DRAFT

Advancing Chemical Science: Future Networking Requirements

David A. Dixon, Pacific Northwest Laboratories
Larry A. Rahn, Sandia National Laboratories

Introduction

The chemistry community is extensive and incorporates a wide range of experimental, computational, and theoretical approaches into the study of chemical problems. Chemistry is one of the base sciences on which many applications are built. There is extensive use of basic chemical measurement techniques in a wide range of areas including atmospheric measurements, geochemical measurements, combustion and chemical process measurements, and cellular observations. Computational chemistry covers a wide range of areas ranging from accurate calculations on small molecules/processes such as heats of formation of radicals and electron scattering to intermediate accuracy calculations for the study of large molecules, separation systems and catalysts, and ultimately to molecular dynamics simulations of complex systems such as biomolecules and materials. There is also an extensive effort to discover of the details of chemical processes as they interact with the unsteady, and often, turbulent fluids that transport and mix the reacting species. Such studies are key to developing understanding that will enable predictive design of complex chemical processes such as combustion or chemical processing in industry. This research includes the production and mining of extensive databases from direct simulations of detailed reacting flow processes.

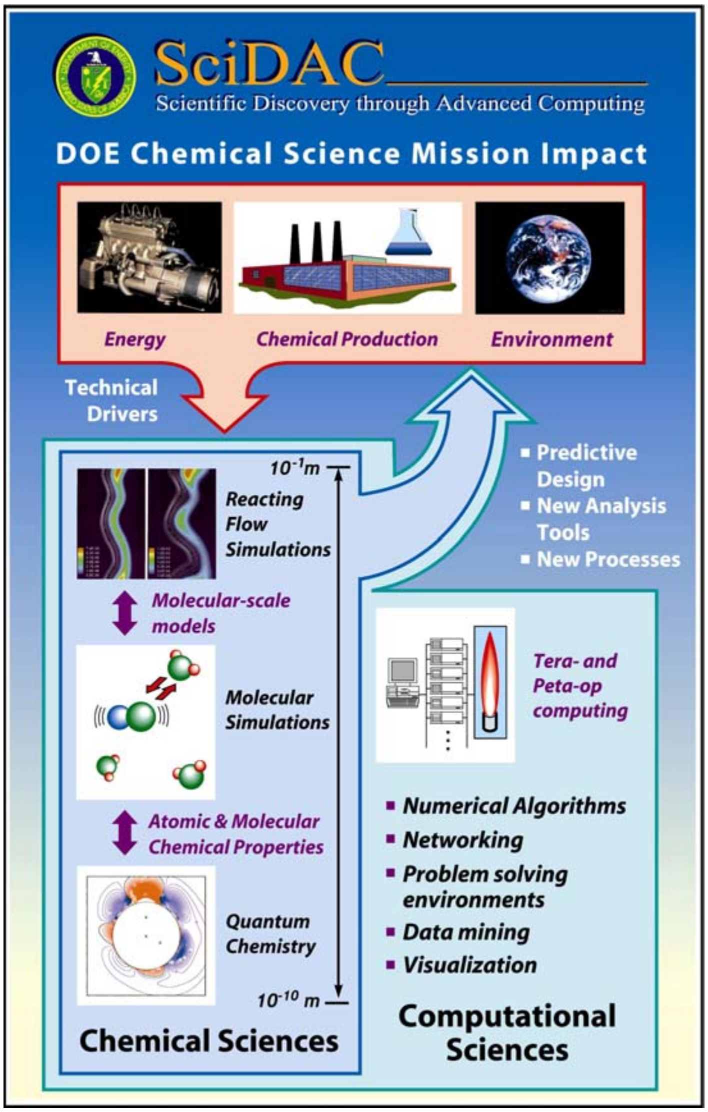

Figure 1. Chemical sciences, and combustion science in particular, require the integration of scientific knowledge over a large range of scales, disciplines, and geographic sites for impact on the DOE mission.

The area of chemical sciences is representative of many DOE programs in that it addresses complex multi-scale phenomena. The situation is similar in earth system studies, fusion research, high-energy physics, biology, and other areas of science. An understanding of environment and device scale phenomena requires more than simply applying one type of computation, with increased computing power, across scales. Different physical phenomena dominate system dynamics at these different scales, leading to a variety of models and experiments relevant in the different regimes. Information from one regime is used as input for the next, essentially "bootstrapping" from the atomistic to the device level. One of the major bottlenecks in such a multi-scale research enterprise is the passing of information from one level to the next in a consistent and validated manner.

The scientific process described above leads to a data- and model-centric view of the communications between sub-disciplines working at different scales. Data at one level is analyzed to develop a model that produces data used in turn by another, repeatedly across the range of scales and types of chemical information required. However, in this process more than just the raw data values need to be communicated. Confidence in a value's accuracy, its uncertainty, dependencies on other data, etc. must all be considered when using it in further computational and experimental research. In the direction of decreasing length and time scales, information about the sensitivity of models on particular data may place a premium on very accurate values for certain fundamental quantities. Enabling the rich bi-directional exchange of both data and metadata between scales is a critical issue in making progress.

Traditionally, this information flow has been accomplished through the research literature and, more recently, through databases of chemical values. The information is fragmented and discovery of new information in these sources remains a manual process. . Determining whether results presented in a paper depend on obsolete values from a different regime may require searching through several papers and databases. These factors make communication difficult and time consuming and increase the likelihood of redundant and irrelevant research.

To overcome current barriers to collaboration and knowledge transfer among researchers working at different scales, a number of enhancements must be made to the information technology infrastructure of the community:

- A collaboration infrastructure is required to enable real-time and asynchronous collaborative development of data and publication standards, formation and communication of inter-scale scientific collaborations, geographically distributed disciplinary collaboration, and project management.
- Advanced features of network middleware are needed to enable management of metadata, user-friendly work flow for web-enabled applications, high levels of security especially with respect to the integrity of the data with minimal barriers to new users, customizable notification, and web publication services.
- Repositories are required to store chemical sciences data and metadata in a way that preserves data integrity and enables web access to data and information across scales and disciplines.
- Tools now used to generate and analyze data at each scale must be modified to enable generation and storage of the required metadata in a format that allows interoperable work flow with other tools and web-based functions, and must be made available for use by geographically distributed collaborators.
- New tools are required to search and query metadata, and to retrieve data across all scales, disciplines, and locations.

The complexities of managing information within such an infrastructure are daunting and the creation, communication and use of the additional information could quickly become unwieldy. However, recent technological advances network protocalls (such as DAV) and the development of the extensible markup language (XML) for defining machine and human readable metadata based on standard schema, have significantly reduced the barriers to creating such a comprehensive informatics environment.

Networking needs

Rapid advances in computational hardware and software along with innovative experimental techniques are revolutionizing the rate at which chemical science research can produce the new information necessary to advance chemical science, for example, combustion technology, biochemistry, environmental chemistry, catalysis and chemical process technology, etc., straining the traditional methods of communication through peer-reviewed literature and static databases. In addition high-throughput technologies promise to revolutionize the science that we do as well as to create huge amounts of data.

Computational Science

Electronic structure – This area is a huge consumer of computational resources in terms of cpu cycles, memory, and I/O but will not require large networking resources for moving data around. For computing across the Grid, the networking needs change significantly in terms of latency and bandwidth.

Molecular dynamics – Based on the new EMSL/MSCF HP supercomputer, a 9+ TFlop system, we estimate that we will be generating ~20 Gbytes/hr leading to 400 to 500 Gbytes/day/~10 Tflop machine.

Simulations of reacting flow systems – Direct numerical simulations (DNS) carried out using currently available computational resources (1 to 3 TeraFLOPS) are limited to two dimensions and either simple processes, or simple chemistry. Recent simulations of the time dependence of auto ignition in hydrogen-air mixtures produced 0.75 Tbytes of data from a 1M node-hour simulation on an IBM SP3. With computational capabilities of 100 TeraFLOPS or more envisioned in the coming decade for DOE science, such simulations could provide insights into complex autoignition phenomena in three dimensions with hydrocarbon chemistry, producing Petabyte-scale data sets.

In principle, a much broader array of reacting flow science will be enabled with the advent of such computing resources. SciDAC supernova researchers are already interested in the DNS tools developed for combustion. Other areas will include chemical processing, multi-phase flows for manufacturing unique particles and coatings, and predictive tools for fire-safety, and possibly inertial confinement fusion.

Example Scenario:

Remote Interactive Visualization and Data Mining of Combustion DNS

The advancement of the DOE mission for efficient, low-impact energy sources and utilization relies upon continued significant advances in fundamental chemical sciences and the effective use of the knowledge accompanying these advances across a broad range of disciplines and scales. This challenge is exemplified in the development of predictive computational models for realistic combustion devices. Combustion modeling requires the integration of computational physical and chemical models that span space and time scales from atomistic processes to those of the physical combustion device itself as illustrated in Figure 1.

Combustion systems involve three-dimensional, time-dependent, chemically reacting turbulent flows that may include multiphase effects with liquid droplets and solid particles in complex physical configurations. Against this fluid-dynamical backdrop, chemical reactions occur that determine the energy production in the system, as well as the emissions that are produced. For complex fuels, the chemistry involves hundreds to thousands of chemical species participating in thousands of reactions. These chemical reactions occur in an environment that is defined by both thermal conduction and radiation. Reaction rates as a function of temperature and pressure are determined experimentally and by a number of methods using data from quantum mechanical computations. The collaborative creation, discovery, and exchange of information across all of these scales and disciplines are required to meet DOE's mission requirements.

As discussed above, this research includes the production and mining of extensive data bases from direct simulations of detailed reacting flow processes. The problem definition, code implementation, and mining of these datasets will necessarily be carried out by a collaborative (and distributed) consortium of researchers. Moreover, this research leads to reduced models that must be validated, for example, Large Eddy Simulations (LES) that can describe geometrical and multi-phase effects along with reduced chemical models. While these reduced model simulations typically will use fewer FLOPS per simulation, more simulations are required to validate models in environments that will enable predictive design codes to be developed.

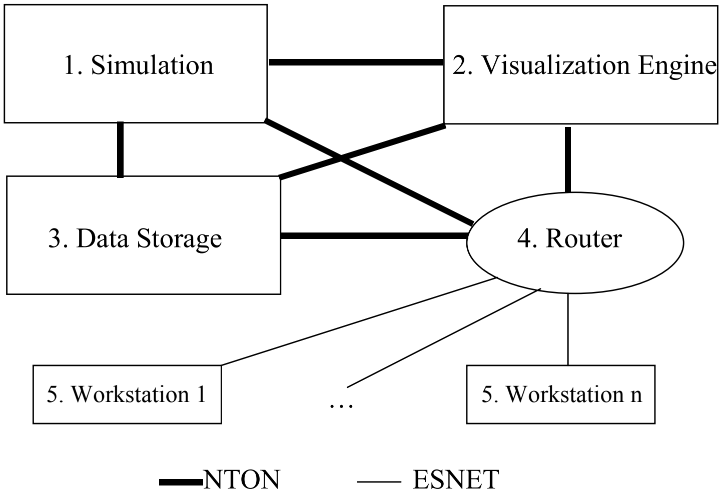

Figure 2. Schematic of a possible coupling of resources via high-speed networks allowing utilization of distributed resources for advanced data visualization and analysis.

Collaborative data mining and analysis of large DNS datasets will require resources well beyond the current practice of using advanced workstations and FTP shipping of data. The factors driving this change are dataset sizes that require unique storage resources, the increased collaborative nature of the research projects, and the great increase in the complexity of the information that can be computed for mining and analysis. Thus, the post-computation evaluation of the data will require unique computing resources for processes such as visualization and feature detection and mining software. Such analysis tools will be remotely operated and benefit from real-time interactions by the research scientist. process.

An example of a computational infrastructure for our visualization needs is shown in Figure 2. Since it is composed of scarce, geographically distributed components, a fast network is needed to unite these resources. The constraints imposed on the network are low latency, high bandwidth, some form of quality-of-service, and reliable multicast fro broadcast to the collaborating scientists. In addition, it would be desirable to provision a fail-over path, which could be used transparently to the user(s) in case of failure along the primary path. The example here was drawn with the National Transparent Optical Network (NTON) in mind to allow coupling of resources in DOE national laboratories in the San Francisco Bay area.

As an example, 244GB are generated from an H2/Air autoignition DNS simulation in 2D (8 species) and 1.3 TB for n-heptane autoignition in 2D (44 species). One H2/air ignition run takes about 50000 to 100000 cpu node hours on an IBM SP system at NERSC depending upon the length of time we need to run (1-2 milliseconds for typical high pressure ignition delay times.) Running on 256 processors, it takes about 2-4 weeks turnaround per run (including time waiting in batch queues.)

Experimental Science

Mass spectroscopy for high throughput proteomics – current estimate is ~100 Gbytes/day/mass spec based on a 50% improvement of the current prototype data collection system. For 20 to 50 mass specs for the entire complex, this leads to 2 to 5 TBytes/day.

Imaging – When running at full speed, a FRET analysis of a cell will be generating a Megapixel/msec of data. A Megapixel is ~24 x10^6 bits leading to 3 x 10^9 bytes for a 1000 step for a 1 sec trajectory. For 50 runs a day, this will generate ~ 150 Gbyte/spectrometer. We are building 6 spectrometers currently leading to ~1 TB at 1 laboratory. If this is spread across multiple labs, this number could easily reach 10 to 20 TB/day.

Examples:

Real-time Dual-Wavelength Confocal Microscopy. This microscope employes a high-speed AOTF system in conjunction with a Nipkow disk confocal imaging and dual Gen III intensified CCD cameras. This allows ratiometric imaging of life cells at 30 fps. Applications, such as FRET in live cells and dynamic colocalization of multiple fluorescent probes can be conducted on timescales of minutes to hours.

Each of the systems below will generate between 0.5 and 1.5 Gbytes of image data/day.

CARS (coherent anti-stokes Raman spectroscopy)/Two-Photon Confocal Combined Microscope. This imaging modality can visualize molecules based on their vibrational properties and simultaneously by their fluorescent properties. CARS is useful for visualizing selected molecular species, such as lipids and deuterated compounds. Spatial resolution is 0.3 microns. The present temporal resolution is about 1 min. per image, but will improve with new laser technology.

MRM/OM is a combined magnetic resonance imaging/spectrograph system that has a simultaneous confocal microscopy capability. The magnetic resonance provides unique physical and chemical information on a distance scale as small as 10 microns, with no limitations on sample opacity, and is both non-destructive and non-invasive. The rather slow MRM imaging time (10s of minutes, depending on resolution) is supplemented with optical fluorescence images. Confocal images are recorded in seconds with micron resolution, simultaneously with MRM imaging. Acoustic microscopy methods to supplement MRM are under development.

AFM-Enhanced Fluorescence Imaging Microscopy. A novel nanoscale characterization combined microscope that uses the Optical/AFM approach and a high-sensitivity far-field microscope to provide unperturbed measurements of reaction rates at sites that have been characterized by atomic force microscopy (AFM) imaging. Work is being done to extend this approach by using the surface enhanced Raman scattering (SERS) method in addition to the fluorescence imaging method to facilitate the application to *both fluorescent and non-fluorescent* substrate species. The strong field enhancement at a sharp AFM metal tip coupled to an external optical field should allow the observation of site specific SERS while providing contrast from fluorescence imaging.

Single-Molecule Spectroscopy and Imaging Microscopy. This instrument is capable of simultaneous structural and spectroscopic analyses of single biomolecular complexes and their reaction (interaction and association) rates, providing insight into the relationship between structure and function of cell signaling proteins and enzymes. The information collected from these approaches is typically lost or hidden in the measurements using current conventional ensemble-averaged methodologies.

Single-Molecule Patch-Clamp/Optical Confocal Imaging Microscopy. This is a unique technology by combining a confocal scanning linear/non-linear optical microscope with state-of-the-art patch clamp technologies. This instrumentation significantly enhanced the diagnostic and investigative capabilities of both methods in the characterization of ion channel/receptor dynamics and mechanism in a living cell. Understanding the correlated conformational dynamics of a single ion channel/receptor will be a significant step in deciphering the molecular mechanisms of ion channels functions and the process of ligand-receptor and receptor-lipid interactions for individual molecules in membranes from living cells.

Nuclear magnetic resonance – A protein structure experiment generates 1 to 1.5 Gbyte/run and one can do 2 to 3 runs/day leading to 3 to 5 Gbyte/spectrometer. Many spectrometers could easily lead to 100's of Gbytes/day.

Cryo electron microscopy(EM)/EM crystallography of single molecules – digitized electron micrograph (photographic film), consists of approximately 6000X6000 pixels, usually digitized in 12 bits and go to 1000 images/day/spectrometer. This is 54 Gbytes/spectrometer.

Synchrotron data – Produces large data sets. Experimentalists have been hampered by the lack of local computing facilities at the synchrotron source which prevents analysis of data and can hold up experiments.

Real time, remote control of experiments – This requires significant networking resources that are always available to the remote user. The experiment transmits data back to the user which enables the user to make decisions about how to control the experiment. The amount of data transmitted *by* the user is usually small but a high integrity network is needed with low latency. The amount of data transmitted *to* the user may be much larger so bandwidth and latency are important for real time control.

General Scenarios

There are two scenarios:
1. Move data around to everyone from a small number of very high performance computers and store data from everyone at a small number of very large data storage centers

2. Place many modest size computers working locally on the data with large amounts of local storage. Computers and storage are co-resident with the data Only summary data is passed around on the network.

Scenario 1 requires high network speeds with reduced manpower needs. Scenario 2 requires significant manpower needs to manage all of the computer systems.

Issues

Challenging networking and other information technology research needed to enable distribution of data, analysis, and collaboration include:

- Multicast service delivered to multiple remote centers with diverse firewall filters
- Network error rate and robustness control *without* impacting the experiment's data acquisition system
- Massive applications software – e.g., 1 million lines of code in NWChem
- Interfaces of databases with the network and storage
- Technology improvements including:
    - Computing technologies
    - Computer science
    - Applied mathematics
    - Software development
    - Networking
    - Computing system-to-network interfaces
    - Fiber technologies

- Data storage
- Data management

The international chemistry research community is increasingly using Grid technologies, an integrated suite of services. The Grid is a set of middleware tools and capabilities that enable seamless end user access to applications, data storage, and compute resources to support high-end modeling. The Globus project (http://www.globus.org/) is one state-of-the-art example of Grid development. Grid middleware faces many hard computer science problems. Vertical integration of existing components to provide Grid services to demanding, well-defined communities is essential to progress on Grid architecture and technologies.

We now present some of the issues as questions that will need to be addressed.

How will the different sets of data be processed across the networks?

For example, for combustion modeling, we envision running the data at large terascale computing centers (NERSC, Oakridge, PNNL), archiving the raw complete data at these centers. On-the-fly data-mining /diagnostiscs visualization will be performed to verify the viability of the runs, with minimal graphics streamed to remote sites (lines, polygons, etc.) Another use of on-the-fly data-mining is to tag and steer the run to redundantly save portions of the data for future mining/analysis of a specified kind (e.g. flamelet analysis, statistical point-wise analysis, etc.) These more limited archives of the raw data can then be shipped to remote sites for further detailed analysis. Alternatively, if bandwidths are big enough, we could perform the analysis on computers where the data was created, and only stream graphics across the network (compressed, non-lossy data). In the interim, we need to have both capabilities. We feel that data-mining, particularly for large complex 3D time-dependent data sets, is essential if information is to be gleaned from the runs efficiently. How the data for reacting flows is mined determines what and how much information needs to be shipped around. Finally, it is impractical/difficult to move large amounts of DNS data around via protocols like FTP. For example, the volume of data for one run currently on NERSC archives (2D, hydrogen/air ignition, 8 reactive species, 1.8 million grid points, each restart file is 200MB, 200 restart files to track ignition evolution, total of 40 Gbytes of raw data per realization. To simulate this in 3D requires minimally 50 million grid points, or 1.11 TB of raw data per realization. Further, a parameter study would include 3 to 5 comparable simulations, varying various flow/chemical parameters. Therefore, a complete ignition study in 3D would generate ~5 TB of raw data. Even if one could move this amount of data to the home institution, it is not clear that there would be adequate storage /computing capability to process/visualize the data.

Additional questions include:

> *How often does one really want to move the data (i.e. will go to a repository once or will it be moved around more often)?*
> *Will only a subset of the data be moved (needing data mining technology)?*
> *Will the data only be visualized remotely?*

These issues are related to tied Grid networking needs with the issue being how much data storage and networking needs need to be close to the machines that are producing the results.

To get the information that network planners need, we need to figure out where the data goes. For example, people do molecular dynamics calculations at EMSL (and presumably NERSC) or a DNS simulation at NERSC. *What fraction of those results are moved to the user's home institution for analysis?*

For a fluid dynamics DNS combustion simulation, 100% of the data generated at NERSC is shipped back to SNL for analysis. For DNS data, there is a group of universities and National Laboratories within the

U.S. (U. Michigan, U. Maryland, North Carolina State U., SNL, Pittsburgh Supercomputing Center) and abroad (Cambridge U.) where it would be desirable to have data moved to for analysis/vis. The U.S. institutions are participants in a Chemical Sciences SciDAC project, URL: //scidac.psc.edu. This is likely to be typical of most large projects in the future, large, distributed groups sharing large amounts of data in a global fashion.

In many respects, solving DOE's networking problems cannot be done without also solving the computing capacity problems. There simply aren't enough computers to go around now, and not enough in the pipeline to handle the work envisioned by these plans. Hence we have trouble deciding whether to analyze data where it's generated (on machines that don't exist), or try to move it (perhaps compressed) to some other place for analysis (where there are likely not enough big computers either).

In the long run, the number of locations with big chemistry data sources will always be larger than the number of places with big computers, especially as physical chemistry is used in biological analysis (e.g., NMR, mass spec). Hence we are always going to need high performance networks to bring big data, computers and people together. Nevertheless, we probably don not do enough to reduce data size before we transmit it, which will require new algorithms broadly supported.