

Title: **Design and Analysis of a Dynamic DWDM-based Multi-Terabit/sec Packet Switching Fabric**
PI: Itamar Elhanany, Assistant Professor
PI Institution: Department of Electrical & Computer Engineering
University of Tennessee

Project Website: http://www.ece.utk.edu/~itamar/DOE_Fabrics

Abstract

During the next few years, complex DOE scientific experiments are expected to generate several petabytes of data that will be transferred to geographically distributed terascale computing facilities. With the evolution of large-scale DOE networks, such as UltraNet, similar engineering challenges to those facing Internet architects are expected to emerge. One of these key challenges is the design of scalable, high-performance packet switching fabrics that reliably facilitate the exchange of data between the network nodes. This project focuses on the design, analysis and prototype demonstration of a novel switch fabric that facilitate the provisioning of flexible, packet-based network services over the existing optical infrastructure.

Project Description

UltraNet currently supports the dynamic bandwidth allocation of circuits residing in optical carriers. Layer-1 bonding has been proposed, and is currently being studied, as means of aggregating multiple wavelength resources whenever additional bandwidth is required. This project addresses what occurs *within the virtual circuits*. In particular, the aim is to offer flexible quality of service (QoS) provisioning, in the form of packet switching services, which are seamlessly implemented over the virtual circuit infrastructure. Unfortunately, due to the unique nature of UltraNet, there are currently no commercial products that offer these desired features. The key approach taken in this project is the design of a customized packet switching fabric. The latter will be cost-efficient, scalable, robust to diverse traffic conditions and flexible in terms of reconfigurable service parameters. Figure 1 depicts the proposed system architecture.

Coarsely speaking, it is proposed to employ the switch fabric between the multiple source ports and the egress optical links. At the core of the fabric is a novel memory management algorithm that facilitates output-queueing emulation. The latter yields work conservation as well as controllable (and guaranteed) QoS provisioning. By supporting multiple layer-2 protocols, the switch core is independent of the type of traffic traversing the network. A comprehensive proof-of-concept of the switching fabric will be designed and evaluated using FPGA devices. These devices have been carefully selected so as to optimally match the performance characteristics of the system.




Fig. 1: The proposed system architecture. A reconfigurable switching fabric, which emulates a pure output-queued switch, will facilitate dynamic per-flow QoS provisioning independent of the layer-2 protocols employed.

The following key accomplishments have been made thus far:

- ❑ A novel high-speed memory management algorithm, which resides at the core of the switching fabric proposed, has been developed and analyzed (published in *IEEE ICC 2005*). This algorithm assumes homogeneous service levels, based on which algorithms that offer heterogeneous QoS guarantees will be derived.
- ❑ An FPGA (Xilinx Virtex 4) device has been carefully selected, with an aim to facilitate the proof-of-concept design. The latter will drive the project's next steps.
- ❑ Coherent with the original goals of this project, a theoretical study has been conducted in which conditions for stability of large switches that issue switching decisions ones every multiple packet times, were obtained (published in *IEEE Communication Letters*).
- ❑ Due to the bursty nature of scientific network traffic, a pragmatic model for flexible traffic generation has been devised and studied. Moreover, the scenario of bursty traffic, distributed over multiple destinations, has been investigated (published in *IEEE Communication Letters*). This result will help evaluate (in simulation as well as hardware platforms) the performance of the switching fabrics developed.

The first steps have been taken towards designing a pragmatic tool for offering advanced packet-switching services in future UltraNet deployments. Initial results indicate that based on available off-the-shelf components (primarily FPGA devices) together with novel algorithms to be executed in custom hardware, a powerful network appliance can be built. During the next two years of this project, the initial results will be expanded and a proof-of-concept platform will be developed.