Macromolecular Crystallography: Networking Requirements and Usage Scenarios

T.N.Earnest, C.W.Cork, G.McDermott, J.R.Taylor
Lawrence Berkeley National Laboratory

1. Introduction

Macromolecular crystallography is an experimental technique that is used to solve structures of large biological molecules (such as proteins) and complexes of these molecules. The current state-of-the-art implementation of this technique requires the use of a source of very intense, tunable, x-rays which are only produced at large synchrotron radiation facilities. There are approximately 18 synchrotron radiation facilities in operation or under construction worldwide: 6 in the United States, 6 in Europe, 2 in Japan, and one each in Brazil, Canada, China, and Taiwan. In the United States alone, there are 36 crystallography stations which are distributed among the synchrotron facilities and dedicated to macromolecular crystallography [1]. Operating costs for each of these crystallography stations is estimated to be approximately $2K - $3K/hr. These stations are also responding to an increasing demand to solve new structures arising from both the national genomics research programs and from commercial drug development R&D. The high operating cost of these facilities, coupled with the heavy demand for their use, has led to an emphasis on increased productivity and data quality which will need to be accompanied by increased network performance and functionality.

Data acquisition for macromolecular crystallography typically involves repeated exposure and readout of imaging detectors while rotating the sample in the x-ray beam. Current systems can produce 10 - 100 MByte images at a maximum rate of 0.5 image/sec. Future systems which should become available within the next 5 - 10 years are expected to reach peak data rates of 50 - 500 MByte/sec. Average data rates are somewhat less due to issues with sample handling and instrument setup; however, new developments in automated sample handling and parallel data processing are promising average data rates which approach 50% of peak rates. A typical dataset consists of 500 - 1000 images which are usually stored uncompressed for fast online analysis (requiring 5 - 100 GByte disk storage).

2. Networking Requirements

As illustrated in Figure 1 below, the data acquisition process involves several interactive online components, data archiving and storage components, and a compute-intensive offline component. Each component has associated networking requirements.

Online process control and online data analysis are real-time, interactive, activities which monitor and coordinate data collection. They require high-bandwidth access to images as they are acquired from the detector. Online data analysis is currently limited primarily to sample quality assurance and to data collection strategy. There is increasing emphasis on the expanding this role to include improved crystal scoring methods and real-time data processing to monitor sample degradation and data quality. Online access to the image datasets is collocated and could make good use of intelligent caching schemes. Datasets from previously exposed samples are not required during online processing.

Offline storage is usually remote from the data collection facility. Local high-performance storage is currently too expensive for long-term storage and data security issues also sometimes warrant the removal of data from shared facilities. Data archiving is currently performed using either FTP/SCP network transfers with conventional data compression, and/or by backup to tape or removable disk with personal transport to the remote storage facility. Removable disk has recently become especially attractive because of its high performance and relatively low cost. End-to-end data transfer rates have been the main limitation in network transfer mechanisms (further discussion below).

In the following sections we will briefly discuss possible networking scenarios for the online and offline data acquisition components.

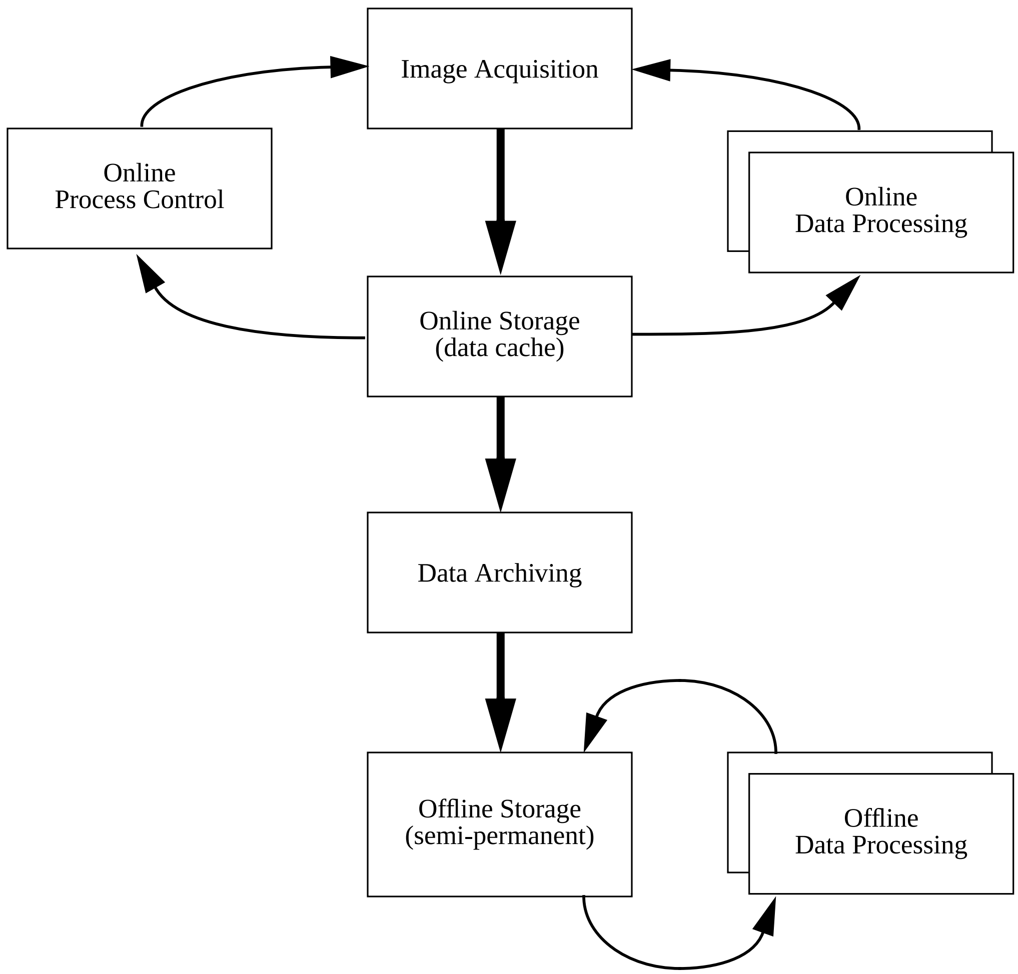

FIGURE 1. A simplified view of the macromolecular crystallography data handling process.

2.1 Online Control and Data Processing

High performance networking can play several roles in online control and data processing. Bob Sweet at BNLs National Synchrotron Light Source has outlined several approaches to remote, networked, collaboratory operation [2]:

1. **Remote Observer.** This involves remote monitoring of data acquisition progress and snapshots of the online data processing status. Data acquisition is coordinated exclusively by local personnel. A limited set of agent programs will provide access and presentation services to the remote client. Network bandwidth and QOS requirements are minimal. This scenario of operation would benefit from middleware utilities which facilitate agent discovery, authentication, and connection.
2. **Collaboratory.** This is a hybrid remote/local control scenario where local operators communicate with remote collaborators to coordinate data acquisition. This scenario will require telepresence and network conferencing software, in addition to the facilities required for the ‘Remote Observer’ approach described above. Network bandwidth requirements are approximately the same as for ‘Remote Observer’, but QOS requirements are increased to guarantee low-latency communications.
3. **Remote Control.** This scenario extends some or all of the online process control function to the remote site. Several variations of this approach are being developed, ranging from remote ‘recipe’ prescription to total remote instrument control. Most of the telepresence facilities required for the ‘Collaboratory’ approach will also be required here for interaction with local support personnel. Network bandwidth requirements are

somewhat greater than for the ‘Remote Observer’ as update rates will need to be increased to improve interactive feedback. QOS requirements are also increased to include additional security and transaction features.

For all of the above scenarios, high bandwidth is not so important as QOS and management services. Especially valuable would be middleware tools which facilitate the initiation, configuration, and monitoring of these services.

2.2 Offline Storage and Data Archiving

As mentioned above, offline storage is generally remote from the local data collection station. The datasets are occasionally stored at shared processing facilities (such as Brookhaven Lab’s ASDP facility [3]). However, the datasets are most often transferred to private institutional storage. This requirement places a large burden on the data archiving process which transfers the data between online and offline storage units. Current requirements for average data transfer rate are 1 - 25 MByte/s per station; it is expected that in 5 - 10 years this will increase by an order of magnitude to 10 - 250 MByte/s per station. This is further exacerbated by the fact that most research facilities have from 4 - 8 stations; this places a future requirement of 40 - 2000 MByte/s per facility. Advanced data compression schemes might be able to reduce these figures by a factor of 5 - 10.

Within this data rate performance envelope there are several scenarios that might be implemented in a high performance networking environment:

1. **Data replication.** This involves variations on data copying between local and remote storage sites. It could be as simple as an automated remote backup tool to a full-fledged data mirroring service. Middleware services which support location, authentication, and data replication will be essential for the success of this scenario.
2. **Virtual storage.** This involves variations of network filesystems. In this scheme online and offline data storage are merged into a single virtual storage facility. If network performance is sufficient, it could involve a scheme as simple as NFS or AFS; however, the heavy random access use by local online data processing applications probably dictates some sort of transparent local cache mechanism. Middleware should implement this virtual storage mechanism and provide simple management tools and API’s for storage location, authentication, and attachment.

High performance networking will only be used for data archiving if it is competitive with other data archiving techniques. The following local data archiving techniques are envisioned as being competitors:

- **Removable disk.** Firewire and/or SCSI removable disk. Maximum write speed ~24 MByte/s.
- **Removable optical/tape.** DVD/R and/or XDLT/Tape. Maximum write speed ~1 - 15 MByte/s.

These systems are competitive with current network performance and are often the preferred data archiving method. The new high performance network service will need to overtake these technologies (and their future extensions).

3. Summary

In addition to increased raw network bandwidth, the next generation high performance networking infrastructure will need to provide tools and services which facilitate object discovery, security, and reliability. These tools are needed both for low-latency applications such as remote control, as well as high throughput data transfer applications such as data replication or virtual storage systems.

References

[1] BIOSYNC - Structural Biology Synchrotron Users Organization (http://biosync.sdsu.edu).

[2] R.M.Sweet, M.Becker, J.M.Skinner, “Collaboratory Tools for Macromolecular Crystallography at the NSLS”, private communication.

[3] NSLS Automated Structure Determination Platform (http://asdp.bnl.gov).