Workflows at experimental facilities: Use cases from the Advanced Photon Source

Ian Foster, Tekin Bicer, Raj Kettimuthu, Michael Wilde, Justin Wozniak, Francesco de Carlo, Ben Blaiszik, Kyle Chard, Francesco de Carlo, Ray Osborn, and others

Argonne National Laboratory. Contact: foster@anl.gov

Increases in the volume and velocity of data from scientific experiments motivate significant changes in experimental protocols. Researchers often want quasi-instant feedback, so that one result can guide future observations. Such needs arise across DOE's Office of Science.

Coupled experiment-computation systems provide a rich source of opportunities and challenges for workflow technologies, as a single experimental session can operate at multiple time scales; engage both distributed systems and tightly coupled parallel computers; and require interactions with data archives and human collaborators: see Figure 1. Timely and reliable execution of the resulting workflows can enable important new experimental modalities (e.g., real-time studies of battery charge-discharge cycles). They can also be essential to efficient use of expensive facilities: without online feedback, a vastly expensive experimental session, booked months in advance, can be entirely wasted.

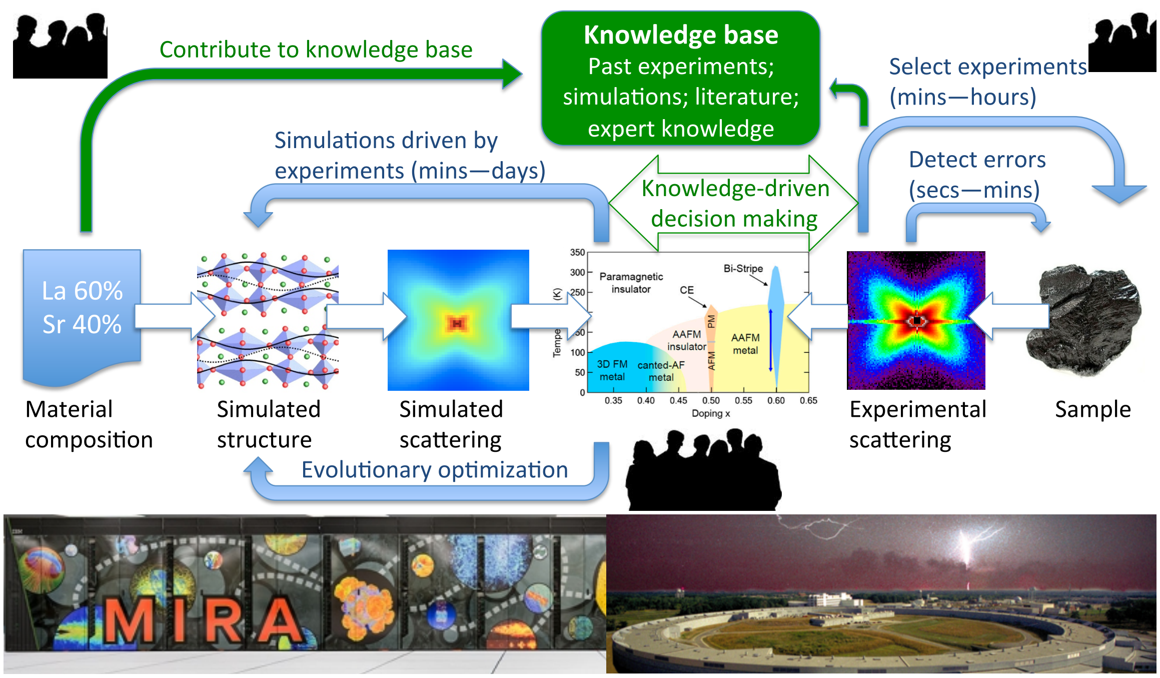

Figure 1: The analysis of defect structures in disordered materials can involve many different workflows with different time scales, computational needs, quality of service requirements, and resources. On the right, experiments are performed on Argonne's Advanced Photon Source (APS); on the left, computation is performed on the Argonne Leadership Computing Facility (ALCF)'s Mira supercomputer.

We use three application pilots from Argonne's Advanced Photon Source (APS) to illustrate some of the opportunities inherent in coupling experiment and computation. (The APS, a synchrotron light source, operates ~65 beamlines that support a wide range of experimental modalities, many scientific communities, and 1000s of users every year.)

Single-crystal diffuse scattering (Osborn, Wozniak, Wilde, et al.) The goal of this work is to understand defect structure in disordered materials. We have developed a range of workflows, illustrated in Figure 1, including rapid reconstruction during individual experiment (100s of cores), analysis of data from multiple experiments (1000s of cores), and evolutionary optimization for inverse modeling (100K+ BG/Q cores; Swift+OpenMP).

X-ray nano/microtomography (Bicer, Gursoy, Kettimuthu, De Carlo, et al.). Rapid image reconstruction enables new applications in bio, geo, and material science imaging, but requires large-scale on-demand computing. In one recent pilot, on-slice parallelization permitted reconstruction of a 360x2048x1024 dataset in ~1 minute, using 32K BG/Q cores, vs. many days on a typical cluster, enabling quasi-instant response.

Near-field high-energy X-ray diffraction microscopy (Almer, Sharma, et al.). This method is used to characterize microstructure in bulk materials. Reconstruction on 10K+ BG/Q cores (Swift + MPI-IO) gives results in ~10 minutes, vs. >5 hours on an O(100) core cluster or months if data taken home. This workflow was recently used to detect errors in experiment configuration that would have otherwise resulted in a total waste of beamtime.

These pilots also illustrate some key technical challenges that require substantial R&D if we are to make effective use of next-generation experimental facilities. For example:

Data-driven analysis: Experimental workflows are concerned above all with organizing, moving, analyzing, sharing, and tracking large quantities of data. In the examples above, we used Globus services to good effect to manage data flows and to track data products and provenance. But much more work is required if we are automate and optimize ever-more-complex workflows, e.g., to track data products generated by large numbers of experiments, process data generated by many such experiments, integrate experiment data with other knowledge sources, and introspect on behaviors of experiments, beamlines, and entire facilities to identify errors and plan future activities.

On-demand computation: Computation can link with experiments at multiple time scales: see Figure 2. The need to support a range of time scales, data volumes, and computational requirements has profound implications for computer systems architecture, scheduling methods, application configuration, and other factors. The computing systems best suited for modern experimental facilities are not yet clear. Given a select of computing resources, the choice of resource(s) for individual workflow components may depend on resource availability, scheduling policies, performance, application characteristics, and other factors.

Workflow implementation: Data- and compute-intensive near-real-time workflows pose simultaneous requirements for high performance (often close-to-bare-metal performance); programmer productivity (because researchers often vary experimental protocols), and system scalability (because many workflows may execute simultaneously).

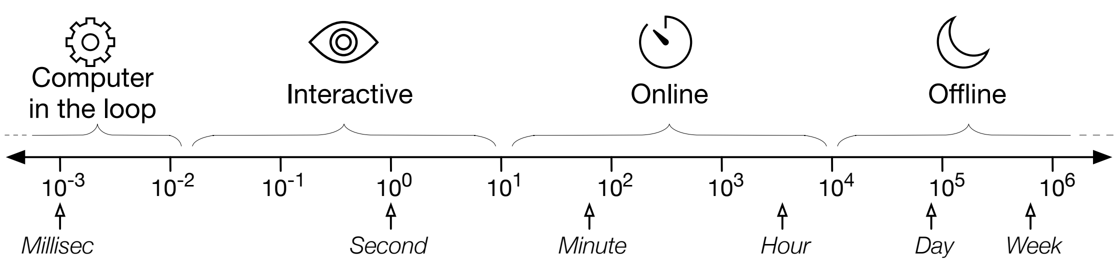

Figure 2: A rough categorization of the different time scales on which computation may be linked with experiments. (In practice, different categories and timescales may overlap.) *Computer-in-the-loop* activities may require timescales of a millisecond or less. For *interactive* activities, such as data subsetting to examine data features, delays of 0.1 to 10 seconds can be acceptable. *Online* computations, used for feedback during data acquisition, for example to detect errors in experiment setup, focus on an area of interest, or tune an experimental parameter, may range from minutes to an hour or more. *Offline* computations, performed periodically to say reprocess an entire data archive, may take hours or days.