XSD/BCDA APS-U End-station/Instrument Control System Strategy

The requirements

The APS Upgrade poses challenging requirements for control and data acquisition, but the requirements generally are not evident from x-ray beam or accelerator parameters, and they generally are not separable from instrument-design or even from operational strategies. Instead, we rely on beamline scientists, engineers, LDRD projects, functional descriptions, and engineering specifications of beamlines, instruments, and vendor products to elucidate actionable controls and data-acquisition requirements that can yield credible end-user data from upgrade instruments. Thus far, our requirements analysis is most complete for microscopes.

X-ray beam-coherence and beam-intensity figures, predigested by beamline scientists (etc.) yield minimum practical focal spot sizes in the range 5-20 nm, which determine controls requirements for positioning and scanning. It has become clear from experimental work that the strategy of controlling focal spots with feedback information from interferometers is likely to be insufficient for upgrade microscopes, in part because suitable interferometers are not accurate enough to direct that control. Periodic errors of around 6 to 30 nm in suitable interferometers not only affect the accuracy and precision of positioning, but also inject audio-frequency noise into scanning motions.

For fluorescence microscopes, the strategy that is evolving is to control focal spots as well as we can, given the available feedback information, and to record interferometers and other position sensors during data acquisition, so that acquired photons can be associated with sensor readings, and the pixel from which a photon was acquired can be calculated at data-analysis time. This strategy allows for correction of sensor signals, if their errors have been measured, and thus motivates the inclusion of calibration instrumentation into the design of upgrade microscopes.

Focal spot size and intensity have additional consequences for control and data acquisition. In the BNP-II engineering specification, the minimum scanning speed of a 10 nm focal spot is calculated (from radiation-damage requirements) to be 5 mm/s. Working from these figures, the control/data-acquisition system must generate data-acquisition triggers and record interferometers at 500 kHz, and the data-latency difference between interferometers and fluorescence detectors must be less than 1 microsecond, in order to assign acquired photons to the correct pixels. The current BNP strategy of generating triggers in a Delta Tau motion-controller servo loop cannot meet this requirement, fluorescence detectors currently in use cannot acquire spectra at this rate, and interferometers have tens of nanometers of noise when read at this rate.

But it's not all bad news. In addition to allowing for sensor correction, determining the focal spot location at data-analysis time allows for a greatly simplified control strategy: we don't need to put the focal spot exactly where we want it in real time, we just have to visit every

pixel. This replaces a stringent two-dimensional, real-time, precision and accuracy control requirement for both the fast and slow axes of a raster scan, with a relative (scan line to scan line) precision-only requirement for the slow axis only. The fast axis just has to move. We'd like it to move smoothly and accurately, but these are no longer requirements. (For clarity, fast axis motion refers to the component of focal spot motion relative to the sample that lies along the fast scanning axis. This is not a requirement on any one stage, but on all of them together.)

All of the upgrade microscopes will be designed to acquire tomography data, and this requires us accurately to measure the location of a rotating object. The plan for doing so involves characterizing a reference cylinder with (we hope) few nanometer accuracy, so it can be used in a microscope. Currently, this characterization effort is centered in the BCDA group, partly because we have developed the instrumentation and software needed to acquire figure-error data and distinguish it from unintended sensor motion and rotation-stage runout.

In addition to programmed and vibrational motion, upgrade microscopes will be vulnerable to sample motion from thermal expansion that may be insufficiently controlled, compensated, or measured. This seems likely — at least for the In-Situ-Nanoprobe — to motivate focusing an interferometer sensor directly on the edge of the sample (typically, a 200 micron thick Si_3N_4 membrane), or within a few millimeters of the sample. We have an optical design intended to do this and still permit samples to scan by many millimeters, but the design has not been tested.

Aside from microscope requirements, we know that both current and upgrade instruments need a flexible capability to synchronize data acquisition with the synchrotron-bunch clock. Current users have also demonstrated a need for a histogramming scaler that is significantly faster than the multichannel scaler currently in use at many APS beamlines.

The solutions

We have demonstrated a data-acquisition capability sufficient to meet most of the above requirements in measurements at the Bionanoprobe. The combination of an FPGA and a Linux processor, in the form of the Xilinx Zynq chip, is the heart of a hardware/software solution we have developed, called softGlueZynq. This system can generate data-acquisition triggers from the difference between sample and zone-plate interferometers, and can record six interferometers and two additional channels (time tag and dwell time), at up to 400 kHz. We also have demonstrated the use of this system to acquire and bin fluorescence spectra with an XIA XMAP detector in list mode, which can tag photons at that pixel rate. An overview talk[(2)] on this hardware is available online.

The control strategy is a shared effort by beamline-controls and mechanical engineers designing the microscopes. It is a work in progress, but some choices have been made. We will test Delta Tau, Advanced Control Systems, PI, National Instruments, and SmarAct positioners and controllers, to assess their suitability and our ability to achieve and measure the scanning

motions that beamline scientists require. Specific microscopes may also employ vendor-integrated positioner/controller solutions for some motions, so we may learn how to develop for those devices as well.

We have an FPGA design that addresses the requirement to synchronize data acquisition with the synchrotron-bunch clock by synchronizing all of the clocks in softGlueZynq with a 44 MHz signal, which is the rf clock divided by eight. Because of the way softGlueZynq is deployed, any copy of softGlueZynq can use this design simply by editing a text file to name the desired design, and rebooting the EPICS IOC. The new design is running at only 4idc, and is currently under evaluation.

To address user needs for a fast histogramming scaler, we have implemented one in softGlueZynq. This component is a two-input, 64-time-bin histogramming scaler that can channel advance at 50 MHz. The scaler is in use at 6idb for piezodiffraction experiments, and at 4idc for laser pump/probe experiments.

Risk mitigation
(softGlue hardware, software and support)

BCDA maintains a hardware/software capability intended to address users' needs for simple, easily modified digital electronics — for things like converting the output of one electronic device into input acceptable to another device. We call this "glue electronics," and this is why our software configurable electronics is called "softGlue." There are now around 30 copies of softGlue on APS beamlines, performing shutter/camera synchronization, encoder tracking and position-based triggering, generating digital waveforms, etc.

From the start it was clear that softGlue's goal, to create an easily deployable and customizable digital logic system, would require a fairly complex logic design running in an FPGA. If you look at an FPGA as simply a chip filled with configurable logic gates, then there are many choices with which to run your design. If you also consider the hardware and software environment in which the design needs to run, and the effort needed to support and maintain it, then the choices narrow down to a few sensible options.

The idea for softGlue or something like it existed for a long time. Several attempts were made at the APS to design and deploy user configurable digital logic. One attempt that pre-dated softGlue, called the "Generic Digital" board, was developed and used at the APS. But it wasn't really user configurable, there was no GUI that allowed run-time changes in the logic. Also, a user (engineer) would need to re-compile the logic in vendor specific design software if a change in the logic was needed. SoftGlue attempted to solve these problems by placing a highly re-configurable set of logic functions down in the FPGA, and tightly coupling these functions to EPICS.

The first attempt to design a prototype of softGlue came from a collaboration with AES/Controls engineers Marty Smith and Eric Norum. Marty had a commercial, off-the-shelf (COTS) board from Acromag with an Altera FPGA on it, working with EPICS. This board is an "Industry Pack" module, which sits on a carrier board in a VME slot. The infrastructure to support a board like this was already in place at almost all beamlines at the APS. With Marty and Eric's help, we were able to put our early design for softGlue down in the Acromag board's FPGA, and create EPICS support and GUI screens that were quickly tried, with success, at several beamlines. This solution was further developed, debugged and deployed over several years, using simple custom hardware allowing the connection of signals on common "Lemo" connectors. The rest of the hardware needed was COTS, commonly available (still is to this day), and easily deployable. As for software, of course we wanted to use EPICS. At the APS, this is the safest most sensible choice on many levels for many reasons. And, the EPICS I/O drivers for softGlue were written using Asyn, which brings a further layer of standardization and enhanced functionality.

The VME/IP softGlue solution was installed and used in many APS beamlines, and is still being used today. But over time it became clear that to maintain the ability to deploy softGlue solutions in an environment that was decreasing its reliance on VME, and to address user needs

for faster and more complex electronics, a new hardware strategy would be needed. We looked at several options. Full custom could be considered but has huge up-front and maintenance costs. National Instruments / LabView was considered, but this choice suffered performance limitations and big hardware costs. We ended up following the lead of other engineers in AES/Controls, XSD/DET, and detector development efforts at NSLS. These engineers were using "development boards," which are cheap, vendor-designed boards on which the FPGA is deployed. These boards are typically more modern, using the latest technology. A big technological benefit was gained by targeting the Xilinx "Zynq" chip for the next softGlue. The Zynq chip has an FPGA along with an ARM processor on a single silicon chip. This is a huge win for softGlue because an EPICS Linux soft ioc can run on the ARM processor, with tight, high-speed coupling to the logic running on the FPGA. The Avnet Zynq development board (MicroZed) was chosen and pursued for what has become "softGlueZynq" today. This solution is higher-speed and has larger capacity, leading to more and enhanced softGlue functionality.

There are nine copies of softGlueZynq running at APS beamlines today. The EPICS software is mostly an extension of what was already written for the original softGlue. The FPGA design, which was originally written in a proprietary Altera language, was translated to Verilog, which is an industry standard, open "Hardware Description Language" making it easier to extend, support, and maintain. Some custom hardware has been built to facilitate I/O connection via the Lemo connectors, and to put the board in a box and deploy as an enhanced 16-channel scaler.

Since we chose a hardware platform for softGlueZynq that other APS engineers were already familiar with and using in their designs, we not only benefitted from their experience in the design phase, but have greatly enhanced our ability to solve problems and move forward with this approach. One of the biggest concerns people have with this approach is, "What if they stop selling your board?" Although Avnet does have a 10 year lifecycle policy[(1)], MOST of the design (EPICS, Verilog) can simply be "ported" to the new FPGA/development-board if needed. This is really what you want to be able to do, since technology, chip complexity, capabilities and applications are always moving forward. Most of the technologies on which softGlueZynq relies are industry standard, and are available from FPGA vendors other than Xilinx. (The direct memory access component may be an exception.)

Reference:

1. http://zedboard.org/sites/default/files/documentations/SOM%20LifeCycle%202017.pdf
2. https://www.youtube.com/embed/zl4Vrh3_s70