Scalable Switching Testbed not "Stopping" the Serial Bit Stream

D. Agrawal[◆], M. Baldi[●], M. Corrà[■], G. Fontana[◆], G. Marchetto[●], V. T. Nguyen[◆],
Y. Ofek[◆], D. Severina[◆], T. H. Truong[◆], O. Zadedyurina[◆]

[◆]Department of Information and Communication Technology
University of Trento, Via Sommarive 14, I-38050 Povo, Trento, Italy
[■] TRETEC S.r.l., Via Solteri 38, 38100 Trento, Italy
[●] Control and Computer Engineering Department, Politecnico di Torino
Corso Duca degli Abruzzi, 24, 10129 Torino, Italy

***Abstract* — In order to achieve ultra scalable IP packet switching it is essential to minimize "stopping" of the serial bit streams. In our recent experimental work we demonstrated how this can be achieved with an ultra-scalable switching architecture reaching multi-terabits per second (10-100 Tb/s) in a single chassis. The implemented testbed uses only off-the-shelf optical and electronic components. The scalability of this architecture is the direct outcome of how global time (i.e., UTC – coordinated universal time) and pipeline forwarding are utilized. The paper presents the design of a prototype switch and experimental activity with it.**

***Index Terms*—optical networks, terabit switching, sub-lambda switching, time-driven switching, pipeline forwarding, streaming media**

I. INTRODUCTION

The Internet has been growing steadily in the past few years. However, services so far deployed over the Internet are nothing compared to the ones that can still be deployed – i.e., we are only at the very beginning. Thus, one may envision that service-wise and traffic-wise growth of the Internet applications is yet to come; one may say that "the best of the Internet is ahead of us". One likely scenario is that the future Internet will be dominated by on-demand applications such as HDTV (perhaps in 3D), video on demand, high quality videoconferencing, distributed gaming, (3D) virtual reality, and many more. These on-demand applications are likely to generate traffic that is either streaming by nature or can be efficiently mapped into streaming and handled as such (e.g., very large file transfer).

This paper presents a testbed demonstrating a method known as *time-driven switching* (TDS) or *fractional λ switching* (FλS) with *pipeline forwarding* (see, for example, [1][2][3][4] and Section II.C) that is particularly suitable to carry streaming media applications that are transported in IP packets (without segmentation and header processing). The necessary condition for pipeline forwarding is having *common time reference* (CTR), which is easily realized with UTC (coordinated universal time) or the standard time-of-day. UTC is globally available via GPS (global positioning system) with accuracy between any two points around the globe well below 1 μs. Specifically, in our implementation we are using the standard UTC second for scheduling the TDS switching operation. Consequently, this design offers:

1. High scalability of network switches (multi-terabit/s in a single chassis, the focus of this paper),
2. Quality of service guarantees (deterministic delay and jitter, no loss) for (UDP-based) constant bit rate (CBR) and variable bit rate (VBR) streaming applications — as needed, while
3. Preserving the support of elastic, TCP-based traffic, i.e., existing applications based on "best-effort" services are not affected in any way.

The testbed implemented in this project is novel network architecture for traffic engineering of (UDP-based) streaming media. Traffic engineering of UDP-based applications is important due to the fact that more and more streaming media applications over the Internet are using UDP. Such applications need the right level of service quality in order to satisfy users' requirements. Note that using over-provisioning in order to solve the service quality problem is not scaleable. Alternatively, the presented solution enables the realization of low complexity, highly scalable IP switches with full utilization.

Section II focuses on the scalable switch design using off-the-shelf components. The architecture and the implementation of the switching system (fabric and switch controller) are presented in Section III. Section IV presents the prototype and some testing results, most significantly with six nodes and 100 km of single mode fiber (4 fiber segments of 25 km) and Section V concludes the presented work.

A. Related Works

First note that traditional solutions for quality of service in asynchronous packet switching, such as [5] (PGPS – packet generalized process sharing), are very different from our approach. Being asynchronous, these solutions do not require a global clock; however, pipeline forwarding of IP packets cannot be implemented. Consequently, asynchronous packet switching approaches require high complexity switches with very large buffers (a lot of "stopping" of the serial bit stream), and thus, they are more expansive, less scalable and with

[1] This work was supported in part by funds from the European Commission (contract N° MC-EXC 002807).

poorer quality of service (e.g., loss and delay).

CTR (which can be realized with UTC) provides **phase synchronization** with **identical frequencies** (derived from the UTC second) everywhere. Traditional TDM (time division multiplexing) systems, such as SONET/SDH, represent another synchronous solution, but they use only frequency synchronization with known bounds on drifts. In order to overcome possible data loss, SONET/SDH is using rather complex overhead information to accommodate: (1) the accumulation of delay uncertainties and (2) continuous clock drifts from the nominal value. Multiplexing in SONET/SDH is based on time slots (TS) organized in a reoccurring structured called frame. Due to the lack of phase synchronization among nodes, a TS incoming from an input interface might be stored up to the duration of a whole frame before being sent out on its output link. In order to keep the delay introduced by each node small, the frame duration is defined as 125 μs and, consequently, the TS is defined to keep small-size data units, i.e., 1 byte. The TS duration in 10Gbps (OC-192/STM-48) is about one nanosecond.

In the past fifteen years there were a number of works on combining WDM with TDM [8]-[15]. None of these works was using CTR with pipeline forwarding and some lack addressing critical timing issues. Specifically, regarding the accumulation of delay uncertainties or jitter and clock drifts, which is solved by using either SONET or CTR with pipeline forwarding, as discussed in Section II.C.

In [8], an optical *time slot interchange* (TSI) utilizing sophisticated optical delay lines is described with no detailed timing analysis. In [9] and [10] two experimental optical systems with in-band master clock distribution and optical delay lines are described, with only limited discussion on timing issues. In [11] a system with constant delays and clocks is described, which can be viewed as a close model to *immediate forwarding* (see Section II.C); however, no timing analysis and no consideration of *non-immediate forwarding* (see Section II.C) analyzed in [1][2] is provided.

In [12] there is a TSI system with a detailed design of an optical slot permuter without using CTR, which consequently has higher implementation complexity. The blocking probability analysis in [13] and [14] models a network with TSI switches while ignoring timing issues, such as, delays and timing errors. Thus, it is not possible to determine whether the analysis is applicable to a network with CTR. The analysis in [15] provides more details than in [13] and [14] regarding timing issues in the network model; specifically, a synchronizer that aligns incoming time slots while ensuring that the delay between nodes is an integer number of time slots is discussed. The issue in this case, is how the alignment is performed, e.g., how to align when one incoming time slot starts exactly in the middle of another incoming time slot. Furthermore, if there are accumulated timing errors and clock drifts the synchronizer operation is complicated and not likely to be implemented in the optical domain (see, for example, the timing analysis in [16]). However, if the network model and analysis in [15] is implicitly assuming both common time reference (CTR) and pipeline forwarding, then the alignment operation with respect to CTR is simple. Consequently, the analysis may be applicable to networks with CTR and pipeline forwarding.

II. Scalable Design with Off-the-shelf Components

A. Why it is scalable

In order to put our testbed in perspective, note that Cisco's top-of-the-line router with a novel network processor design, CRS-1, has only 640Gb/s switching capacity per chassis [the announcement of 92Tb/s aggregated input/output capacity should be divided by 2 (for taking into account that bits exiting the switch outputs are the same that previously entered from its inputs and they were switched only once) and then by 72 chassis's], which represents a factor of 2 improvement after 5 years of development and 500 million dollars of investment with respect to the previous Cisco 12000. So, if Internet traffic is doubling, say, every 18 months there is a real switching bottleneck at the horizon.

This example suggests that existing asynchronous IP switching architectures are limited to about one Tb/s in a single chassis. On the other hand, scalability of all-optical switching was successful only for whole lambda (optical channel) switching. This implies capacity provisioning of an entire (whole) optical channel capacity or nothing. Thus, this approach also suffers from poor scalability. Thus, in order to maximize scalability the following three implementation principles were used (the motivations and justifications will be presented in the following sections):

1. Electronic switching components, with
2. Optical interconnection of the electronic switching components, and with
3. Global pi*peline forwarding* (PF) with *time-driven switching* (TDS) and control.

Today, there are single-chip, low-cost, high-capacity electronic cross-points (for example, 144-by-144 switch matrix with 0-11 Gb/s per input/output ports, i.e., more that 1.5 Tb/s aggregate switching capacity!). However, constructing a large switching matrix based on such cross-points may require optical interconnects (see, for example, [6][7]). Optical interconnects allow, at least in principle, any desired interconnection topology, while minimizing various noise and interference sources. Finally, by using global time it is possible to minimize the following switching complexity components:

1. Input buffer size while eliminating the need for output buffers all together (i.e., space complexity measure),
2. Number of switching elements in the electronic switching fabric (i.e., space complexity measure), and
3. Number of operation required by the electronic switch controller in order to continuously configure the switching fabric input ports to output ports permutation (i.e., time complexity measure).

To summarize the above discussion, our design was guided by the clear limitations of both "all-optical" and "all-electrical" switching approaches. Consequently, our current conclusion is that a hybrid or "best-of-breed" switching solution in order to achieve 10-100 Tb/s switching capacity in a single chassis.

Specifically, our current design approach is based on electronic switching with optical interconnects and without "stopping" the serial bit stream. This may change as optical switching components evolve.

It is worth noting that no physical breakthrough is required in order to design dense electrical switching with optical interconnections, although there are many technical challenges. In other words, our switch design is base on using electrons for switching data, and photons for transporting data between electronic switching elements.

B. Optimal switch design

The switch design presented in this work is characterized by the following optimized components:

1. Input ports with optimal memory access speedup of 1 (s=1), where speedup is defined as the ratio between the link bandwidth and the memory access bandwidth.
2. Input ports with a single queue/buffer (note that typical input buffered switches have N queues – one queue per output in order to prevent head-of-the-line blocking).
3. Optimal output port design without buffers, i.e., the data packets are forwarded from the switching fabric directly onto the out-going link.
4. Optimal switching fabric complexity by using multistage Banyan interconnection networks [17], with switching complexity of O(a*N*lg_aN), where N is the number of input/output and a is the size of each switching element or block.
5. Switch control complexity which is equal to the fabric complexity, O(a*N*lg_aN), where N is the number of input/output ports and a is the size of each switching element or block.

C. UTC with pipeline forwarding

Pipeline forwarding is a known optimal method that is extensively used in manufacturing (e.g., automotive assembly line) and computer architecture (e.g., reduced instruction set computer). In this work pipeline forwarding is used as the basic operating principle for *time-driven switching* (TDS) (also known as sub-lambda or *fractional lambda switching* (FλS)). As in other pipeline forwarding implementations, the necessary operating requirement for realizing TDS is having *common time reference* (CTR). In the context of global network CTR is realized by using UTC (coordinated universal time) that is globally available via GPS and Galileo (in the near future).

For implementing TDS with pipeline forwarding the UTC second is partitioned into a predefined number of time-frames (TFs), as shown in Figure 1. TFs can be viewed as virtual containers for multiple IP packets that are switched at every TDS switch, which are UTC synchronized. In Figure 1, a group of K TFs forms a time-cycle (TC); L contiguous time-cycles are grouped into a super cycle (SC) or one UTC second.

The necessary condition of pipeline forwarding over a sequence of TDS switches is that the input-to-input delay between any pair of switches should be an integer number of TFs.

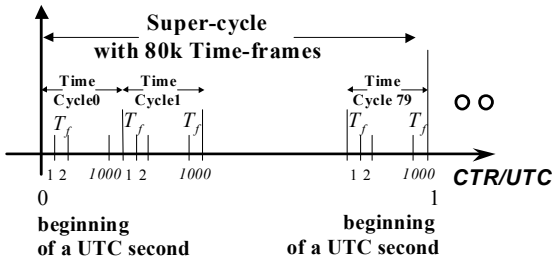

Figure 1: CTR/UTC for pipeline forwarding in TDS

Time driven switching enables the switching of IP packets in time frames (TFs) without decoding or processing the IP packet headers. It is based on the setting of a switching schedule of IP packets in TFs along a predefined route/path in the TDS network.

In TDS there are two main types of TF forwarding, wherein each TF contains multiple packets:

1. Immediate forwarding (IF): upon the arrival of each TF to an FλS switch, the content of the TF (e.g., IP packets) is scheduled to be "immediately" switched and forwarded to the next switch during the next TF. Hence, the buffer that is required is bounded to one TF and the end-to-end transmission delay is minimized. Clearly, if the IP packets in TFs arrive to the switch at the exact time no buffers are required. Consequently, the amount of alignment buffer for one TF of 12.5μs at 10Gb/s is only 15KB.
2. Non-immediate forwarding (NIF): TFs may be delayed at the input of the switch for a predefined number of additional TFs, typically, one or two (e.g., *2-frame forwarding* or *k-forwarding*). The main advantage of NIF is reduction in blocking probability, which is the probability that there are available TFs but not in the right sequence or schedule.

D. Switching scalability

Although Banyan [17] is known to suffer from *space blocking* (where there is no available path (route) through the switch even though the input port and output port are available) with UTC-based pipeline forwarding of data packets *space blocking* has only a diminishing effect [1]. Intuitively, the multiple TFs in each time cycle provide an additional degree of freedom for scheduling, i.e., if there are 1000 TFs in each time cycle there are 1000 different possible switch instances that can be used for scheduling on every switch. In addition, non-immediate forwarding of TFs can be used, which implies that a TF (as virtual container with multiple IP packets) may be delayed one or more TFs before being switched and forwarded. In summary, the combination of plurality of TFs in each time cycle together with non-immediate forwarding enable the efficient use Banyan-based switching fabric with optimal switching complexity.

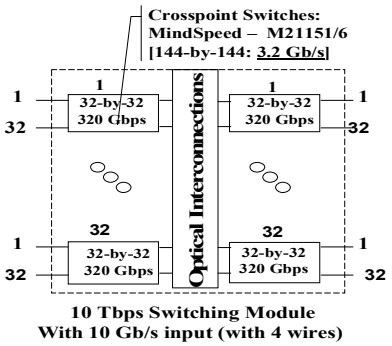

Figure 2: 10 Tb/s Switching Module with off-the-shelf M21151/M21156 implemented in the current testbed

Figure 2 and Figure 3 are two possible Banyan-based switching fabric configurations. The first configuration, Figure 2, is based on the Mindspeed M21151/M21156 cross-point switches that was used in our current prototype, as describe in the following sections. The Banyan design using M21151/6 has aggregate switching capacity of 10 Tb/s. The switch design in Figure 3 is based on a state-of-the-art cross-point, VSC3040, made by Vitesse. This cross-point has serial input/output rate of up to 11Gb/s with 144 input and output ports (the M21151/6 has only 3.2Gb/s). Using this new cross-point it is possible to construct a Banyan-based switching fabric with aggregate switching capacity of 160 Tb/s (128*128*10Gb/s), as shown in Figure 3.

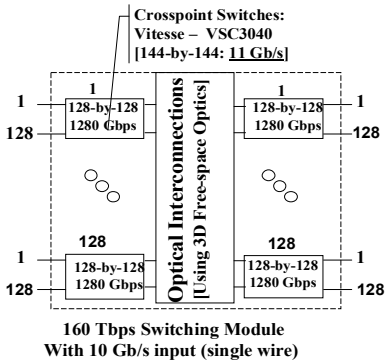

Figure 3: 160 Tb/s Switching Module with off-the-shelf VSC3040 for future testbed

Although the Mindspeed M21151/6 and Vitesse VSC3040 electronic cross-point switches provide effective switching solution in the electronic domain, interconnecting such a large switching matrix electrically is not trivial due to noise, cross talk and various other interference problems. Therefore, as shown in Figure 2 and Figure 3, the best way to construct such large switching fabrics is by using optical interconnects. One interesting optical interconnection configuration is using free-space optic in three dimensions. More specifically, using high-density (10 per square centimeter) short-range (50 centimeters) VCSELs (vertical cavity surface emitting laser) which shine their light upwards, from the outputs of the first stage of the Banyan cross-point switches, then the light is reflected downwards, by micro-mirrors, to the inputs of the second stage of the Banyan cross-point switches. However, details related to this solution are left for further research.

In comparison, optical switches with below 1μs switching time have lower capacity, much larger physical size (due to the high complexity of electronic circuitry required for controlling the tiny optical switching device), and consequently, are much more expansive. The main challenge in constructing the large switching fabrics shown in Figure 2 and Figure 3 is interconnection. As was mentioned, the most promising approach is a hybrid opto-electronic switching fabric.

III. The TDS Switch Architecture and Implementation

The functional diagram of the TDS (time-driven switching) design, implemented in our testbed, is shown in Figure 4. It has three major parts: field programmable gate-array (FPGA) switch controller, switch fabric constructed with off-the-shelf M21151/M21156 cross-point switches, and a GPS time receiver. The switch fabric is controlled by the switch controller, which is responsive to timing signals from the GPS receiver: 1PPS (1 pulse per (UTC) second) and 10MHz.

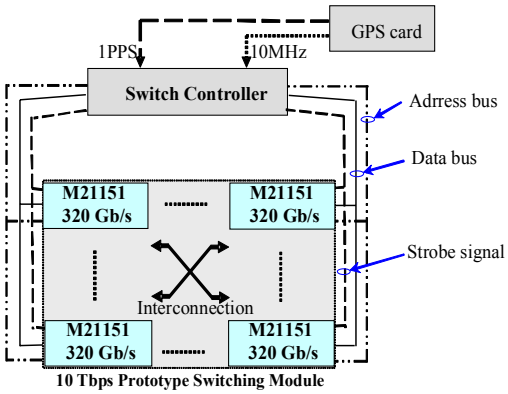

Figure 4: Block diagram of TDS switch: 2-stage Banyan implemented with M21151/M21156 cross-point switches

A. General description

The GPS receiver provides 1 PPS (pulse per (UTC) second) and 10MHz signals to the FPGA switch controller. The communication between the switch controller and the cross-point switches can be parallel or serial. There are three main signal types that connect the switch controller with the cross-point switches: address, data and strobe signals, as shown in Figure 4. The cross-point configuration information, to which input each output should be connected, for all 144 outputs, and for each TF within the time cycle. This information is stored in the memory table of the FPGA switch controller. The data and address signals are used for writing the switch configuration for the next TF onto the cross-point switches. This writing process should end before the falling edge of the next strobe signal, which corresponds to start of the next TF. At the falling edge of the next strobe signal the next switching configuration is latched onto the cross-points. The next cross-point switch configuration is ready in less than ten nanoseconds.

B. GPS receiver

GPS receiver, EPSILON Board OEM II, provides accurate and stable time and frequency signals for UTC (coordinated universal time) synchronization. It provides 1PPS (pulse per second) and 10MHz sine wave, and UTC time-of-day output.

Furthermore, the 10MHz frequency reference is cycle locked to the 1PPS, which is the standard UTC second. This implies that within 1PPS there are exactly 10,000,000 cycles of the 10 MHz output from the GPS card.

C. Mindspeed cross-point switch

Mindspeed cross-point M21151/M21156; the primary component of this switch implementation, which is a low-power CMOS, high-speed 144-by-144. Each input can receive serial signal from 0 to 3.2 GHz (thus, its maximum aggregate capacity is 460 GB/s). The serial input signal goes through a sequence of internal inverters inside the M21151/M21156, and therefore, the serial signal doesn't have any power loss. The crosspoint switch ports are equipped with input equalization (IE) and output integrated clock data recovery (CDR), which further preserves the serial signal quality. Each CDR is preceded by a programmable IE. Implementation and design of switch control card is described in following section.

IV. PROTOTYPE OVERVIEW AND TESTING

In the presented testbed, shown in Figure 5, two video streams are generated by a video server (to the left), transported, with *deterministic quality of service*, through a network of one router and two multi-terabit/s switches (all implementing pipeline forwarding) and delivered to two different video clients. IP packets carrying video samples are transported unchanged, as a whole, end-to-end. Namely, no change can be seen by observing packets flowing on any link of the testbed as only conventional IP packets encapsulated into Ethernet frames travel across the network testbed.

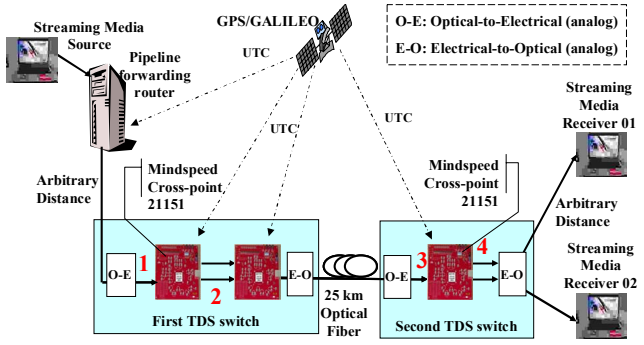

Figure 5: Initial testbed setup without "stopping" the bits

A. Initial testbed set-up

The initial switch prototype experimental setup is shown in Figure 5. The main components are streaming media sources, a network interface for scheduling packet entrance into the first switch, a 25km single mode optical fiber, two stage source side switch, and single stage receiver side switch. The deployed cross-point switches have 144 channels with capacity of 3.2 Gb/s each. Two streaming media flows for two different receivers — a DVD movie with soundtrack and subtitles and an animation movie with soundtrack — are transmitted from one PC using VLC media player. Asynchronous packets are sent to the network interface (details about the network interface are outside of the scope of this paper) that, based on the 1PPS signal from GPS, forwards the streaming media packets from two sources to the first switch in a periodic fashion, namely during predefined TFs of each time cycle. The first cross-point switch splits packets into two streams forwarded to second cross-point on separate channels. At second cross-point both streams are multiplexed again into one channel that is transmitted via the 25km single mode optical fiber link to the second switch in Figure 5.

B. Advanced testbed testing

To evaluate the robustness of the switch prototype and to know its limits we have extended our testing to a 100km metro-like network shown in Figure 6. It consists of six nodes, each including either one or two cross-point switches. Five of the six nodes are connected optically through four segments of 25km optical fiber as shown in Figure 6.

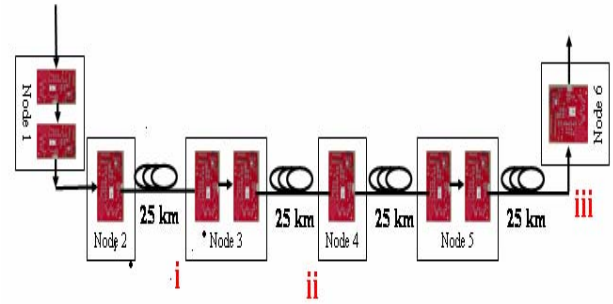

Figure 6: 100 km, 6-node metropolitan network

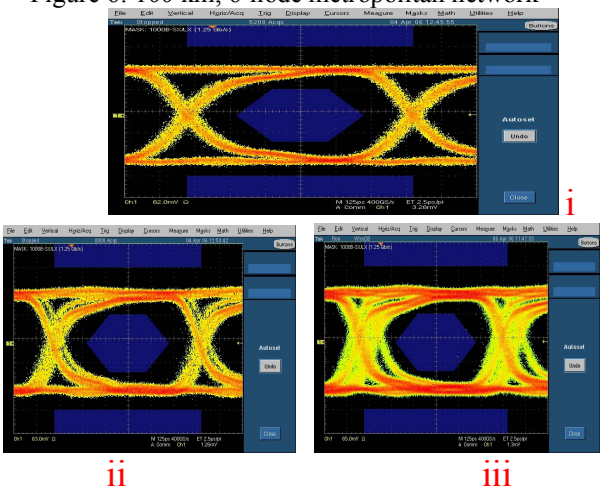

Figure 7: Eye Pattern at locations (i), (ii) and (iii)

The eye pattern test is a quick method for visually examining the quality of serial signals, e.g., the amount of timing jitter and amplitude variation in a serial data stream. All possible transitions of the signal — positive-going, negative going, leading, and trailing — are displayed in a single diagram. This single diagram provides information about the eye opening, noise, jitter, rise and fall times, and amplitude. Figure 7 shows actual eye diagrams as captured on a real-time Tektronics, 6GHz digital storage oscilloscope. The eye pattern test is performed and measurements were taken at locations: (i), (ii) and (iii), shown in Figure 6. Figure 7 shows the eye pattern measurements after 25km (location i), after 50km

(location ii) and after 100km (location iii). The oscilloscope has an inbuilt standard Gigabit Ethernet 1000 base SX-LX mask test. It is apparent from the eye pattern that there is reduction of safety margin as the optical fiber length increases. However, the signal is compliant to the 1000 base SX-LX standard even after having traveled for 100km.

V. Conclusion

Implementing UTC-based pipeline forwarding in a real testbed that is scalable to multi-terabit/s switching capacity has been a rewarding experience. The implementation success is a direct outcome of the simplicity of the pipeline forwarding method. The beauty is that the simplicity of this realization did not compromise two most desired performance properties for the future Internet: (1) sub-lambda switching scalability to 10 and 160 Tb/s in a single chassis and (2) predictable quality of service (QoS) performance for streaming media and large file transfers – as possibly needed in content distribution and Grid computing. Cisco's top-of-the-line router, CRS-1, has only 640 Gb/s per chassis, which represent a factor of 2 improvement, after 5 years of development. So if the Internet traffic is doubling every, say, 18 months there is a real sub-lambda switching bottleneck on the horizon, for which this experimental work is a viable solution.

Although the presented prototype is based on optically interconnected electronic switches, the serial bit streams are never "stopped", which is critical in our current plans for an all-optical implementation. In recent experiments it was possible to transmit the serial bit streams through six nodes and 100 km of fiber without "stopping" them. This is significant as it demonstrates that the deployed technology is suitable for providing (1) ultra-scalable switching capacity and (2) sub-lambda reconfigurable optical add-drop multiplexing (ROADM) in metropolitan networks while minimizing the provider cost per end-user.

References

[1] M. Baldi and Y. Ofek, "Fractional Lambda Switching - Principles of Operation and Performance Issues", *SIMULATION: Transactions of The Society for Modeling and Simulation International*, Vol. 80, No. 10, Oct. 2004.

[2] D. Grieco, A. Pattavina and Y. Ofek, "Fractional Lambda Switching for Flexible Bandwidth Provisioning in WDM Networks: Principles and Performance", *Photonic Network Communications*, Issue: Volume 9, Number 3, Date: May 2005.

[3] M. Baldi, Y. Ofek, "Fractional lambda switching," *Proc. of ICC 2002*, New York, USA, Vol.5, pp 2692-2696.

[4] A. Pattavina, M. Bonomi, Y. Ofek, "Performance evaluation of time driven switching for flexible bandwidth provisioning in WDM networks," *Proc. of Globecom 2004*, Texas, USA, Vol. 3.

[5] A. K. Parekh and R. G. Gallager, "A generalized processor sharing approach to flow control – the multiple node case," IEEE/ACM Transactions on Networking, 2(2):137–150, 1994.

[6] W. Gibbs, "Computing at the Speed of Light" Scientific American, Nov. 2004, pp. 80-87.

[7] "Silicon Photonics," Edited by L. Pavesi and D. J. Lockwood. Springer, 2004.

[8] D. K. Hunter and D. G. Smith, "New architectures for optical TDM switching," IEEE/OSA Journal of Lightwave Technology, vol. 11, no. 3, pp. 495-511, Mar. 1993.

[9] I. P. Kaminow et al., "A wideband all-optical WDM network," IEEE Journal on Selected Areas in Communications, vol. 14, no. 5, pp. 780-799, June 1996.

[10] P. Gambini et al., "Transparent optical packet switching: network architecture and demonstrators in the KEOPS project," IEEE Journal on Selected Areas in Communications, vol. 16, no. 7, pp. 1245-1257, Sept. 1998.

[11] Nen-Fu Huang, Guan-Hsiung Liaw, and Chuan-Pwu Wang, "A novel all-optical transport network with time-shared wavelength channels," IEEE Journal on Selected Areas in Communications, vol. 18, no. 10, pp. 1863-1875, Oct. 2000.

[12] H. F. Jordan, D. Lee, K. Y. Lee, and S. V. Ramanan, "Serial array time slot interchangers and optical implementations," IEEE Transactions on Computers, vol. 43, no. 11, Nov. 1994.

[13] R. Srinivasan, A. K. Somani, "A Generalized Framework for Analyzing Time–Space Switched Optical Networks," IEEE Journal on Selected Areas in Communications, vol. 20, no. 1, pp. 202-215, January 2002.

[14] S. Subramaniam, E. J. Harder, and H. Choi, "Scheduling Multirate Sessions in Time Division Multiplexed Wavelength-Routing Networks," IEEE Journal on Selected Areas in Communications, vol. 18, no. 10, October 2000.

[15] B. Wen and K. Sivalingam "Routing, Wavelength and Time-Slot Assignment in Time Division Multiplexed Wavelength-Routed Optical WDM Networks," IEEE INFOCOM 2002.

[16] Y. Ofek, "Generating a Fault Tolerant Global Clock using High-speed Control Signals for the MetaNet Architecture," IEEE Transactions on Communications, May 1994.

[17] L. R. Goke, G. J. Lipovski, "Banyan Networks for Partitioning Multiprocessor Systems,"1st Annual Symposium on Computer Architecture, Dec. 1973, pp. 21-28.