Chapter 9

Conclusions

This dissertation addressed the problem of building a novel architecture for High-Speed Networks (HSNs) that solves two fundamental limitations of existing architectures:

- Scaling with respect to link speeds;
- Strict guarantees of the Quality of Service (QoS) in the communications.

Isochronets is the architecture that solves both limitations. The main advantages of Isochronets are:

- *Support of end-to-end transport of multi-protocol frames.* Isochronets do not rely on nor need to adapt to any frame structure because frames do not need to be processed inside the network to be switched. Multiple frame structures (for example, IP, ATM, etc.) can coexist in the network without adaptation.
- *Support of tunable guaranteed QoS of end-to-end transport.* Through band allocation, sources can be given priority in a tree. When transmitting in this mode, traffic has guaranteed bandwidth and does not experience any jitter, loss, or delay introduced by contention with other sources. Consequently, QoS is strictly guaranteed through tunable bands.
- *Support of asynchronous, synchronous and isochronous traffic streams.* Synchronization signals are forwarded to network interfaces when bands begin and can

be used to synchronize traffic motions. Nevertheless, within a band, motion is asynchronous. Depending on the kind of service a band provides (priority or contention), the end-to-end motion may attain every synchronization need.

- *Minimization of processing and queueing latency in the transport path.* Switching is accomplished by re-configuring nodes over time. Frames are not processed to be switched. As a consequence, nodes do not need to process frames at incoming high-speed rates. Rather, nodes need only re-configure themselves periodically. Additionally, queueing only happens to access outgoing links and can be finely controlled, and even avoided completely by assigning priorities to sources in a band.

The following challenges in designing, building, and analyzing Isochronets were addressed in this work.

1. Initial protocols for tree allocation, band allocation, and band synchronization were developed.
    - All protocols can be implemented off-line and execute at slow speed relative to transmissions.
    - The tree allocation problem is solved by an exhaustive search algorithm that searches all possible combinations of non-interfering trees to find the one that leads to the minimal number of bands. This is a reasonable solution because tree allocations happen only when the network topology changes.
    - The band synchronization problem is solved by adjusting the link propagation delays in the network to make the final delay a multiple of the clock period.

This fact trivializes synchronization because all that is necessary, after link delay adjustment, is to start and end each band exactly at the same moment within each cycle.

- The band allocation problem sizes the bands by using a dynamic algorithm based on the Least Recently Used (LRU) memory allocation policy. It operates by shrinking band sizes for destinations that are not accessed frequently and then increasing band sizes of destinations that need higher demands.

2. A novel scheme for real-time service provision at periphery protocol stacks was developed.

   - The novel Loosely-synchronous Transfer Mode (LTM) was created to support real-time services. It issues loosely-synchronous signals to periphery nodes indicating current network status, that is, destinations that are reachable and their associated QoS. As a consequence:
     * Sources and destinations are synchronized with each other and with the network using the loosely-synchronous signals.
     * Multiple protocol frames can be transferred without fragmentation and reassembly because there is no pre-defined frame structure in the communication.
     * QoS can be controlled and guaranteed because the loosely-synchronous signals coordinate interactions between sources to avoid contention at intermediate nodes.
   - The Synchronous Protocol Stack (SPS) was defined. SPS is any stack ex-

tended with the loosely synchronized signals provided by the LTM. An SPS forwards LTM synchronization signals through its layers up to the applications. In the SPS:

* Existing protocol stacks are supported unchanged.
* Existing stacks are extended with novel source and destination synchronization to support a variety of application synchronization needs.

3. The feasibility of designing and implementing an electronic and an optical Isochronet switch (Isoswitch) was demonstrated.

- An electronic Isoswitch has been designed and implemented. It has four input and four output ports each operating at 1 Gb/s. The design addressed the following challenges:
  * The design is scaleable with respect to number of channels and link speeds.
  * The implementation is easy to integrate with existing hardware.
  * The Isoswitch can provide novel LTM synchronization services to periphery protocol stacks.
- An interface card between the Isoswitch and a Sun SPARC 1 machine has been built. The card has a nominal throughput of 22 Mb/s. The card signals the beginning of band to the SPARC processor and can thus be directly used to implement LTM and SPS.
- A preliminary design of an all-optical switch implementation is proposed. The design proposed uses Wavelength Division Multiplexing (WDM) to allocate one wavelength for each band. The design has a few advantages:

- No conversion of en-route signals to the electronic domain, which makes possible implementations potentially at hundreds of terabits per second.
- Fewer wavelengths are necessary than in pure WDM, namely one per band. In the worst case, the number of necessary wavelengths is the number of destinations in the network.

4. An initial performance study of Isochronets was developed.
   - An analytical model was developed to study mean delay and loss in RDMA. It was assumed that frame arrivals happen according to a Poisson process.
     - The mean loss in RDMA- was approximated by an approximate closed-form expression. It was used to show that the maximum loss in RDMA- is bound to 50%.
     - The mean delay under RDMA++ when the input load is low to medium was approximated by a closed-form expression.
     - The mean delay in RDMA+ was approximated by a closed-form expression. The loss rate was shown to be negligible if the band size is large compared to the mean frame size.
   - Simulation studies validated the analytical expressions and compared RDMA end-to-end delay performance with Packet Switching (PS) and Circuit Switching (CS). The general performance trend is the following. For low input loads, PS outperforms RDMA and CS because there is no admission delay waiting for a band or circuit. When the load increases, the processor becomes saturated and then RDMA outperforms PS. CS always performs worst than

RDMA because of the penalty incurred waiting for a circuit, which is larger than the one waiting for a band.