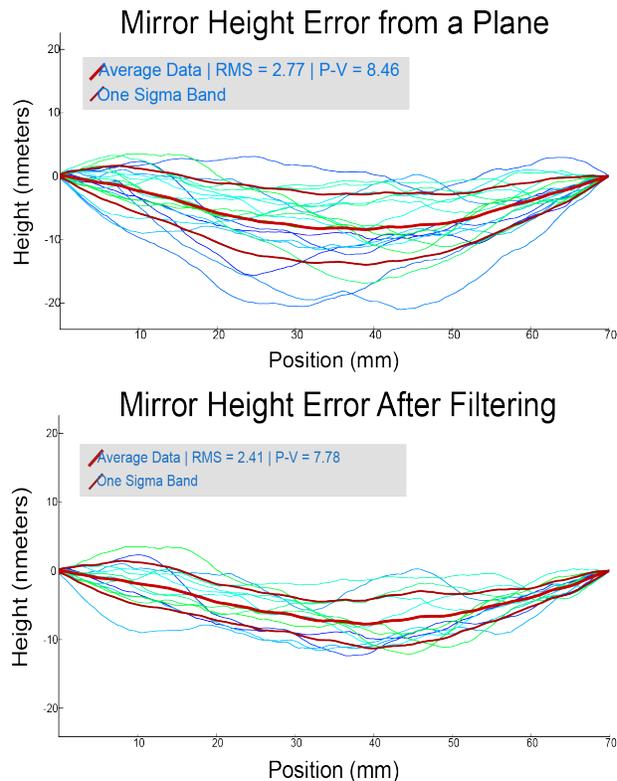


Figure 2: Example of filtering based on deviation from the average mirror height profile

the residual profile after subtracting the curves. OMEN includes a variety of curve fitting options, currently it can do up to fifth degree polynomials, circles, and it can fit the average profile from a different scan. This is demonstrated in figure 3, with a parabola fit to a height profile. The former two fitting options are handled entirely within Raven using the scipy package to do fitting. Fitting on another scan's profile is handled using a combination of Raven methods and Lightning methods in order to prevent redundancy, as the latter is designed to fit data sets to each other. Once a fit is selected, the plot will update to reflect the fit to the average run. It will also generate the residual data for each run after subtracting the fit to that run, or at the operator's discretion, it can subtract the average fit instead. If the residual data is plotted, it assumes the same plot options as the

regular data, and if the average is included, it will mention in the legend the root-mean-square and peak-to-valley values. Once a suitable fit is found, statistics about it are available in a few forms. There is an option to read the fit out into a pop up window for easy access, along with an option to save it to a text document with statistics about the fit, and finally within the legend of the plot the root-mean-square and peak-to-valley of the residual data are listed alongside the fit.

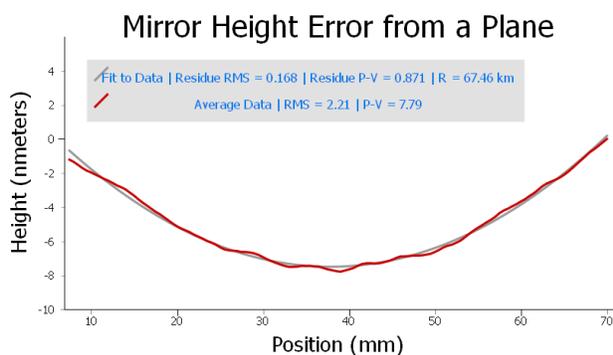


Figure 3: OMEN basic interface with twenty scans of a flat, 70 mm mirror loaded

## Region of Interest

Focusing on a region of interest can be a necessity for analysis, so it is given a menu in OMEN. This functionality allows the operator to manipulate cursors overlaying the plot, as shown in figure 4. These cursors specify two x-positions clearly visible on the screen. When the region of interest is applied, Raven will remove all data outside the region defined by the cursors. By default, the cursors cover the center third of the active data set.

## Calibration

The autocollimator in the AC-LTP has a limited angular measurement range and non linear response to large angles. Its response also depends

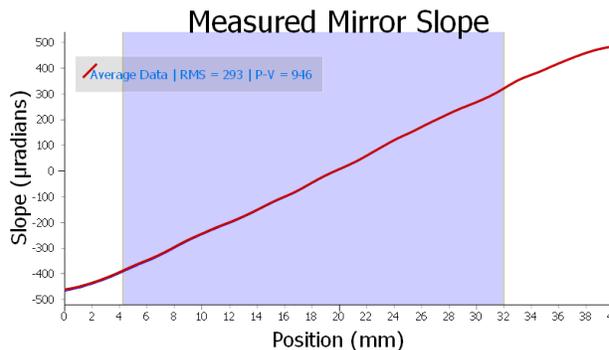


Figure 4: Selecting a region of interest, here focusing where spherical aberrations are small.

of the probe optical path length. Therefore calibration is necessary to correct for measurement errors. OMEN incorporates this functionality, allowing a user to add a calibration run and apply it to the data. The calibration run is read as a new Raven object, which is inverted and fed into a Lightning object for interpolation. This forms a look-up table that is used to re-evaluate all the slope values in the current Raven object. When in segments mode, it will apply the calibration to all segments, taking care to adjust the original data, without the offset induced to match up segments to each other.

## Stitching Partial Scans

Stitching partial scans is conducted by partitioning the mirror into small segments, over each of which the slope varies little and the AC-LTP is known to be accurate. For each segment, the stage is tilted until the AC-LTP reads the center of the segment as zero (horizontal), and then a scan is conducted. By default, each segment has a 50% overlap with its neighboring segments. This allows for a series of measurements that together cover the mirror, and have sufficient overlap with each other to stitch them together to make a proper full scan available, as shown in figure 5.

The stitching of such segments is the main pur-

pose for Lightning. As of the first time you enter Segments Mode, any further data loaded will be considered a partial scan, or segment, and appended to the running list. There is, however, accounting within Lightning that will remove anything it thinks might be redundant with the added segment (namely anything that either completely contains or is completely contained in the new segment). Whenever a segment is added or modified, Lightning will go through all the segments from left to right, and fit each one by vertical translation, and to a small degree horizontal translation, to the one on its left. It then stores a conglomerate run that is the average at each point of the fitted segments. The translations found by the fits is read by OMEN and forwarded to the Raven objects storing each segment, and a new Raven object is made using the data in the conglomerate file, representing the full scan stitched together from the available segments. This allows the operator to do any of the analysis they would do to a single-scan profile to the stitched profile. Whenever any change is made to the stitched data, all of the segments are updated to reflect the change to their region of the scan, and the respective Raven objects are modified appropriately. Further, by returning to the default, Single-Run Mode, the operator can continue to conduct any analysis they would do on individual segments, and the full stitched version will update automatically to reflect the changes.

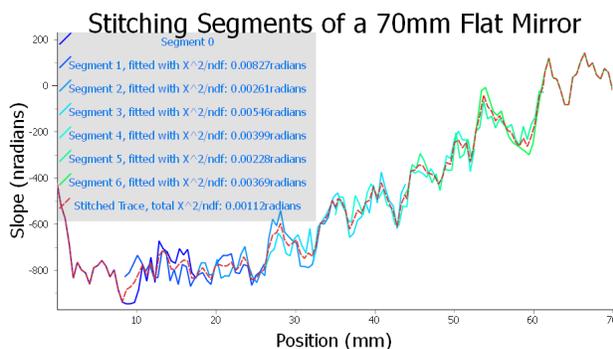


Figure 5: Stitching together partial scans of a mirror.

While in Segment Mode, OMEN also modifies the available menus to reflect the possible options for plotting the stitched version. This includes changing the plot menu dramatically, and changing the save function to save files for all the segments along with the full stitched run. All of the functions that allow for live-updating the full stitched data do come at a significant cost to speed. For complete rigor, every time a segment is added or modified, the full collection is re-fitted and the full stitched scan, re-calculated. For large files or large numbers of segments, this can be a significant time sink as the code repeatedly does large numbers of fitting.

## Advanced Functions

There are a few options that are desirable in an analysis engine, but don't fit well with the other methods, so are grouped together as "advanced" functions. Specifically there are six of these:

1. Re-Calculate Height: Sends a command through Raven to re-do the integration that found the data on the mirror height.
2. Re-Calculate Slope: Sends a command through Raven to differentiate the height data to find new data on the mirror's slope
3. Zero the Position: Sets the center of the active data set to zero along x
4. Use Average Fit: Sets the fitting scripts to only pay attention to the fit to the average scan, rather than finding residuals of fits to each individual scan
5. Apply Calibration: This is described earlier, but is included in the Advanced Functions menu
6. Select Segment to Left/Right: If in Single-Run Mode, and Segment Mode has been initiated, it makes the active segment the one left or right of the active one.

## Known Systematic Error Sources

The integration and differentiation methods for calculating height from slope data and vice versa respectively, are possible sources for error. They are limited by the sampling rate given in the data, and in the former case do a simple left-handed riemann summation without any attempt at interpolation. Differentiation has similar limitations, as it just finds the slope between two height data points. Further these two methods are not exact inverses of each other, so there is cumulative distortion as the user recalculates one from the other and back. This is alleviated currently by having frequent recalculation be unnecessary, as Raven stores both sets of data and will only recalculate on command, and it is rare a user would have a need to recalculate height or slope more than once.

The fitting method for the circle is still fairly inaccurate. It will frequently get caught on local minima for the chisquare, and usually the parabola fit is more accurate and useful. There are occasionally similar problems with the fitting involved in stitching and in fitting one profile to another. The code will output a warning if a problem has occurred in the latter case, but the circle has no such failsafes. It is currently only managed by the user looking at the fit plot and seeing whether it looks reasonable.

The averaging script assumes that each scan has identical sets of x-coordinates. Since this is not strictly true, it could be a source of error as the averaging script is then averaging data from different sample locations, as opposed to repeated measurements of the same point. This is likely a very small source of error though.

Finally, interpolation in Lightning is done by simple linear interpolation. This is used for calibration, mirror profile comparisons, and for segment stitching. There exist more faithful interpolation techniques that can be added later by simply changing out the LinearInterpolation

method of the Lightning class. Due to time constraints, these have not yet been implemented.

## OMEN Application

### Data Collection

A few typical scans were run on flat mirrors to acquire data shown earlier in the paper. This demonstrated the power of the new analysis software, and allowed for testing and refining it.

The AC-LTP was then set up to run a segmented scan as described in the Stitching Partial Scans section. It was run over a flat, an elliptical, and a spherical mirror to test the effectiveness of this method. Each was partitioned into seven segments, which each had 50% overlap with each other. These were fed into OMEN and each compared with a continuous, complete scan of their respective mirrors. The results of these measurements, and their comparisons, are shown in figure 6.

The optics group also recently acquired a high precision stage to tilt mirrors with a precision on the order of 10 nrad[5]. This allows for an accurate investigation into the large angle regime against which the AC-LTP can be calibrated. After installation of this stage, a stability scan was run simulating the conditions of a mirror measurement. The scan showed an measurement RMS near 50 nrad, which is within the normal operating parameters of the equipment. This shows that the high precision stage can be used to investigate the large angle regime and create a calibration curve. This was tested, making a calibration curve for the range from 0  $\mu$ rad to 360  $\mu$ rad. The curve was measured in half microradian increments, and averaged over twenty scans. In order to get a better handle on the data, a linear fit was subtracted from this curve, leaving the residue shown in figure 7. The residue has significant oscillations, which are a known side effect of the type of autocollimator used [6].

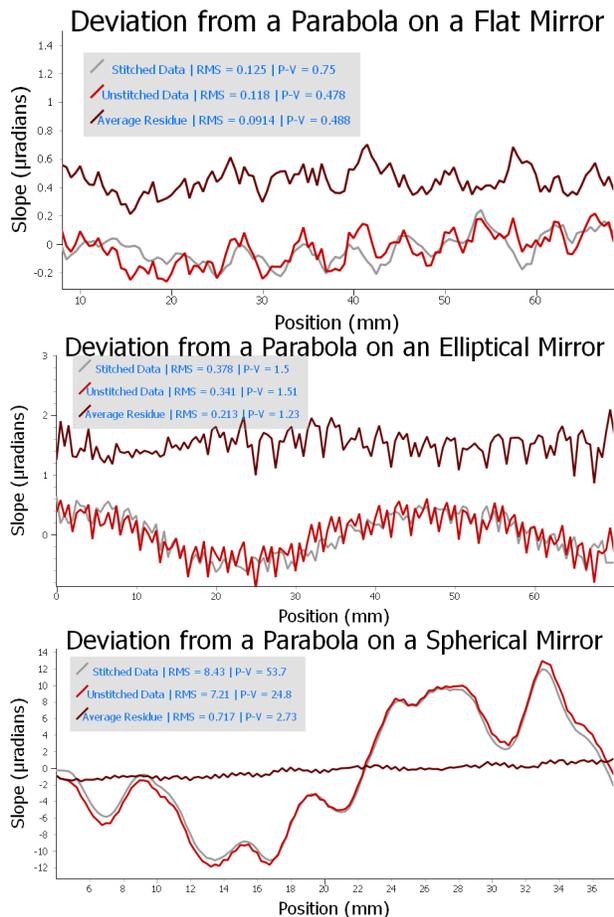


Figure 6: Stability scan simulating a typical run, taking simulated scans and averaging

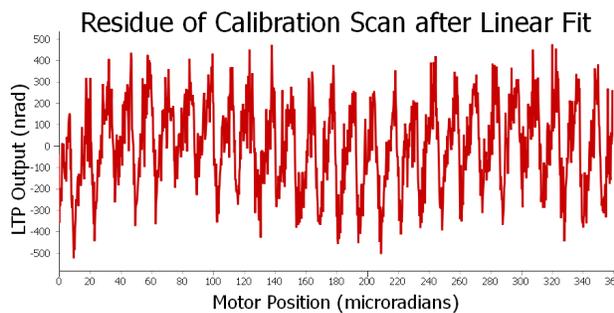


Figure 7: Residue of calibration scan after subtracting a linear fit

The stitched measurement for the spherical mirror was then compared to the calibrated region of the full scan. The scan first had the calibration curve applied and the overall linear fit was subtracted to more easily see the effects. This residue was compared to equivalent residue of the stitched scan, as above, with the region of interest specified to be the area where the calibration curve applied. The two plots are available for comparison in figure 8.

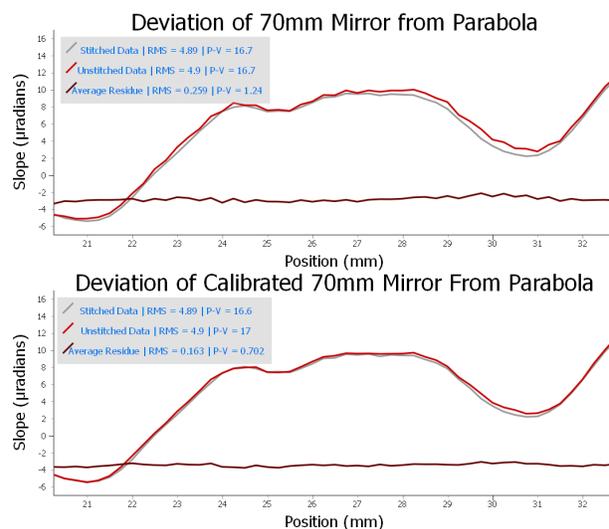


Figure 8: Comparing calibrated and non-calibrated data to stitched scans

## Conclusions and Future Direction

OMEN is a useful data analysis GUI for use with slope measurement devices to do metrology on high precision mirrors. It allows for a variety of standard functions, and it is fast and effective at analysis. It is also relatively easy to modify if new functions are needed in the future. It can create good plots for publishing results, and allows for a degree of adjustment to achieve a desired aesthetic.

OMEN also can conduct new investigations into procedures to measure high curvature mirrors and explore the high angle regime of AC-LTP measurement. It allows for the use of partial scans to make a slow but accurate measurement of curved mirrors, which worked fairly well. The deviation between full measurements of the mirrors and stitched partial scans was small. This deviation also appeared to correlate with the measured slope as well, especially in the case of the spherical mirror, as shown in figure 6. When a calibration curve was applied, the full scan was significantly closer to the stitched data, as shown in figure 8. This implies that benefit would come of future comparison between full scans after applying a full calibration curve.

This software opens up a variety of possibilities for future research. It makes analysis of AC-LTP output simple and fast, which allows for more complicated or involved projects in the future. There is also room for improvement within the software itself. Many of the problems listed in Known Systematic Error Sources are theoretically solvable, but have not yet been addressed.

## Acknowledgements

I would like to thank the Lee Teng Undergraduate Fellowship for their generous sponsorship and for granting me the opportunity to take part in this research.

I would like to thank my mentor, Dr. Assoufid for welcoming me into his group, and being willing to meet me almost every day to discuss my project. I would also like to thank Jun Qian and Joe Sullivan for their work building, programming, and operating the equipment, and for helping me understand what the data was and how to work with it.

Finally I would like to thank the Optics Group in the X-ray Science Division of the Advanced Photon Source at Argonne National Laboratory, and Argonne National Laboratory itself for hosting me for the summer.

## References

- [1] Assoufid, Lahsen et al., Nucl. Instrum. Methods A 710, 31-36 (2013).
- [2] Qian, J., J. Sullivan, M. Erdmann, A. Khounary, and L. Assoufid, Nucl. Instrum. Methods A 710, 48-51 (2013).
- [3] Siewert, F. et al. AIP Conference Proceedings, Volume 705, p. 847-850
- [4] Elcomat 3000, Möller Wedel Optical, <http://www.moeller-wedel-optical.com>.
- [5] Shu, D., J. Qian, W. Liu, S. Kearney, J. Anton, J. Sullivan, and L. Assoufid, Proc. SPIE 9206, Advances in Metrology for X-Ray and EUV Optics V, 920601 (October 7, 2014); doi: 10.1117/12.2084726Proc. of SPIE Vol. 9206
- [6] Geckeler, Ralf D., and Andreas Just. "Distance-dependent Influences on Angle Metrology with Autocollimators in Deflectometry." Advances in X-Ray/EUV Optics and Components III (2008): n. pag. Web