Atlanta Area BioRing Project

Technical Overview Working Draft

March 4, 2003

Russell Clark rjc@cc.gatech.edu
Matt Sanders msanders@cc.gatech.ddu

Document Status

This is a working draft describing the ongoing efforts to develop the Atlanta Area BioRing. This document is intended to collect ideas regarding the BioRing project goals as well as the possibilities for deployment.

Project Summary

The BioRing project is about leveraging the availability of high speed, highly available networking services connecting the core Bio-X researchers in the state of Georgia. The BioRing will support researchers in Biology, BioMedicine, BioInformatics, etc. from several key institutions.

- Georgia Tech
- Emory University
- Morehouse School of Medicine
- Clark Atlanta University
- Georgia State
- University of Georgia
- Medical College of Georgia, Athens

The initial collaboration study is with researchers at Georgia Tech and Emory. We are actively pursuing involvement of researchers at the other institutions.

Applications - The Big Four

An important aspect of the BioRing project is the identification and in some cases development of the enabling applicatons that will lead to real advantage for the supported Bio-X researchers. These applications can be classified in the following four key areas.

1. Access To Large Data Sets - e.g. genomics

   Bio-medical researchers rely on a large set of databases around the world. These databases include such resources as genome and protein maps, imaging libraries and pharmeceutical interactions. The volume and breadth of the data is increasing many-fold. Scientists need to be able to access this data quickly and repeatedly to do their work.

However, simply connecting fast networks to access the data will not be sufficient. The sheer volume of data as well as the variations in type, format and quality require that increasingly powerful data management tools be available. Automated data mining and management is necessary to filter out the information which is not currently relevant to the immediate task. It is also critical to automate the discovery of related data without requiring that scientists manually read each new article or analyze each new dataset.

2. Access To Computing Horsepower - e.g. Oak Ridge Supercomputers

   Much scientific discovery today is done using computer simulations to model the ideas that are either too expensive or not possible to perform in a real laboratory. Some applications to consider are: Genomic sequence alignment studies, Simulations of drug interactions, Reconstruction and modeling of evolutionary trees. Visualizations of genomic data.

   Reference grid-style CPU sharing, etc. as other examples here

3. Remote Instrument Access - Imaging, experimental systems

   Much modern scientific research requires access to sophisticated scientific instrumentation. Imaging systems such as MRI and Electron Microscopes are expensive and often inaccessible to scientists outside of the larger research labs. With the introduction of high capacity networking such as the Atlanta Bio-ring, scientists can now access these tools remotely, control their operation and receive the resulting images or reports without physically travelling to the instrument location.

   Another powerful example of this benefit is in the area of medical diagnostic equipment. A physician or medical researcher can view and even remotely control the diagnostic imaging equipment with a patient at another site. This allows physicians to easily call in outside, expert consultation on a patient's diagnoses and course of treatment.

4. Collaboration Tools

   A significant area of interest in high-speed networking is in the support of scientific collaboration. This area has many different aspects.

   Supporting collaboration across different distances: same campus, across town, around the world

   The BioRing is not focused on providing yet another video-conferencing tool for remote meetings. While video-conferencing applications may be supported, they are seen as one small part of a larger collaborative framework.

   Researchers need to be able to interact with each other while using the tools of their trade. For instance,

   - manipulating a genome model and highlighting a particular aspect for others to see.
   - Several researchers, viewing a MRI image and all able to point to aspects of the image and change viewing angle, scaling etc.
   - Bio-chemists conducting a simulation experiment, all participating in the parameter decisions and viewing the results in real-time.

Network Technology

This section will provide a bit of background on the current state of network technologies to support this project. We are proposing to develop a metro area fiber ring that operates at 1 and 10 Gigabit rates. We will try to identify the feasibility of obtaining NIC's for switches, routers, and endhosts that can operate at these speeds. For the near term, we do not expect to see very many end-systems that can utilize a sustained 10 Gbps data stream.

Current/Near term hardware offerings:

- Cisco routing gear
- Movaz optical switching gear
- Ciena optical switching gear

Possible Levels of Deployment

Figure 1 presents four possible deployment scenarios for the BioRing research network.

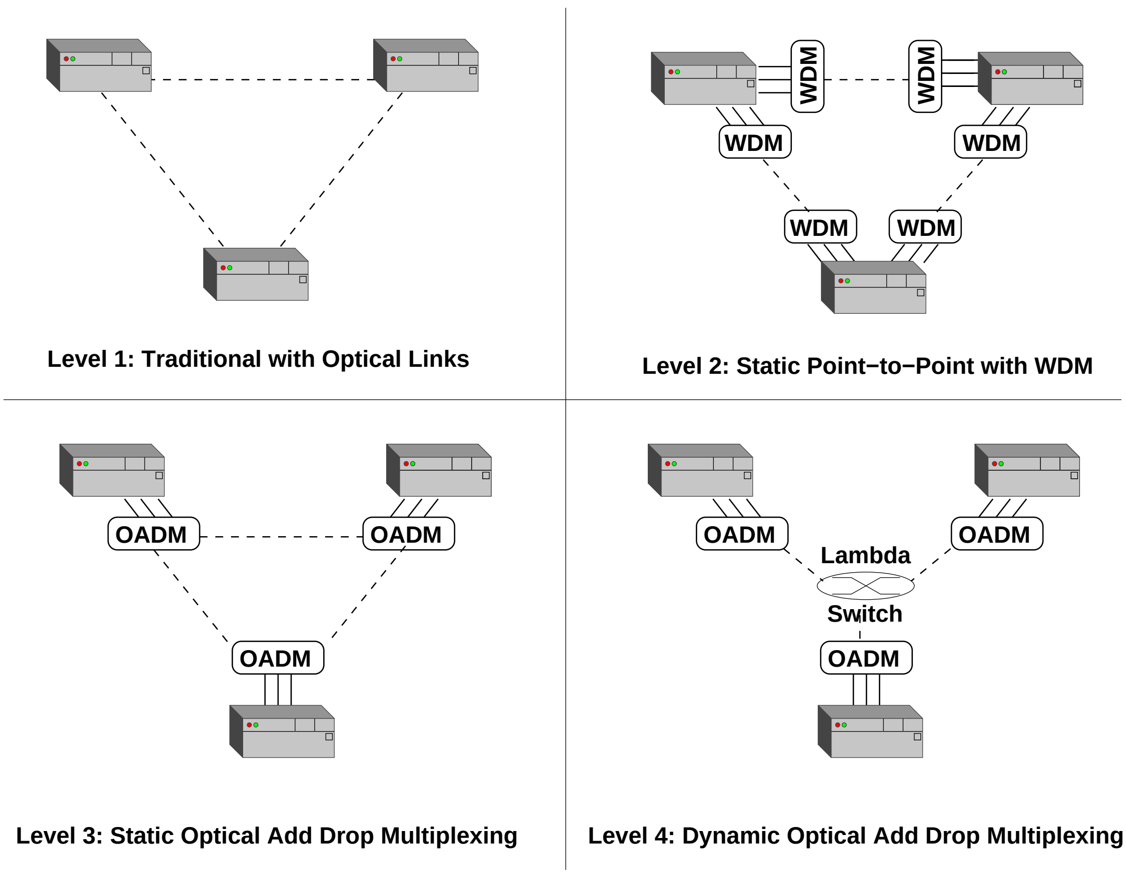

Figure 1: Optical Deployment Scenarios

- Level 1 Network Routed Architecture Using Single Channel Optical Links

  In perhaps the simplest deployment scenario, the Bio-Ring fiber will be used to provide high speed interconnect links between router ports at each of the participating sites. The maximum capacity (e.g. 10 Gbps) is shared by all traffic traversing the link. All traffic forwarding is performed at the IP layer.

This approach is easy to manage and deploy but is a fairly uninteresting architecture. Aside from the limited IP QOS services, there is little that can be done to isolate traffic and reserve bandwidth for real-time needs such as instrument control or collaboration support. The resulting network is essentially a fast Internet, dedicated to the BioRing project.

The primary advantage of this approach is that it can be deployed today with standard hardware and network management that is well understood.

- Level 2 Static WDM on Interconnect Links

  As in Level 1, the Bio-Ring fiber will be used to provide high speed interconnect links between router ports at each of the participating sites. However, the fiber links will be subdivided using WDM and allocated for specific usages.

  This allows for partitioning of the link capacity for different traffic priorities. The partitioning is static and must be pre-determined for a link. This greatly limits the flexibility of the fiber infrastructure for supporting specific, dedicated services.

  Additional hardware is required at the router/fiber interfaces to access the separate wavelengths. This raises the price of the deployment over Level 1 but provides some means of prioritizing traffic on the physical links.

- Level 3 Static Optical Add Drop Multiplexing (OADM)

  In Levels 1 and 2, all transit traffic (going through a site) is routed at the IP level. In the Level 3 scenario, the ring consists of an all optical pathway between all sites. This reduces the latency for transit traffic across the ring by avoiding the Optical-Electronic-Optical conversion. However, the wavelength allocation is static. Also, the exchange of traffic from the ring to/from the end stations is via routers at each site.

  This approach involves additional hardware and configuration complexity. The Lambdas (wavelength circuits) must be statically allocated. Traffic from point A to point B on the ring must be sent on a Lambda that terminates at point B. This static allocation requires significant lead time and money for new interfaces in order to add or re-route Lambdas in the network.

- Level 4 All Optical Switching Architecture

  Level 4 adds optical switching to the all optical network path. This overcomes the static configuration limits of Level 3. Also, network layer routing is no longer required. Layer 2 routing is no longer performed as the architecture is completely optical with a Lambda switch at the core. Light paths are configured from the edge of the ring to the Lambda switch. The switch can dynamically alter the direction of a light path and thus deliver packets to different destinations without pre-configuring the paths. In this scenario, allocation and reallocation of network capacity could be done on the order of minutes rather than hours or days.

Another aspect of deployment is how far out to the end user the high capacity fiber network will be delivered. The above discussion centers around the devices which form the ring between campuses. One can think of each router representing a single campus presence on the network. It is also important to look at the possible deployment of optical network directly to a laboratory, computer cluster, or a single high end instrument.

Timeline

We are currently developing a timeline for the BioRing project. The initial fiber plant connecting GT to Emory using AGL fiber will be available for Summer of 2003.