

Scaling Internet Routers Using Optics*

Isaac Keslassy, Shang-Tse Chuang, Kyoungsik Yu,
David Miller, Mark Horowitz, Olav Solgaard, Nick McKeown
Stanford University

ABSTRACT

Routers built around a single-stage crossbar and a centralized scheduler do not scale, and (in practice) do not provide the throughput guarantees that network operators need to make efficient use of their expensive long-haul links. In this paper we consider how optics can be used to scale capacity and reduce power in a router. We start with the promising *load-balanced switch architecture* proposed by C-S. Chang. This approach eliminates the scheduler, is scalable, and guarantees 100% throughput for a broad class of traffic. But several problems need to be solved to make this architecture practical: (1) Packets can be mis-sequenced, (2) Pathological periodic traffic patterns can make throughput arbitrarily small, (3) The architecture requires a rapidly configuring switch fabric, and (4) It does not work when linecards are missing or have failed. In this paper we solve each problem in turn, and describe new architectures that include our solutions. We motivate our work by designing a 100Tb/s packet-switched router arranged as 640 linecards, each operating at 160Gb/s. We describe two different implementations based on technology available within the next three years.

Categories and Subject Descriptors

C.2 [Networking]: Routers

General Terms

Algorithms, Design, Performance.

Keywords

Load-balancing, packet-switch, Internet router.

*This work was funded in part by the DARPA/MARCO Center for Circuits, Systems and Software, by the DARPA/MARCO Interconnect Focus Center, Cisco Systems, Texas Instruments, Stanford Networking Research Center, Stanford Photonics Research Center, and a Wakely Stanford Graduate Fellowship.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SIGCOMM'03, August 25–29, 2003, Karlsruhe, Germany.
Copyright 2003 ACM 1-58113-735-4/03/0008 ...\$5.00.

1. INTRODUCTION AND MOTIVATION

This paper is motivated by two questions: First, how can the capacity of Internet routers scale to keep up with growths in Internet traffic? And second, can optical technology be introduced *inside* routers to help increase their capacity?

Before we try to answer these questions, it is worth asking if the questions are still relevant. After all, the Internet is widely reported to have a glut of capacity, with average link utilization below 10%, and a large fraction of installed but unused link capacity [1]. The introduction of new routers has been delayed, suggesting that faster routers are not needed as urgently as we once thought.

While it is not the goal of this paper to argue *when* new routers will be needed, we argue that the capacity of routers must continue to grow. The underlying demand for network capacity (measured by the amount of user traffic) continues to double every year [2], and if this continues, will require an increase in router capacity. Otherwise, Internet providers must double the number of routers in their network each year, which is impractical for a number of reasons: First, it would require doubling either the size or the number of central offices each year. But central offices are reportedly full already [3], with limited space, power supply and ability to dissipate power from racks of equipment. And second, doubling the number of locations would require enormous capital investment and increases in the support and maintenance infrastructure to manage the enlarged network. Yet this still would not suffice; additional routers are needed to interconnect other routers in the enlarged topology, so it takes more than twice as many routers to carry twice as much user traffic with the same link utilization. Instead, it seems reasonable to expect that router capacity will continue to grow, with routers periodically replaced with newer higher capacity systems.

Historically, routing capacity per unit volume has doubled every eighteen months (see Figure 1).¹ If Internet traffic continues to double every year, in nine years traffic will have grown eight times more than the capacity of individual routers.

Each generation of router consumes more power than the

¹Capacity is often limited by memory bandwidth (defined here as the speed at which random packets can be retrieved from memory). Despite large improvements in I/O bandwidths, random access time has improved at only 1.1-fold every eighteen months. Router architects have therefore made great strides to introduce new techniques to overcome this limitation.




Figure 1: The growth in router capacity over time, per unit volume. Each point represents one commercial router at its date of introduction, normalized to how much capacity would fit in a single rack. Best linear fit: 2.05-fold increase every 18 months [4].

last, and it is now difficult to package a router in one rack of equipment. Network operators can supply and dissipate about 10 kW per rack, and single-rack routers have reached this limit. There has therefore been a move towards multi-rack systems, with either a remote, single-stage crossbar switch and central scheduler [5, 6, 7, 8], or a multi-stage, distributed switch [9, 10]. Multi-rack routers spread the system power over multiple racks, reducing power density. For this reason, most high-capacity routers currently under development are multi-rack systems.

Existing multi-rack systems suffer from two main problems: Unpredictable performance, and poor scalability (or both). Multi-rack systems with distributed, multistage switching fabrics (such as buffered Benes or Clos networks, hypercubes or toroids) have unpredictable performance. This presents a problem for the network operators: They don't know what utilization they can safely operate their routers at; and if the throughput is less than 100%, they are unable to use the full capacity of their expensive long-haul links. This is to be contrasted with single-stage switches for which throughput guarantees are known [11, 12].

However, single-stage switches (e.g. crossbars with combined input and output queueing) have problems of their own. Although arbitration algorithms can theoretically give 100% throughput,² they are impractical because of the complexity of the algorithm, or the speedup of the buffer memory. In practice, single-stage switch fabrics use sub-optimal schedulers (e.g. based on WFA [13] or \mathcal{S} SLIP [14]) with insufficient speedup to guarantee 100% throughput. Future higher capacity single-stage routers are not going to give throughput guarantees either: Centralized schedulers don't scale with an increase in the number of ports, or with an increase in the line-rate. Known maximal matching algorithms for centralized schedulers (PIM [15], WFA [13], \mathcal{S} SLIP [14]) need at least $O(N^2)$ interconnects for the arbitration process, where N is the number of linecards. Even if arbitration is distributed over multiple ASICs, intercon-

²For example WFA [13] with a speedup of 2, MWM with a speedup of one [12].

nect power still scales with $O(N^2)$. The fastest reported centralized scheduler (implementing maximal matches, a speedup of less than two and no 100% throughput guarantees) switches 256 ports at 10Gbps [5]. This design aims to maximize capacity with current ASIC technology, and is limited by the power dissipation and pin-count of the scheduler ASICs. Scheduler speed will grow slowly (because of the $O(N^2)$ complexity, it will grow approximately with \sqrt{N}), and will continue to limit growth.

In summary, multi-rack systems either use a multi-stage switch fabric spread over multiple racks, and have unpredictable throughput; or they use a single-stage switch fabric in a single rack that is limited by power, and use a centralized scheduler with unpredictable throughput. If a router is to have predictable throughput, its capacity is currently limited by how much switching capacity can be placed in a single rack. Today, the limit is approximately 2.5Tb/s, and is constrained by power consumption.

Our goal is to identify architectures with predictable throughput and scalable capacity. In this paper we'll explain how we can use optics with almost zero power consumption to place the switch fabric of a 100Tb/s router in a single rack, without sacrificing throughput guarantees. This is approximately 40 times greater than the electronic switching capacity that could be put in a single rack today. We describe our conclusion that the Load-Balanced switch, first described by C-S. Chang *et al.* in [16] (which extends Valiant's method [17]), is the most promising architecture. It has provably 100% throughput. It is scalable: It has no central scheduler, and is amenable to optics. It simplifies the switch fabric, replacing a frequently scheduled and re-configured switch with two identical switches that follow a fixed sequence, or are built from a mesh of WDM channels.

In what follows we will start by describing Chang's Load-Balanced switch architecture in Section 2, and explain how it guarantees 100% throughput without a scheduler. We then tackle four main problems with the basic Load-Balanced switch that make it unsuitable for use in a high-capacity router: (1) It requires a rapidly configuring switch fabric, making it difficult, or expensive to use an optical switch fabric, (2) Packets can be mis-sequenced, (3) Pathological periodic traffic patterns can make throughput arbitrarily small, and (4) It does not work when some linecards are missing or have failed. In the remainder of the paper we find practical solutions to each: In Section 4 we show how novel buffer management algorithms can prevent mis-sequencing and eliminate problems with pathological periodic traffic problems. The algorithms also make possible multiple classes of service. In Section 5 we show how problem (3) can be solved by replacing the crossbar switches by a fixed optical mesh — a powerful and perhaps surprising extension of the load-balanced switch. And then in Section 6 we explain why problem (4) is the hardest problem to solve. We describe two implementations that solve the problem: One with a hybrid electro-optical switch fabric, and one with an all-optical switch fabric.

2. LOAD-BALANCED ARCHITECTURE

2.1 The Architecture

The basic load-balanced switch is shown in Figure 2, and consists of a single stage of buffers sandwiched by two identical stages of switching. The buffer at each intermediate




Figure 2: Load-balanced router architecture

input is partitioned into N separate FIFO queues, one per output (hence we call them virtual output queues, VOQs). There are a total of N^2 VOQs in the switch.

The operation of the two switch fabrics is quite different from a normal single-stage packet switch. Instead of picking a switch configuration based on the occupancy of the queues, both switching stages walk through a fixed sequence of configurations. At time t , input i of each switch fabric is connected to output $[(i + t) \bmod N] + 1$; i.e. the configuration is a cyclic shift, and each input is connected to each output exactly $\frac{1}{N}$ -th of the time, regardless of the arriving traffic. We will call each stage a *fixed, equal-rate* switch. Although they are identical, it helps to think of the two stages as performing different functions. The first stage is a load-balancer that spreads traffic over all the VOQs. The second stage is an input-queued crossbar switch in which each VOQ is served at a fixed rate.

When a packet arrives to the first stage, the first switch immediately transfers it to a VOQ at the (intermediate) input of the second stage. The intermediate input that the packet goes to depends on the current configuration of the load-balancer. The packet is put into the VOQ at the intermediate input according to its eventual output. Sometime later, the VOQ will be served by the second fixed, equal-rate switch. The packet will then be transferred across the second switch to its output, from where it will depart the system.

2.2 100% Throughput

At first glance, it is not obvious how the load-balanced switch can make any throughput guarantees; after all, the sequence of switch configurations is pre-determined, regardless of the traffic or the state of the queues. In a conventional single-stage crossbar switch, throughput guarantees are only possible if a scheduler configures the switch based on knowledge of the state of all the queues in the system. In what follows, we will give an intuitive explanation of the archi-

ture, followed by an outline of a proof that it guarantees 100% throughput for a broad class of traffic.

Intuition: Consider a single fixed, equal-rate crossbar switch with VOQs at each input, that connects each input to each output exactly $\frac{1}{N}$ -th of the time. For the moment, assume that the destination of packets is *uniform*; i.e. arriving packets are equally likely to be destined to any of the outputs.³ (Of course, real network traffic is nothing like this — but we will come to that shortly.) The fixed, equal-rate switch serves each VOQ at rate R/N , allowing us to model it as a GI/D/1 queue, with arrival rate $\lambda < R/N$ and service rate $\mu = R/N$. The system is stable (the queues will not grow without bound), and hence it guarantees 100% throughput.

Fact: If arrivals are uniform, a fixed, equal-rate switch, with virtual output queues, has a guaranteed throughput of 100%.

Of course, real network traffic is *not* uniform. But an extra load-balancing stage can spread out non-uniform traffic, making it sufficiently uniform to achieve 100% throughput. This is the basic idea of the two-stage load-balancing switch. A load-balancing device spreads packets evenly to all the inputs of a second, fixed, equal-rate switch.

Outline of proof: The load-balanced switch has 100% throughput for non-uniform arrivals for the following reason. Referring again to Figure 2, consider the arrival process, $a(t)$ (with N -by- N traffic matrix Λ) to the switch. This process is transformed by the sequence of permutations in the load-balancer, $\pi_1(t)$, into the arrival process to the second stage, $b(t) = \pi_1(t) \cdot a(t)$. The VOQs are served by the sequence of

³More precisely, assume that when a packet arrives, its destination is picked uniformly and at random from among the set of outputs, independently from packet to packet.




Figure 3: Possible system packaging for a 100 Tb/s router with 640 linecards arranged as 40 racks with 16 linecards per rack.

permutations in the switching stage, $\pi_2(t)$. If the inputs and outputs are not over-subscribed, then the long-term service opportunities exceed the number of arrivals, and hence the system achieves 100% throughput:

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T (b(t) - \pi_2(t)) = \frac{1}{N} e\Lambda - \frac{1}{N} e < 0,$$

where e is a matrix full of 1's.

In [16] the authors prove this more rigorously, and extend it to all sequences $\{a(t)\}$ that are stationary, stochastic and weakly mixing.

3. A 100TB/S ROUTER EXAMPLE

The load-balanced switch seems to be an appealing architecture for scalable routers than need performance guarantees. In what follows we will study the architecture in more detail. To focus our study, we will assume that we are designing a 100Tb/s Internet router that implements the requirements of RFC 1812 [18], arranged as 640 linecards operating at 160Gb/s (OC-3072). We pick 100Tb/s because it is challenging to design, is probably beyond the reach of a purely electronic implementation, but seems possible with optical links between racks of distributed linecards and switches. It is roughly two orders of magnitude larger than Internet routers currently deployed, and seems feasible to build using technology available in approximately three years time. We pick 160Gb/s for each linecard because 40Gb/s linecards are feasible now, and 160Gb/s is the next logical generation.

We will adopt some additional requirements in our design: The router must have a guaranteed 100% throughput for any pattern of arrivals, must not mis-sequence packets, and should operate correctly when populated with any number of linecards connected to any ports.

The router is assumed to occupy multiple racks, as shown in Figure 3, with up to 16 linecards per rack. Racks are connected by optical fibers and one or more racks of optical switches. In terms of optical technology, we will assume that it is possible to multiplex and demultiplex 64 WDM channels onto a single optical fiber, and that each channel can operate at up to 10Gb/s.

Each linecard will have three parts: An Input Block, an Output Block, and an Intermediate Input Block, shown in Figure 4. As is customary, arriving variable length packets will be segmented into fixed sized packets (sometimes called




Figure 4: Linecard block diagram

“cells”, though not necessarily equal to a 53-byte ATM cell), and then transferred to the eventual output, where they are reassembled into variable length packets again. We will call them fixed-size packets, or just “packets” for short. The Input Block performs address lookup, segments the variable length packet into one or more fixed length packets, and then forwards the packet to the switch. The Intermediate Input Block accepts packets from the switch and stores them in the appropriate VOQ. It takes packets from the head of each VOQ at rate R/N and sends them to the switch to be transferred to the output. Finally, the Output Block accepts packets from the switch, collects them together, reassembles them into variable length packets, and delivers them to the external line. Each linecard is connected to the external line with a bidirectional link at rate R , and to the switch with two bidirectional links at rate R .

Despite its scalability, the basic load-balanced switch has some problems that need to be solved before it meets our requirements. In the following sections we describe and then solve each problem in turn.

4. SWITCH RECONFIGURATIONS

4.1 Fixed Mesh

While the load-balanced switch has no centralized scheduler to configure the switch fabric, it still needs a switch fabric of size $N \times N$ that is reconfigured for each packet transfer (albeit in a deterministic, predetermined fashion). While optical switch fabrics that can reconfigure for each packet transfer offer huge capacity and almost zero power consumption, they can be slow to reconfigure (e.g. MEMS switches that typically take over 10ms to reconfigure) or are expensive (e.g. switches that use tunable lasers or receivers).⁴ Below, we’ll see how the switch fabric can be replaced by a fixed mesh of optical channels that don’t need reconfiguring.

Our first observation is that we can replace each fixed, equal-rate switch with N^2 fixed channels at rate R/N , as illustrated in Figure 5(a).

Our second observation is that we can replace the two switches with a single switch running twice as fast. In the basic switch, both switching stages connect every (input, output) pair at fixed rate R/N , and every packet traverses both switching stages. We replace the two meshes with a

⁴A glossary of the optical devices used in this paper appears in the Appendix.




Figure 5: Two ways in which the load-balanced switch can be implemented by a single fixed-rate uniform mesh. In both cases, two stages operating at rate R/N , as shown in (a), are replaced by one stage operating at $2R/N$, and every packet traverses the mesh twice. In (b), the mesh is implemented by N^2 fibers. In (c), the mesh is N^2 WDM channels interconnected by an AWGR. λ_w^i is transmitted on wavelength λ_w from input i and operates at rate $2R/N$.

single mesh that connects every (input, output) pair at rate $2R/N$, as shown in Figure 5(b). Every packet traverses the single switching stage twice; each time at rate R/N . This is possible because in a physical implementation, a linecard contains an input, an intermediate input and an output. When a packet has crossed the switch once, it is in an intermediate linecard; from there, it crosses the switch again to reach the output linecard.

The single fixed mesh architecture leads to a couple of interesting questions. The first question is: Does the mesh need to be uniform? i.e. so long as each linecard transmits and receives data at rate $2R$, does it matter how the data is spread across the intermediate linecards? Perhaps the first stage linecards could spread data over half, or a subset of the intermediate linecards. The answer is that if we don't know the traffic matrix, the mesh must be uniform. Otherwise, there is not a guaranteed aggregate rate of R available between any pair of linecards. The second question is: If it is possible to build a packet switch with 100% throughput that has no scheduler, no reconfigurable switch fabric, and buffer memories operating without speedup, where does the packet switching actually take place? It takes place at the input of the buffers in the intermediate linecards — the linecard decides which output the packet is destined to, and writes it to the correct VOQ.

4.2 When N is Large

A mesh of links works well for small values of N , but in practice, N^2 optical fibers or electrical links is impractical or

too expensive. For example, a 64-port router, with 40Gb/s lines (i.e. a capacity of 2.5Tb/s) would require 4,000 fibers or links, each carrying data at 1.25Gb/s. Instead, we can use wavelength division multiplexing to reduce the number of fibers, and increase the data-rate carried by each. This is illustrated in Figure 5(c). Instead of connecting to N fibers, each linecard multiplexes N WDM channels onto one fiber, with each channel operating at $2R/N$. The $N \times N$ arrayed waveguide grating router (AWGR) in the middle is a passive data-rate independent optical device that routes wavelength w at input i to output $[(i + w - 2) \bmod N] + 1$. The number of fibers is reduced to $2N$, at the cost of N wavelength multiplexers and demultiplexers, one on each linecard. The number of lasers is the same as before (N^2), with each of the N lasers on one linecard operating at a different, fixed wavelength. Currently, it is practical to use about 64 different WDM channels, and AWGRs have been built with more than 64 inputs and outputs [19]. If each laser can operate at 10Gb/s,⁵ this would enable routers to be built up to about 20Tb/s, arranged as 64-ports, each operating at $R = 320\text{Gb/s}$.

Our 100Tb/s router has too many linecards to connect directly to a single, central optical switch. A mesh of WDM channels connected to an AWGR (Figure 5(c)) would require 640 distinct wavelengths, which is beyond what is practical today. In fact a passive optical switch cannot interconnect 640 linecards. To do so inherently requires the switch to take data from each of the 640 linecards and spread it back over all 640 linecards in at least 640 distinct channels. We are not aware of any multiplexing scheme that can do this. If we try to use an active optical switch instead (such as a MEMS switch [21], electro-optic [22] or electro-holographic waveguides [23]), we must reconfigure it frequently (each time a packet is transferred), and we run into problems of scale. It does not seem practical to manufacture an active, reliable, frequently reconfigured 640-port switch from any of these technologies. And so we need to decompose the switch into multiple stages. Fortunately this is simple to do with a load-balanced switch. The switch does not need to be non-blocking; it just needs a path to connect each input to each output at a fixed rate.⁶ In Section 6, we will describe two different three-stage switch fabric architectures that decompose the switch fabric by arranging the linecards in groups (corresponding, in practice, to racks of linecards).

5. PACKET MIS-SEQUENCING

In the basic architecture, the load-balancer spreads packets without regard to their final destination, or when they will depart. If two packets arrive back to back at the same input, and are destined to the same output, they could be spread to different intermediate linecards, with different occupancies. It is possible that their departure order will be reversed. While mis-sequencing is allowed (and is common) in the Internet,⁷ network operators generally insist that routers do not mis-sequence packets belonging to the same applica-

⁵The modulation rate of lasers has been steadily increasing, but it is hard to directly modulate a laser faster because of wavelength instability and optical power ringing [20]. For example, 40Gb/s transceivers use external modulators.

⁶Compare this with trying to decompose a non-blocking crossbar into, say, a multiple stage Clos network.

⁷Internet RFC 1812 “Requirements for IP Version 4 Routers” [18] does not forbid mis-sequencing.

tion flow. In its current version, TCP does not perform well when packets arrive to the destination out of order because they can trigger un-necessary retransmissions.

There are two approaches to preventing mis-sequencing: To prevent packets from becoming mis-sequenced anywhere in the router [24]; or to bound the amount of mis-sequencing, and use a re-sequencing buffer in the third stage [25]. None of the schemes published to date would work in our 100Tb/s router. The schemes use schedulers that are hard to implement at these speeds, need jitter control buffers that require N writes to memory in one time slot [25], or require the communication of too much state information between the linecards [24].

5.1 Full Ordered Frames First

Instead we propose a scheme geared toward our 100Tb/s router. Full Ordered Frames First (FOFF) bounds the difference in lengths of the VOQs in the second stage, and then uses a re-sequencing buffer at the third stage.

FOFF runs independently on each linecard using information locally available. The input linecard keeps N FIFO queues — one for each output. When a packet arrives, it is placed at the tail of the FIFO corresponding to its eventual output. The basic idea is that, ideally, a FIFO is served only when it contains N or more packets. The first N packets are read from the FIFO, and each is sent to a different intermediate linecard. In this way, the packets are spread uniformly over the second stage.

More precisely, the algorithm for linecard i operates as follows:

1. Input i maintains N FIFO queues, $Q_1 \dots Q_N$. An arriving packet destined to output j is placed in Q_j .
2. Every N time-slots, the input selects a queue to serve for the next N time-slots. First, it picks round-robin from among the queues holding more than N packets. If there are no such outputs, then it picks round-robin from among the non-empty queues. Up to N packets from the same queue (and hence destined to the same output) are transferred to different intermediate linecards in the next N time-slots. A pointer keeps track of the last intermediate linecard that we sent a packet to for each flow; the next packet is always sent to the next intermediate linecard.

Clearly, if there is always at least one queue with N packets, the packets will be uniformly spread over the second-stage, and there will be no mis-sequencing. All the VOQs that receive packets belonging to a flow receive the same number of packets, so they will all face the same delay, and won't be mis-sequenced. Mis-sequencing arises only when no queue has N packets; but the amount of mis-sequencing is bounded, and is corrected in the third stage using a fixed length re-sequencing buffer.

5.2 Properties of FOFF

FOFF has the the following properties which are proved in [26].

- *Packets leave the switch in order.* FOFF bounds the amount of mis-sequencing inside the switch, and requires a re-sequencing buffer that holds at most $N^2 + 1$ packets.

- *No pathological traffic patterns.* The 100% throughput proof for the basic architecture relies on the traffic being stochastic and weakly mixing between inputs. While this might be a reasonable assumption for heavily aggregated backbone traffic, it is not guaranteed. In fact, it is easy to create a periodic adversarial traffic pattern that inverts the spreading sequence, and causes packets for one output to pile up at the same intermediate linecard. This can lead to a throughput of just R/N for each linecard.

FOFF prevents pathological traffic patterns by spreading a flow between an input and output evenly across the intermediate linecards. FOFF guarantees that the cumulative number of packets sent to each intermediate linecard for a given flow differs by at most one. This even spreading prevents a traffic pattern from concentrating packets to any individual intermediate linecard. As a result, FOFF generalizes the 100% throughput to *any* arriving traffic pattern; there are provably no adversarial traffic patterns that reduce throughput, and the switch has the same throughput as an ideal output-queued switch. In fact, the average packet delay through the switch is within a constant from that of an ideal output-queued switch.

- *FOFF is practical to implement.* Each stage requires N queues. The first and last stage hold at most $N^2 + 1$ packets per linecard (the second stage holds the congestion buffer, and its size is determined by the same factors as in a shared-memory work-conserving router). The FOFF scheme is decentralized, uses only local information, and does not require complex scheduling.
- *Priorities in FOFF are practical to implement.* It is simple to extend FOFF to support k priorities using $k \cdot N$ queues in each stage. These queues could be used to distinguish different service levels, or could correspond to sub-ports.

We now move on to solve the final problem with the load-balanced switch.

6. FLEXIBLE LINECARD PLACEMENT

Designing a router based on the load-balanced switch is made challenging by the need to support non-uniform placement of linecards. If all the linecards were always present and working, they could be simply interconnected by a uniform mesh of fibers or wavelengths as shown in Figure 5. But if some linecards are missing, or have failed, the switch fabric needs to be reconfigured so as to spread the traffic uniformly over the remaining linecards. To illustrate the problem, imagine that we remove all but two linecards from a load-balanced switch based on a uniform mesh. When all linecards were present, the input linecards spread data over N center-stage linecards, at a rate of $2R/N$ to each. With only two remaining linecards, each must spread over both linecards, increasing the rate to $2R/2 = R$. This means that the switch fabric must now be able to interconnect linecards over a range of rates from $2R/N$ to R , which is impractical (in our design example $R = 160Gb/s$).

The need to support an arbitrary number of linecards is a real problem for network operators who want the flexibility

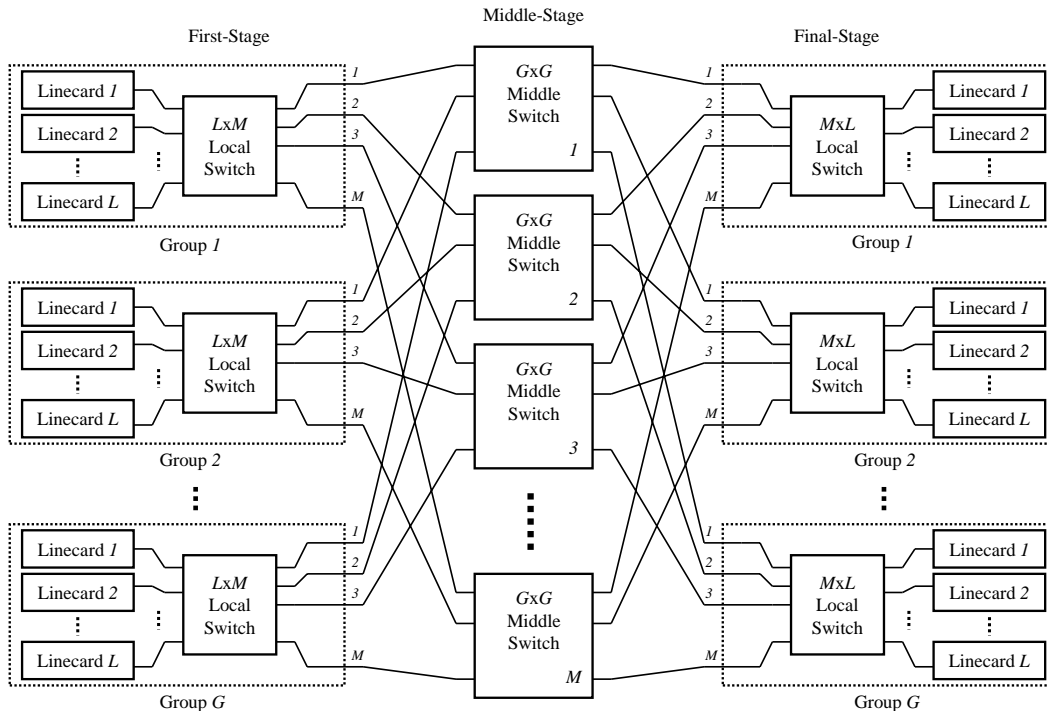


Figure 6: Partitioned switch fabric.

to add and remove linecards when needed. Linecards fail, are added and removed, so the set of operational linecards changes over time. For the router to work when linecards are connected to arbitrary ports, we need some kind of re-configurable switch to scatter the traffic uniformly over the linecards that are present. In what follows, we'll describe two architectures that accomplish this. As we'll see, it requires quite a lot of additional complexity over and above the simple single mesh.

6.1 Partitioned Switch

To create a 100Tb/s switch with 640 linecards, we need to partition the switch into multiple stages. Fortunately, partitioning a load-balanced switch is easier than partitioning a crossbar switch, since it does not need to be completely non-blocking in the conventional sense; it just needs to operate as a uniform fully-interconnected mesh.

To handle a very large number of linecards, the architecture is partitioned into G groups of L linecards. The groups are connected together by M different $G \times G$ middle stage switches. The middle stage switches are statically configured, changing only when a linecard is added, removed or fails. The linecards within a group are connected by a local switch (either optical or electrical) that can place the output of each linecard on any one of M output channels and can connect M input channels to any linecard in the group. Each of the M channels connects to a different middle stage switch, providing M paths between any pair of groups. This is shown in Figure 6. The number M depends on the uniformity of the linecards in the groups. For uniform linecard placement, the middle switches need to distribute the output from each group to all the other groups, which requires

G middle stage switches.⁸ In this simplified case $M = G$, i.e. there is one path between each pair of groups. Each group sends $1/G$ -th of its traffic over each path to a different middle-stage switch to create a uniform mesh. The first middle-stage switch statically connects input 1 to output 1, input 2 to output 2, and so on. Each successive switch rotates its configuration by one; for example, the second switch connects input 1 to output 2, input 2 to output 3, and so on. The path between each pair of groups is subdivided into L^2 streams; one for each pair of linecards in the two groups. The first-stage local switch uniformly spreads traffic, packet-by-packet, from each of its linecards over the path to another group; likewise, the final-stage local switch spreads the arriving traffic over all of the linecards in its group. The spreading is therefore hierarchical: The first-stage allows the linecards in a group to spread their outgoing packets over the G outputs; the middle-stage interconnects groups; and the final-stage spreads the incoming traffic from the G paths over the L linecards.

The uniform spreading is more difficult when linecards are not uniform, and the solution is to increase the number of paths M between the local switches.

THEOREM 1. *We need at most $M = L + G - 1$ static paths, where each path can support up to $2R$, to spread traffic uniformly over any set of $n \leq N = G \times L$ linecards that are present so that each pair of linecards are connected at rate $2R/n$.*

The theorem is proved formally in [26], but it is easy to show an example where this number of paths is needed. Consider the case when the first group has L line cards, but all

⁸Strictly speaking, this requires that $G \geq L$ if each channel is constrained to run no faster than $2R$.

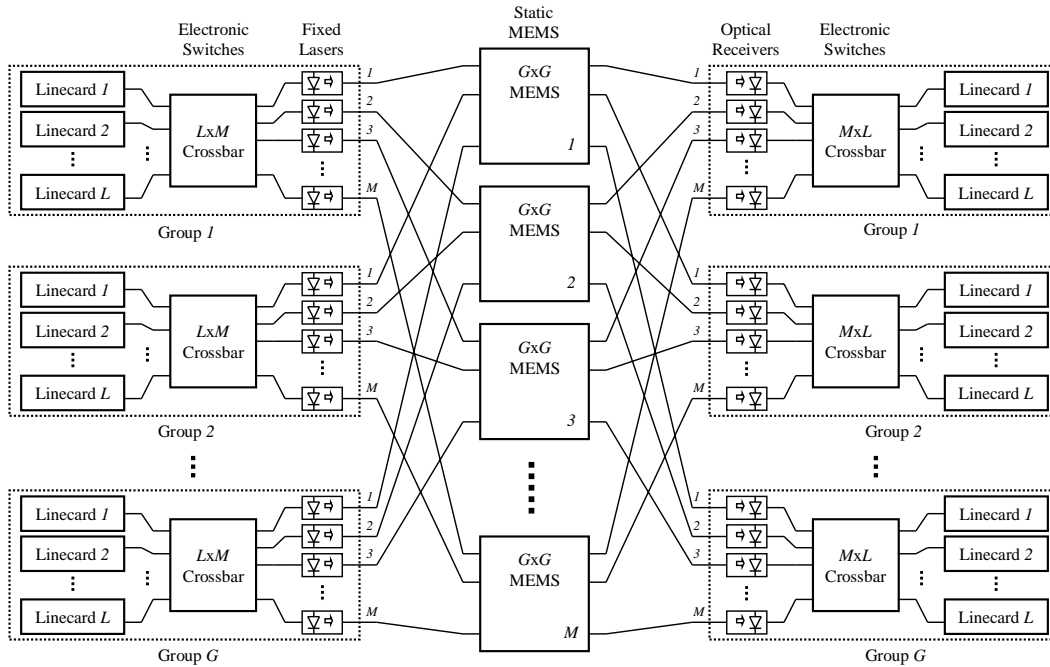


Figure 7: Hybrid optical and electrical switch fabric.

the other groups have just one linecard. A uniform spreading of data among the groups would not be correct. The first group needs to send and receive a larger fraction of the data. The simple way to handle this is to increase the number of paths, M , between groups by increasing the number of middle-stage switches, and by increasing the number of ports on the local switches. If we add an additional path for the each linecard that is out of balance, we can again use the middle-stage switches to spread the data. Since the maximum imbalance is $L - 1$, we need to have $M = L + G - 1$ paths through the middle switch. In the example given, the extra paths are routed to the first group (which is full), so now the data is distributed as desired, with $L/(L + G - 1)$ of the data arriving at the first group.

The remaining issue is that the path connections depend on the particular placement of the linecards in the groups, so they must be flexible and change when the configuration of the switch changes. There are two ways of building this flexibility. One uses MEMS devices as an optical patch-panel in conjunction with electrical crossbars, while the other uses multiple wavelengths, MEMS and optical couplers to create the switch.

6.2 Hybrid Electro-Optical Switch

The electro-optical switch is a straightforward implementation of the design described above. As before, the architecture is arranged as G groups of L linecards. In the center, M statically configured $G \times G$ MEMS switches interconnect the G groups. The MEMS switches are reconfigured only when a linecard is added or removed and provide the ability to create the needed paths to distribute the data to the linecards that are actually present. This is shown in Figure 7. Each group of linecards spreads packets over the MEMS switches using an $L \times M$ electronic crossbar. Each output of the electronic crossbar is connected to a different MEMS switch

over a dedicated fiber at a fixed wavelength (the lasers are not tunable). Packets from the MEMS switches are spread across the L linecards in a group by an $M \times L$ electronic crossbar.

We need an algorithm to configure the MEMS switches and schedule the crossbars. Because the switch has exactly the number of paths we need, and no more, the algorithm is quite complicated, and is beyond the scope of this paper. A description of the algorithm, and proof of the following theorem appears in [26].

THEOREM 2. *There is a polynomial-time algorithm that finds a static configuration for each MEMS switch, and a fixed-length sequence of permutations for the electronic crossbars to spread packets over the paths.*

6.3 Optical Switch

Building an optical switch that closely follows the electrical hybrid is difficult since we need to independently control both of the local switches. If we used an AWGR and wavelengths as the local switches, they could not be independently controlled. Instead, we modify the problem by allowing each linecard to have L optical outputs, where each optical output uses a tunable laser. Each of the $L \times L$ outputs from a group goes to a passive star coupler that combines it with the similar output from each of the other groups. This organization creates a large $(L \times G)$ number of paths between the linecards; the output fiber on the linecard selects which linecard in a group the data is destined for and the wavelength of the light selects one of the G groups. It might seem that this solution is expensive, since it multiplies the number of links by L . However, the high line rates ($2R = 320Gb/s$) will force the use of parallel optical channels in any architecture, so the cost in optical components is smaller than it might seem.

Once again, the need to deal with unbalanced groups

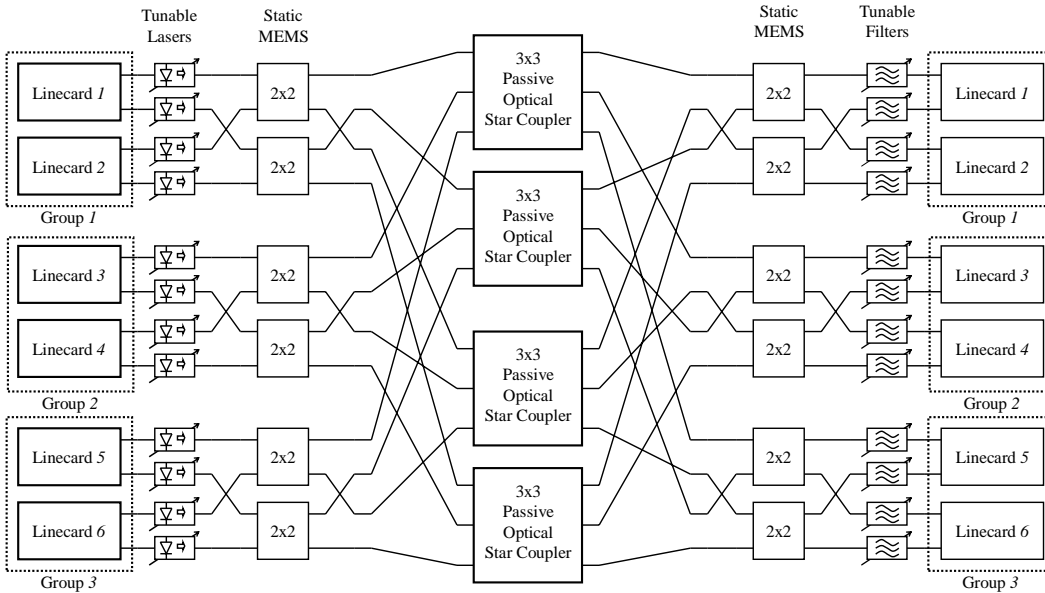


Figure 8: An optical switch fabric for $G = 3$ groups with $L = 2$ linecards per group.

makes the switch more complex than the uniform design. The large number of potential paths allows us to take a different approach to the problem in this case. Rather than dealing with the imbalance, we logically move the linecards into a set of balanced positions using MEMS devices and tunable filters. This organization is shown in Figure 8. Again, consider our example in which the first group is full, but all of the other groups have just one linecard. Since the star couplers broadcast all the data to all the groups, we can change the effective group a card sits in by tuning its input filter. In our example we would change all the linecards not in the first group to use the second wavelength, so that effectively all the single linecards are grouped together as a full second group. The MEMS are then used to move the position of these linecards so they do not occupy the same logical slot position. For example, the linecard in the second group will take the 1st logical slot position, the linecard in the third group will take the 2nd logical slot position, and so on. Together these rebalance the arrangement of linecards and allows the simple distribution algorithm to work.

7. PRACTICALITY OF 100TB/S ROUTER

It is worth asking: Can we build a 100Tb/s router using this architecture, and if so, could we package it in a way that network operators could deploy in their network?

We believe that it is possible to build the 100Tb/s hybrid electro-optical router in three years. The system could be packaged in multiple racks as shown in Figure 3, with $G = 40$ racks each containing $L = 16$ linecards, interconnected by $L + G - 1 = 55$ statically configured 40×40 MEMS switches.

To justify this, we will break the question down into a number of smaller questions. Our intention is to address the most salient issues that a system designer would consider when building such a system. Clearly our list cannot be complete. Different systems have different requirements, and must operate in different environments. With this caveat, we consider the following different aspects.

7.1 The Electronic Crossbars

In the description of the hybrid electro-optical switch, we assumed that one electronic crossbar interconnects a group of linecards, each at rate $2R = 320Gb/s$. This is too fast for a single crossbar, but we can use bit-slicing. We'll assume W crossbar slices, where W is chosen to make the serial link data-rate achievable. For example, with $W = 32$, the serial links operate at a more practical 10Gb/s. Each slice would be a 16×55 crossbar operating at 10Gb/s. This is less than the capacity of crossbars that have already been reported [27].

Figure 9 shows L linecards in a group connected to W crossbar slices, each operating at rate $2R/W$. As before, the outputs of the crossbar slices are connected to lasers. But now, the lasers attached to each slice operate at a different, fixed wavelength, and data from all the slices to the same MEMS switch are multiplexed onto a single fiber. As before, the group is connected to the MEMS switches with M fibers. If a packet is sent on the n -th crossbar slice, it will be delivered to the n -th crossbar slice of the receiving group. Apart from the use of slices to make a parallel datapath, the operation is the same as before.

Each slice would connect to $M = 55$ lasers or optical receivers. This is probably the most technically challenging, and interesting, design problem for this architecture. One option is to connect the crossbars to external optical modules, but might lead to prohibitively high power consumption in the electronic serial links. We could reduce power if we could directly connect the optical components to the crossbar chips. The direct attachment (or "solder bumping") of III-V opto-electronic devices onto silicon has been demonstrated [28], but is not yet a mature, manufacturable technology, and is an area of continued research and exploration by us, and others. Another option is to attach optical modulators rather than lasers. An external, high powered continuous wave laser source could illuminate an array of




Figure 9: Bit-sliced crossbars for hybrid optical and electrical switch. (a) represents the transmitting side of the switch. (b) represents the receiving side of the switch.

integrated modulators on the crossbar switch. The array of modulators modulate the optical signal and couple it to an outgoing fiber [29].

7.2 Packaging 100Tb/s of MEMS Switches

We can say with confidence that the power consumption of the optical switch fabric will not limit the router's capacity. Our architecture assumes that a large number of MEMS switches are packaged centrally. Because they are statically configured, MEMS switches consume almost no power, and all 100Tb/s of switching can be easily packaged in one rack using commercially available MEMS switches today. Compare this with a 100Tb/s electronic crossbar switch, that connects to the linecards using optical fibers. Using today's serial link technology, the electronic serial links alone would consume approximately 8kW (assume 400mW and 10Gb/s per bidirectional serial link). The crossbar function would take at least 100 chips, requiring multiple extra serial links between them; hence the power would be much higher. Furthermore, the switch needs to terminate over 20,000 optical channels operating at 10Gb/s. Today, with commercially available optical modules, this would consume tens of kilowatts, would be unreliable and prohibitively expensive.

7.3 Fault-Tolerance

The load-balanced architecture is inherently fault-tolerant. First, because it has no centralized scheduler, there is no electrical central point of failure for the router. The only centrally shared devices are the statically configured MEMS switches, which can be protected by extra fibers from each linecard rack, and spare MEMS switches. Second, the failure of one linecard will not make the whole system fail; the MEMS switches are reconfigured to spread data over the correctly functioning linecards. Third, the crossbars in each group can be protected by an additional crossbar slice.

7.4 Building 160Gb/s Linecards

We assume that the address lookup, header processing and buffering on the linecards are all electronic. Header processing will be possible at 160Gb/s using electronic tech-

nology available within three years. At 160Gb/s, a new minimum length 40-byte packet can arrive every 2ns, which can be processed quite easily by a pipeline in dedicated hardware. 40Gb/s linecards are already commercially available, and anticipated reductions in geometries and increases in clock speeds will make 160Gb/s possible within three years.

Address lookups are challenging at this speed, but it will be feasible within three years to perform pipelined lookups every 2ns for IPv4 longest-prefix matching. For example, one could use 24Mbytes of 2ns SRAM (Static RAM)⁹ and the brute force lookup algorithm in [30] that completes one lookup per memory reference in a pipelined implementation.

The biggest challenge is simply writing and reading packets from buffer memory at 160Gb/s. Router linecards contain 250ms or more of buffering so that TCP will behave well when the router is a bottleneck, which requires the use of DRAM (dynamic RAM). Currently, the random access time of DRAMs is 40ns (the duration of twenty minimum length packets at 160Gb/s!), and historically DRAMs have increased in random access speed by only 10% every 18 months. We have solved this problem in other work by designing a packet buffer using commercial memory devices, but with the speed of SRAM and the density of DRAM [31]. This technique makes it possible to build buffers for 160Gb/s linecards.

7.5 Packaging 16 Linecards in a Rack

Network operators frequently complain about the power consumption of 10Gb/s and 40Gb/s linecards today (200W per linecard is common). If a 160Gb/s linecard consumes more power than a 40Gb/s linecard today, then it will be difficult to package 16 linecards in one rack ($16 \times 200 = 3.2kW$). If improvements in technology don't solve this problem over time, we can put fewer linecards in each rack, so

⁹Today, the largest commercial SRAM is 4Mbytes with an access time of 4ns, which suggests what is feasible for on-chip SRAM. Moore's Law suggests that in three years 16Mbyte SRAMs will be available with a pipelined access time below 2ns. So 24Mbytes can be spread across two physical devices.

long as $G \times L \geq 640$. For example, we could halve the number of linecards per rack and double the number of groups. This comes at the expense of more MEMS switches ($M \geq L + G - 1$).

8. CONCLUSION

Our main conclusion is that we achieved what we set out to do: To solve the problems that allow us to design a 100Tb/s load-balanced router, with guaranteed 100% throughput under all traffic conditions. We believe that the electro-optic router we described, including switch fabric and linecards, can be built using technology available within three years, and fit within the power constraints of network operators.

To achieve our capacity requirement, optics are necessary. A 100Tb/s router needs to use multiple racks regardless of the architecture, because packets need to be processed and stored electronically, and the power consumption of the electronics is too much for a single rack. Optical links are needed to interconnect the multiple racks, but we don't have to use an optical switch fabric. A distributed switch fabric could spread the switching function over all the racks. The key problem is that a distributed switch fabric could not guarantee the system throughput. But if throughput guarantees are not needed, a distributed switch fabric (such as a torus or hypercube) might be a reasonable alternative.

Providing throughput guarantees would normally require a complicated centralized scheduler to configure the switch fabric. The main advantages of the load-balanced switch are that: (1) It allows us to eliminate the centralized scheduler without sacrificing the throughput guarantees, and (2) It allows us to use a switch fabric that doesn't need to be frequently reconfigured. More precisely, in the hybrid electro-optic switch fabric only the lower-capacity local switches in each group need to be reconfigured frequently — the high capacity MEMS switches change only when linecards are added or removed. The switch fabric is essentially transparent, consumes no power, and eliminates power-hungry conversions between the electrical and optical domain. All the central switching can be packaged in a single rack.

In future, as optical technology matures, it will be possible to replace the hybrid electro-optical switch with an all-optical fabric. This has the potential to reduce power further by eliminating many electronic crossbars and serial links.

9. REFERENCES

[1] A. M. Odlyzko, "The current state and likely evolution of the Internet," *Proc. Globecom'99*, pp. 1869-1875, 1999.

[2] A. M. Odlyzko, "Comments on the Larry Roberts and Caspian Networks study of Internet traffic growth," *The Cook Report on the Internet*, pp. 12-15, Dec. 2001.

[3] Pat Blake, "Resource," *Telephony*, Feb. 2001, available at http://telephonyonline.com/ar/telecom_resource/index.htm

[4] Router Capacity: Raw data
<http://www.stanford.edu/~nickm/RouterCapacity.xls>

[5] N. McKeown, C. Calamvokis, S.-T. Chuang, "A 2.5Tb/s LCS switch core," *Hot Chips XIII*, Aug. 2001.

[6] K. Y. Yun, "A terabit multiservice switch," *IEEE Micro*, Vol. 21, pp. 58-70, Jan.-Feb. 2001.

[7] Alcatel, "Alcatel's packet-based digital cross-connect switch application," July 2002, available at <http://www.alcatel.com>.

[8] PMC-Sierra Inc., "Tiny-Tera one chip set," April 2000, available at <http://www.pmc-sierra.com/pressRoom/chess.html>.

[9] Juniper Networks, "The essential core: Juniper Networks T640 Internet routing node with matrix technology," April 2002, available at <http://www.juniper.net/solutions/literature/solutionbriefs/351006.pdf>.

[10] W. J. Dally, "Architecture of the Avici terabit switch router," *Proc. Hot Interconnects XIII*, Aug. 1998.

[11] N. McKeown, A. Mekkittikul, V. Anantharam and J. Walrand, "Achieving 100% throughput in an input-queued switch," *IEEE Trans. on Comm.*, Vol.47, No.8, Aug. 1999.

[12] J.G. Dai and B. Prabhakar, "The throughput of data switches with and without speedup," *Proc. of the IEEE INFOCOM*, Vol. 2, pp. 556-564, Tel Aviv, Israel, March 2000.

[13] Y. Tamir and H.C. Chi, "Symmetric crossbar arbiters for VLSI communication switches," *IEEE Trans. on Parallel and Distributed Systems*, vol. 4, no. 1, pp. 13-27, 1993.

[14] N. McKeown, M. Izzard, A. Mekkittikul, B. Ellersick, and M. Horowitz, "The Tiny Tera: A packet switch core," *Proc. of Hot Interconnects V*, Aug. 1996.

[15] T.E. Anderson, S.S. Owicki, J.B. Saxe, and C.P. Thacker, "High speed switch scheduling for local area networks," *ACM Trans. on Computer Systems*, Vol. 11, No. 4, pp. 319-352, Nov. 1993.

[16] C.-S. Chang, D.-S. Lee and Y.-S. Jou, "Load balanced Birkhoff-von Neumann switches, Part I: one-stage buffering," *Computer Comm.*, Vol. 25, pp. 611-622, 2002.

[17] L.G. Valiant and G.J. Brebner, "Universal schemes for parallel communication," *Proc. of the 13th ACM Symposium on Theory of Computation*, pp. 263-277, 1981.

[18] F. Baker, "Requirements for IP Version 4 Routers", RFC 1812, June 1995, available at <http://www.faqs.org/rfcs/rfc1812.html>.

[19] P. Bernasconi, C. Doerr, C. Dragone, M. Capuzzo, E. Laskowski and A. Paunescu, "Large N x N waveguide grating routers", *Journal of Lightwave Technology*, Vol. 18, No. 7, pp. 985-991, July 2000.

[20] K. Sato, "Semiconductor light sources for 40-Gb/s transmission systems," *Journal of Lightwave Technology*, vol.

[21] R. Ryf *et al.*, "1296-port MEMS transparent optical crossconnect with 2.07 petabit/s switch capacity," *Optical Fiber Comm. Conf. and Exhibit (OFC) '01*, Vol. 4, pp. PD28 -P1-3, 2001.

[22] G.I. Papadimitriou, C. Papazoglou, and A.S. Pomportsis, "Optical switching: switch fabrics, techniques, and architectures," *Journal of Lightwave Technology*, vol. 21, no. 2, pp. 384-405, Feb. 2003.

[23] B. Pesach *et al.*, "Free-space optical cross-connect switch by use of electroholography," *Applied Optics*, Vol. 39, No. 5, pp. 746-758, Feb. 2000.

- [24] I. Keslassy and N. McKeown, "Maintaining packet order in two-stage switches," *Proc. of the IEEE Infocom*, June 2002.
- [25] C.-S. Chang, D.-S. Lee and C.-M. Lien, "Load balanced Birkhoff-von Neumann switches, Part II: multi-stage buffering," *Computer Comm.*, Vol. 25, pp. 623-634, 2002.
- [26] I. Keslassy, S-T. Chuang, and N. McKeown, "Architectures and algorithms for a load-balanced switch," *Stanford University HPNG Technical Report - TR03-HPNG-061501*, Stanford, CA, June 2003.
- [27] B. Dally, Velio Communications, "A single chip terabit switch," *Hot Chips XIII*, Stanford, Aug. 2001.
- [28] A. Krishnamoorthy, "Optoelectronics flip-chip-bonded to CMOS VLSI circuits," *IMAPS Advanced Technology Workshop on Next Generation IC and Package Design*, Solvang, CA, July 1999.
- [29] A. L. Lentine *et al.*, "Arrays of optoelectronic switching nodes comprised of flip-chip-bonded MQW modulators and detectors on silicon CMOS circuitry," *IEEE Photonics Technology Letters*, Vol. 8, pp. 221-223, Feb. 1996.
- [30] P. Gupta, S. Lin, and N. McKeown, "Routing lookups in hardware at memory access speeds," in *Proc. of IEEE INFOCOM '98*, pp. 1240-1247, San Francisco, CA, April 1998.
- [31] S. Iyer, R. R. Kompella, and N. McKeown, "Designing buffers for router line cards," *Stanford University HPNG Technical Report - TR02-HPNG-031001*, Stanford, CA, Mar. 2002.
- [32] J. Gripp *et al.*, "Demonstration of a 1.2 Tb/s optical packet switch fabric" *27th European Conf. on Optical Comm. (ECOC) '01*, Vol. 6, Amsterdam, The Netherlands, Oct. 2001.
- [33] N. Kikuchi, "Error-free signal selection and high-speed channel switching by monolithically integrated 64-channel WDM channel selector," *Electronics Letters*, Vol. 38, No. 15, pp. 823-824, July 2002.

APPENDIX

A. GLOSSARY OF OPTICAL COMPONENTS

- *MEMS Switches* - optical equivalent of a crossbar using micromirrors to reflect optical beams from inputs to outputs and are transparent to wavelength and data-rate. Typical reconfiguration times are 1-10ms [21, 22].
- *Tunable Lasers* - lasers that can transmit light at different wavelengths. Tuning times of 10ns have been demonstrated using commercial devices [32].
- *Tunable Filters* - optical detectors that detect channel information at selected wavelengths. Devices have been demonstrated with tuning times of 10ns [33].
- *Passive Optical Star Coupler* - light from N incoming fibers is combined and broadcast to N outgoing fibers with an intrinsic $1/N$ power loss.