An extended comparison of slotted and unslotted deflection routing.

Thierry Chich, Pierre Fraigniaud

Laboratoire de l'Informatique du Parallélisme

Ecole Normale Supérieure de Lyon
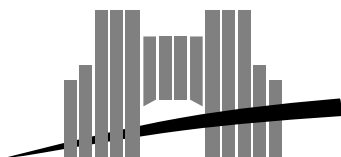Unité de recherche associée au CNRS n°1398

An extended comparison of slotted and unslotted deflection routing

Thierry Chich and
Pierre Fraigniaud

March 20, 1997

An extended comparison of slotted and unslotted deflection routing

Thierry Chich and
Pierre Fraigniaud

March 20, 1997

Abstract

In this paper, we have experimentally compared synchronized versus asynchronized all-optical deflection networks. The originality of our approach is first that we have included a model of bursty traffic: it is simulated by a bi-Poissonian emission. Second, we have compared four routing modes: *synchronous* mode, *partially synchronous* mode, and *asynchronous* modes with *fixed* and *bounded* size packets. All modes were considered under the same emission protocols. More precisely, we have run the several experiments with a careful attention to the several time scalings related to these different modes. Our experiments mainly show that the natural decrease of the performances of the asynchronous mode, compared to the synchronous mode, can be balanced in a significant way by the use of a sophisticated routing algorithm. Moreover, we have also shown that asynchronous routing is not very sensitive to bursty traffic. These results, and the fact that asynchronous networks are easier to design, and cheaper to build than synchronous networks, show the practical interest of asynchronous deflection routing.

Keywords: All-optical networks, deflection routing.

Résumé

Nous avons expérimentalement comparé l'effet du synchronisme ou de l'asynchronisme pour le routage par déflexion dans des réseaux tout-optiques. L'originalité de notre approche est d'avoir intégré un modèle de trafic sporadique, sous la forme d'émissions bipoissonniennes. Elle est aussi dans le fait que nous avons comparé quatre modes de routage : mode *synchrone*, mode *partiellement synchrone*, et modes *asynchrones* avec messages de *taille fixe* ou de *taille bornée*. Tout ces modes ont été examinés à partir du même protocole d'émission. Plus précisément, nous avons veillé attentivement au respect des différentes échelles de temps que nous avons dues considérer. Nos expérimentations montrent essentiellement que la dégénérescence naturelle des performances du mode asynchrone comparées à celles du mode synchrone peut être contrebalancée par un mode de routage asynchrone plus astucieux. De plus, nous montrons que les réseaux asynchrones sont moins sensibles aux trafics sporadiques. Ces résultats, accentués par le fait que les réseaux asynchrones sont plus simples et moins chers à construire, montrent l'intérêt pratique des réseaux à déflexion asynchrone.

Mots-clés: Réseaux transparents, routage par déflexion.

An extended comparison of slotted and unslotted deflection routing[*]

Thierry Chich and Pierre Fraigniaud

March 20, 1997

1 Introduction

As their names clearly indicate, all-optical networks exploit photonic technologies for the implementation of both switching, and transmission [11]. When some data that are stored in an electronic format must be sent through the network, the network-application interface constructs a *message* whose *payload* (i.e. message's data) is in optic format. The message is then transmitted through the network toward its destination without any conversion of the payload from optics to electronics in order to eliminate the electronic bottleneck, and to exploit the enormous capacity of optics. When the message arrives at its destination, its payload is again converted from optics to electronics by the network-application interface.

Two large classes of networks can be considered, depending on the way messages are routed from their sources to their destinations: *single-hop*, and *multi-hop* [1, 3]. Both have their advantages, and drawbacks. Single-hop routing provides end-to-end transparent channels. The current implementations of single-hop routing use wavelength division multiplexing (WDM) systems [15]. Single-hop routing is non (locally) adaptive. Multi-hop routing can also be implemented by WDM systems. It allows adaptive choices of the routes, up to the price of reintroducing, to some extend, the electronic bottleneck. Indeed, taking fast routing decisions requires a computational power which is difficult to obtain optically using the current technology. Hence, multi-hop routing requires that the *header* of the message must be converted in an electronic format at every hop.

The high bandwidth of optics allows us to deal with enhanced networking characteristics as, in particular, an integration of the services at a town level. For this purpose, *Metropolitan Area Networks* (MAN) [13] have been proposed. MAN are often opposed to LAN (local Area Networks), and WAN (Wide Area Networks) in term of both size, and structure. A LAN is generally of modest size, and is connected as a ring or a star. A WAN is generally defined as a large loosely coupled network. A MAN is of intermediate size, and can be viewed as a tightly coupled network. The relatively large size of a MAN does not allow global control of the network, and multi-hop routing is an adapted switching mode for that kind of networks. MAN is supposed to integrate many different services as RTC for telephone and video, and X25 or IP for data transfers between LANs. Such multi-media applications generate very different data flows in the network. Hence, it must support sporadic traffics [10, 14].

Routing in electronic networks (of any kind) is highly based on the use of a large number of buffers. This is not possible on all-optical networks for which deflection routing [4, 7] (also called hot-potato routing) is frequently prefered. Indeed, optical buffers are difficult to build with the current technology.

Routing messages among nodes of an interconnection network requires to satisfy many constraints as small latency (for interactive applications), regular traffic (for constant-bandwidth applications), ordered delivery (for on-line applications), and, of course, fault-tolerance. Therefore, many parameters must be carefully chosen. The *topology* models the way the nodes are connected together. The need of fast routing decisions, and major technological constraints, limit the degree of the network. Networks as meshes (wrapped around or not) [12], or *de Bruijn* [16] are possible candidates. *Synchronous networks* provide efficient routing since a routers can consider all its inputs globally in order to optimize its routing decision. On the other hand, *asynchronous networks* are more easy to build. An application does not generate a single message in general, but a sequence of (possibly ordered) messages. Also, it might be useful to decompose a large message in sub-messages of smaller size. Actually, the network-application interface produces *packets*. A packet is the smallest entity subject to routing decisions. A packet can be of fixed size, or just of bounded size. This choice is mainly forced by the type of network-synchronization. The *routing strategy* defines the ways source-destination routes are constructed. For instance, XY-routing is the most common rule to route in meshes. A *shortest-path routing* is a routing strategy which insures that, in absence of other traffic, any message will always follow a shortest path from its source toward its destination. *Bufferisation* is possible in all-optical networks, but it requires a lot of optical fibers, and possibly amplifiers. Also, a message cannot be bufferized for an arbitrarily large period of time. The main consequence of the buffer limitation is that the routing decision must be taken on-line, and very fast.

In this paper, we will study the influence of the level of synchronism on routing in deflection networks. A similar study has already be done in [6]. However, our approach and our conclusion differ in many aspects. For instance, although we noticed an important degradation of the throughput of unslotted deflection networks, we did not observe a situation where severe congestion occurs. Several slight differences between [6] and this paper may explain this divergence (insertion policy, topology, etc.). The paper is organized as follows. Section 2 accurately describes our model of all-optical networks, and our simulation protocol. Then Section 3 presents experimental measures on the throughput of the network, on the load of the network, and on the number of lost packets. Section 4 deals with local parameters, as the length of the queues, and the time spent by a packet in a queue before being sent through the network. Section 4 presents also results on the number of time a packet can be deflected. Finally, Section 5 contains some concluding remarks. In particular, we conclude that asynchronous routing deserves to be considered with attention. Indeed, we have noticed that the performance degradations of asynchronous routing versus synchronous routing are mainly due to algorithmic problems rather than intrinsic problems related to asynchronism.

2 Model and experiments

2.1 Routing in all-optical networks

We have considered one of the most famous types of topology: the *bidirectional Manhattan street network* without wrapped around links, that is the symmetrically oriented mesh (see Figure 1(a)). Each node is connected to the its router by a single bidirectional channel. We did not add optical buffers (i.e. fiber loop of a given length) to routers since all-optical routing should try to avoid the use of buffers. Each router is therefore supposed to be a 5×5 crossbar (see Figure 1(b)).

In our model, a message is composed of its payload and its header. The payload contains the data (files, images, sounds, etc.), and the header includes useful information for the routing function (destination label, packet number, source label, etc.). The payload generally consists of a
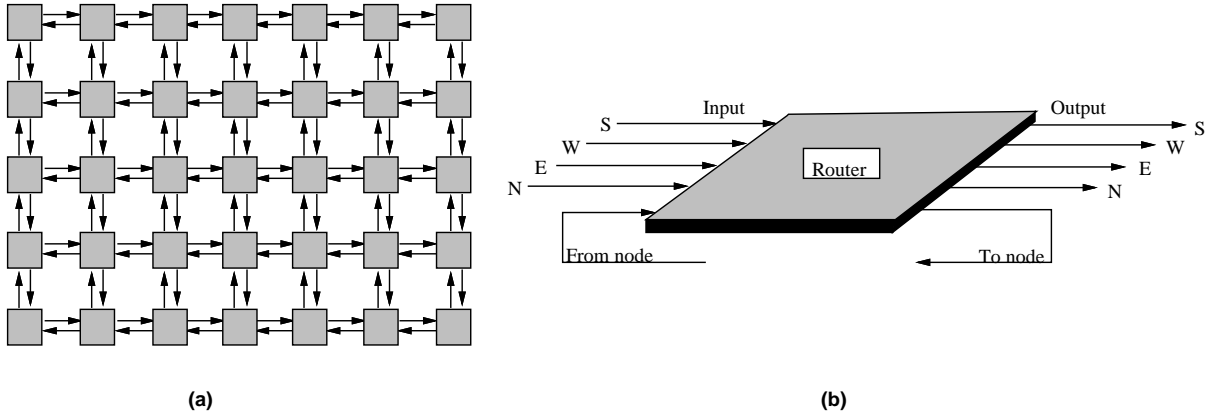
2

(a)

(b)

Figure 1: (a) A 5×7 bidirectional Manhattan street network, and (b) one of its routers (i.e. a 5×5 crossbar).
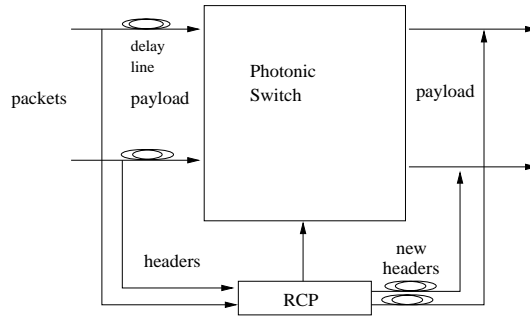
Figure 2: Deflection routing in a 2×2 crossbar: the routing control processor in denoted by RCP.

large amount of data which must be kept always in optical form so that it circulates at the photonic rate. Headers and payloads can be transmitted on different wavelength. The bandwidth allocated to headers is very limited because the size of the header is limited by the electronic bottleneck. For example, the *Mixed Rate* technique was reported at 100 Mb/s for the header and 700 Mb/s for the payload [5].

When a message arrives at a given router (see Figure 2), its header is converted in electronic format, and it is decoded by the router which takes the routing decision. Once the decision has been taken according to some simple rules, that is when a single output port has been selected, the router connects the input port to the output port so that the payload can cut through the router. The payload is just slightly delayed in a loop while the routing control processor (RCP) is computing the route. The RCP can generate a new header for the message.

We use a shortest path routing. More precisely, a message is always routed on a shortest path from its source to its destination in absence of other traffic. If many possible output ports correspond to shortest paths, then one of then is chosen uniformly at random. If all the output ports corresponding to shortest paths are busy, i.e., if they are already used by other messages, then one of the other output ports is chosen uniformly at random, and the message is deflected from its
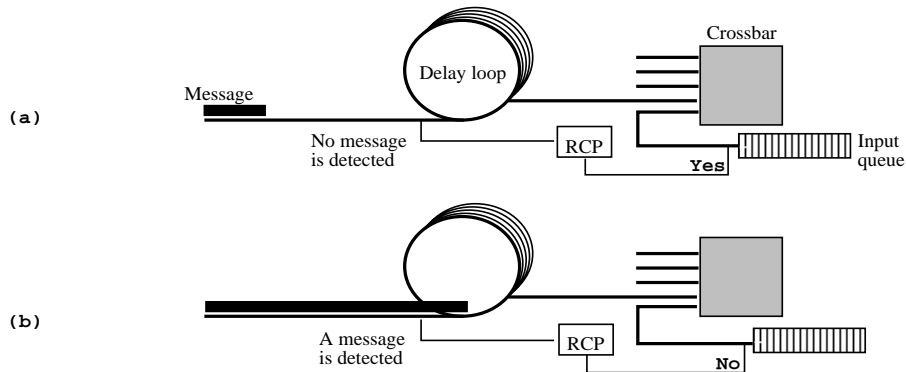
Figure 3: In an asynchronous network, a packet is inserted only if there is enough time before the receipt of another message (case (a)), otherwise it waits in the FIFO input queue (case (b)).

shortest path. A comparison between store-and-forward and deflection routing in Shuffle networks shows that deflection routing offers relatively good behavior compared to store-and-forward [2]. Deflection routing requires much less bufferisation than store-and-forward routing.

In deflection routing, messages are assumed to be inserted in the network under some conditions (otherwise $n + 1$ messages could compete for only n output ports). Therefore, storage facilities are placed at all inputs of the network to store packets which cannot be inserted at the time they are created. In our model, these buffers are running in FIFO mode. If the network is synchronous, one can insert a new packet in the network if one of the four input ports is empty. Indeed, four simultaneous inputs can always be routed. In asynchronous network, the problem is a bit more tricky since there is no condition on the time interval during which an input port will be empty. A possible solution [8] is to check "in advance" on each optic fiber to figure out whether or not a message can be inserted (see Figure 3). This can be done by adding a fiber loop to each input channel in order to delay the arrival of packets. This is possible as soon as packets are of bounded length. This is typically the case of IP packets.

Another important difference between synchronous and asynchronous networks is that, in synchronous networks, one can handle globally the several packets arriving in a router at the same time. This fact allow each router to allocate optimally the packets to their routes, and, therefore, one can minimize locally the number of deflections. We call *locally–matched* deflection routing this kind of strategy.

2.2 Traffic generation

We have fixed the size of the mesh at 12×12, with input queue of size 100 requests at each node. The bandwidth of the links is supposed to be $10Gb/s$, and each link is supposed to have a length of $2km$. A packet is supposed to be a maximum size $1Ko$, that is $1\mu s$, that is also $200m$.

2.2.1 Time scaling

We have considered two time scaling in order to separate the behavior of the network from the behavior of the applications using the network:

1. Simulation tick, or network tick; and

4

2. Processor tick, or emission tick.

The interval between two simulation ticks is $0.1\mu s$. At each tick, we consider possible emission of packets at each node, and we route packets in the network. The traffic demand is simulated as follows. At each node, the decision to introduce or not a packet in the input queue is taken according to a probabilistic law, and destinations are chosen randomly uniformly among the other nodes. Each node follows the same law. Processor tick is fixed at $20ns$. At this speed, and at each processor, the emission follows a Bernoulli law (when a processor sends, it sends exactly one packet). This Bernoulli law is in turn simulated by a Binomial law at the simulation tick. We denote by t_n (resp. t_p), the tick of the network (resp. of the processor).

Most of our experimental results are presented as function of the load offered to the network. The offered load is expressed in packet per *slot*. In synchronous networks, the slot is the packet size. In asynchronous networks, the slot is an abstract measure expressing the maximum size of a packet. The slot is denoted by t_s. We have fixed the slot at $t_s = 1\mu s$. Hence, to get a fixed offered load $\mathcal{L}$, we have forced the parameter of the Bernoulli law $B(\lambda)$ followed by the processors emissions to be $\lambda = \frac{\mathcal{L}}{t_s/t_p}$. Then the emission law of a network is $B(\lambda, t_n/t_p)$. This protocol produces the same emission law for both synchronous and asynchronous networks. Note that it would not have been the case if we would have followed the naive approach consisting of setting $t_n = t_s$ for synchronous simulation.

2.2.2 Sporadic traffic

We mainly consider two different emission laws for two different kinds of experiments: Poissonian traffic, and sporadic traffic. In the Poissonian traffic, every processor follows the same Bernoulli law. This is the most commonly studied traffic in the literature.

In order to simulate a sporadic traffic, we have used an emission law denoted by $S(\mathcal{L}_g, p, \mathcal{L}_b, p')$ [9]. In this case, each node is in two possible states called *ground* and *bursty*. These states alternate according to two probabilities p and p'. From the ground state, the probability to enter the bursty state is p. From the bursty state, the probability to enter the ground state is p'. In the ground state, the emission law is Poissonian, thence it is similar to the one previously described. In the bursty state, we allow processors to send a large number of packets within one slot (such packets will be stored in the input queue). When a processor is in the bursty state, its offered load is of average $\mathcal{L}_b > 1$ (to be compared with the global offered load in the Poissonian traffic which is always strictly less than 1).

We have considered, as in [14], that bursty traffics are mainly caused by ftp-data-like applications. Moreover, whatever is the load of the network, a burst offers the same characteristic. Thus, we decide to set $\mathcal{L}_b = cst$, independently from the global load $\mathcal{L}$. For the same reasons, the probability p' to get out of a bursty application is not related to the global load, and thus it is set as a constant. The ground emission rate $\mathcal{L}_g$ is defined as a linear function of the offered load of the network $\mathcal{L}$. Indeed, the ground traffic is induced by telnet-like connections [14] whose number grows linearly with the number of running applications. We have set $\mathcal{L}_g = \mathcal{L}_g^{\max}\mathcal{L}$, where $\mathcal{L}_g^{\max}$ is the ground traffic saturating the network (to be fixed later according to the experimental results on the Poissonian traffic). For a given offered load, the probability p is fixed to $p'\frac{\mathcal{L}(\mathcal{L}_g^{\max}-1)}{\mathcal{L}-\mathcal{L}_b}$ so that the mean of the law $S(\mathcal{L}_g, p, \mathcal{L}_b, p')$ is $\mathcal{L}$. Thence, in our sporadic model, an increase of the load will be induced by an increase of the frequency at which we enter in the bursty state.

We have fixed (somewhat arbitrarily) the value of $\mathcal{L}_b$ at 5 (smaller bursts would not be significant, and larger bursts would saturate the input queues). We also set $p' = 0.02$. This value was fixed according to the one of $\mathcal{L}_b$. It implies that a burst will fill up the input queue with 25 packets

in the average, that is with 25% of the size of the queue. Finally, we set $\mathcal{L}_g^{\max} = 0.3$ (see Figure 5). One can see on Figure 4 that Poissonian and sporadic traffics are indeed very different.
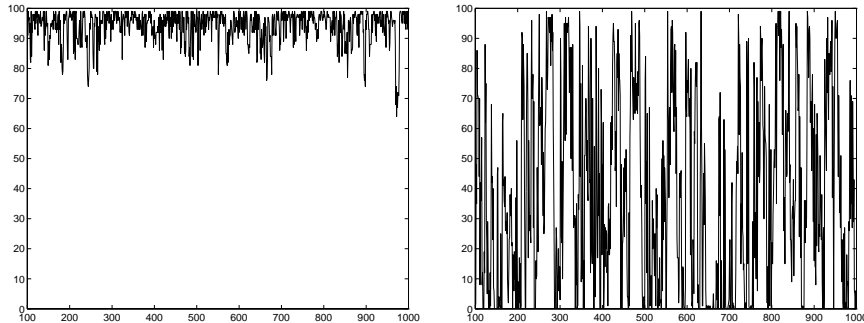
Figure 4: Poissonian traffic versus sporadic traffic: number of packets in a queue as function of the time.

2.2.3   Message of variable length

Asynchronous networks support messages of different lengths. We have considered that the length of messages L follows a bimodal law polarized at (1) the length of the acknowledgment packets, and (2) the length of the slot. We have fixed the minimum size of a packet at $200ns$ (the size of a header). We recall that a slot is $1\mu s$. We set $P(L = 200ns) = 0.3$, and $P(L = 1\mu s) = 0.4$. The other message lengths are chosen as multiple of $0.1\mu s$, uniformly in the interval $[300, 900]$. The average of such a law is $642ns$.

2.3   Experimental measures

All measurements are performed at the steady state. We have measured the *throughput* of the network as a function of the input demand. More precisely, we have counted the number of packets which arrive at destination at each slot in average, divided by the number of nodes (that is 144). This number is in $[0, 1]$ for synchronous networks. The input demand is the average number of packets each node sends in average at each step. Note that since input queues are of bounded size, packets can be lost when the network approaches the saturation. The number of *loss packets* is inversely proportional to the throughput. We have also considered the *load* of the network. It is the average number of packets simultaneously in the network divided by the number of links. Again this number is in $[0, 1]$ for synchronous networks.

We have also considered local measurements. More precisely, we have created a specific traffic, called *spy traffic*, between two given nodes. In our experiments, node $(2, 2)$ sends packets to node $(9, 9)$ according to a Poissonian law of mean 0.01 (i.e., at a low rate). We have reported the distribution of the number of times spy packets are deflected. We have also measured the distribution of the number of packets in the queue of $(2, 2)$ when a spy packet enters the queue, and the distribution of the number of simulation ticks a spy packet has to wait in the queue of $(2, 2)$ before being sent to the network.

3 Global measurements

In this section, we present experimental results on the throughput of the network, on the load of the network, and on the number of loss packets. The two first subsections deal with synchronous networks, whereas the last subsection deals with asynchronous networks. Subsection 3.2.1 presents two intermediate models, somewhat in between synchronous and asynchronous networks. These models are introduced for the purpose of a clear comparison between these two types of networks.

3.1 Synchronous routing under Poissonian traffic

Figure 5 presents the well know behavior of synchronous routing under Poissonian traffic. Figure 5(a) shows the two states of the network: a linear increase of the throughput until the network gets saturated. When the network saturates, the throughput becomes constant, and the number of loss packets increases (whereas no packets are lost for a low offered load). One can check that the network starts to saturate for an offered load larger than 0.3 (either by comparison with the diagonal line, or by looking at the number of loss packets).

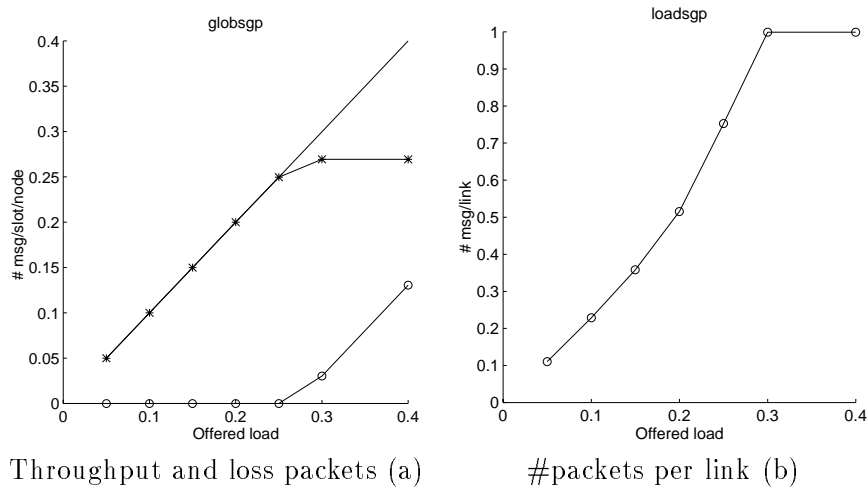Throughput and loss packets (a)   #packets per link (b)

Figure 5: Synchronous routing under Poissonian traffic

Figure 5(b) presents the average number of messages per link. Again, the form of the results is not surprising. When the offered load increases, the number of messages per link increases super linearly. This is due to the interactions between the network load on one hand, and the number of messages deflections on the other hand. When the network reaches the saturation, all the bandwidth of the network is used. This is always the case for synchronous networks.

3.2 Synchronous routing under sporadic traffic

Figure 6(a) and (b) show the influence of a sporadic traffic on synchronous routing.

Figure 6(a) shows that, when the network is not yet saturated, the number of loss packets is larger under the sporadic traffic than under the Poissonian traffic. This is due to the large standard deviation of the bi-Poissonian traffic. When the network is saturated, the routings of the two types of traffic offer the same behavior. As one can check on Figure 6(b), the loss of packets under a
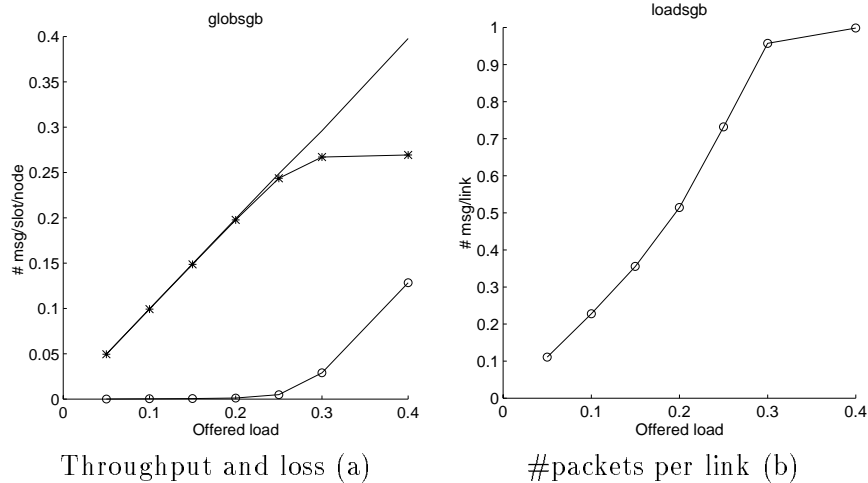
Throughput and loss (a) #packets per link (b)

Figure 6: Synchronous routing under sporadic traffic

sporadic traffic imply that the links saturate for a larger offered load. The number of loss packets is the major difference between Poissonian and sporadic traffic. However, for a same number of packets inside the network, the behavior of these packets is roughly the same for both traffics.

3.2.1 Two intermediate routing modes

As we said before, one can use synchronism to improve the routing algorithm by handling simultaneously the several packets entering a router at the same time. This global handling of the input packets allow us to minimize the number of deflections for the considered group of packets. Another advantage of the synchronism is that one can introduce more packets in the network than for asynchronous networks with fixed size packets. Indeed, there is often not enough space between two consecutive packets to insert an entering packet in asynchronous networks with fixed size packets. In order to study separately the influences of these two good properties of synchronous routing over asynchronous routing, we have introduced two intermediate models. The first model is called *partially synchronous*: the network is synchronous but packets arriving simultaneously in a router are handled in any order without optimization of the number of deflections. The second model is called *fixed size asynchronous*: the network is asynchronous but packets are of fixed size.

Figure 7 presents the behavior of the partially synchronous model under a Poissonian traffic. One can notice a large degradation of the performances in comparison with synchronous routing. For instance, the network get saturated for a much smaller offered load (roughly 0.2 rather than 0.3). Regarding asynchronous routing with fixed size packets, one can check on Figure 8 that the network saturates before the bandwidth is totally used. Actually, the bandwidth cannot be totally used in average in an asynchronous mode. A simultaneous observation of Figures 5, 7, and 8 makes clear the way performances decrease from synchronous routing to asynchronous routing with fixed size packets.
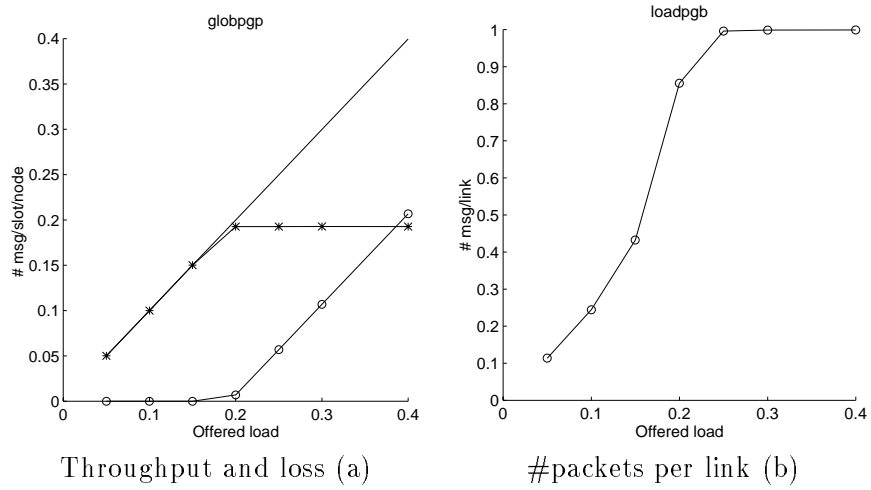
8

Throughput and loss (a)   #packets per link (b)

Figure 7: Partially synchronous routing under the Poissonian traffic

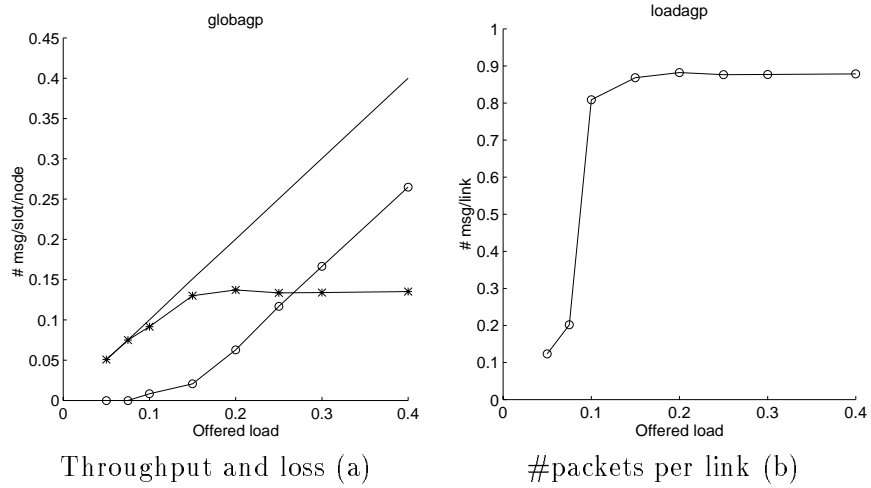Throughput and loss (a)   #packets per link (b)

Figure 8: Fixed size packet asynchronous routing under the Poissonian traffic

9

3.3 Asynchronous routing

This section focuses on a totally asynchronous routing, that is with packets of arbitrary size. Compared to Figure 8, Figure 9 shows that allowing packets of different size improve the performances. This is straightforward because, for a same number of packets, the asynchronous mode requires much less bandwidth than the synchronous mode (smaller average size packet).

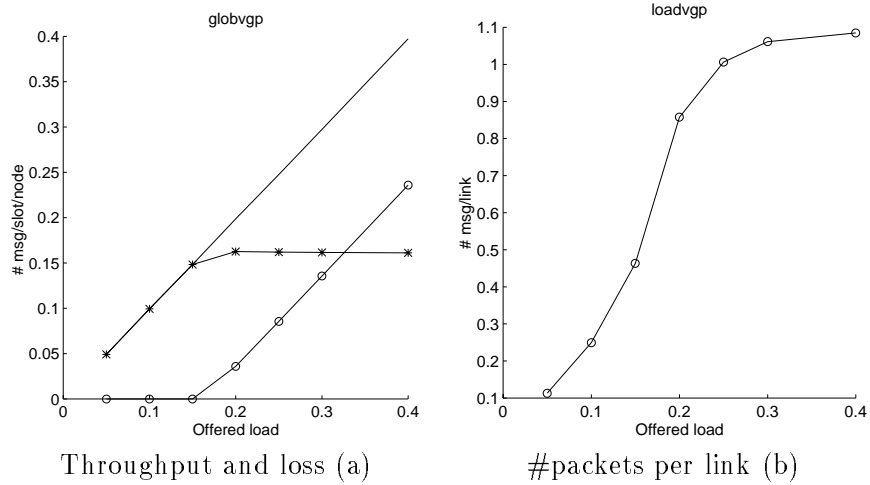Throughput and loss (a)   #packets per link (b)

Figure 9: Asynchronous routing under Poissonian traffic

Figure 9 and 10 show that the behavior of the network looks the same for both Poissonian and sporadic traffics in asynchronous routing. Note that this phenomenon cannot be caused by a different occupation of the queues. Indeed, we have fixed the maximum number of packets allowed in a queue, and not the maximum capacity of a queue. Therefore it is an intrinsic property of asynchronous routing.
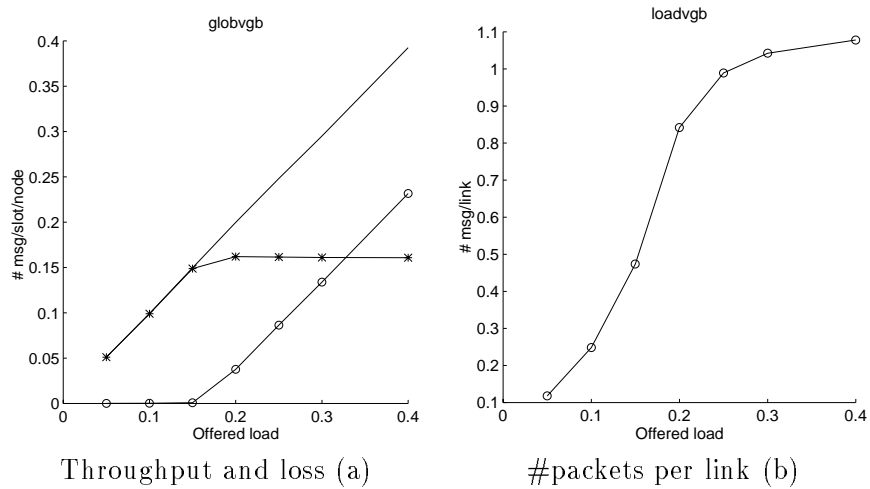

Throughput and loss (a)   #packets per link (b)

Figure 10: Asynchronous routing under sporadic traffic

4 Local measurements

In this section, we report measures on the spy traffic described in Section 2.3. In particular, we have measured the latency (message delay) of the spy traffic. We have also considered the behavior of the queue of the source of this spy traffic. There will be seven different distributions that will be presented in the forthcoming figures, one for each value of offered load: 0.05, 0.10, 0.15, 0.20, 0.25, 0.3, and 0.4. To figure out the one-to-one correspondence between the curves and the offered load, one can look at the y-axis. Values on this axis are (naturally) ordered in a decreasing order of the offered load. To draw the behavior of the queue, we had to use two scales for the presentation of the results. One is dedicated to low offered loads, and the other is dedicated to the saturated state. Indeed, moving from one state to the other produces many changes in the behavior of the queue sizes.

4.1 Message delays

We have measured the number of deflections of spy messages sent by node $(2,2)$ to node $(9,9)$. For a sake of uniformity, we have normalized the results as a function of the number of received messages.

4.1.1 Synchronous routing

Figure 11 shows that, under low traffic condition (that is for an offered load at most 0.2), the number of deflections is relatively small, but for a load of 0.25. When the traffic load increases, the shape of the distributions changes: the median increases a bit, and both maximum and standard deviation of the distribution strongly increase.
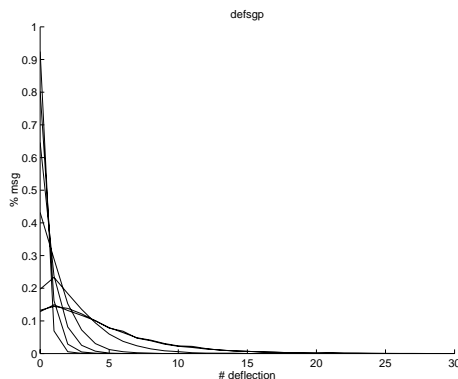

Figure 11: Distribution of the number of deflections for synchronous routing under the Poissonian traffic

There is no significant difference between sporadic and Poissonian traffic when looking at the distribution of the delays (see Figure 12). The tiny improvements under the sporadic traffic come from the smaller average number of messages per link under this mode. This confirms the fact that, as we already pointed out in Section 3.2, the internal behavior of a deflection network is roughly independent of the traffic nature.
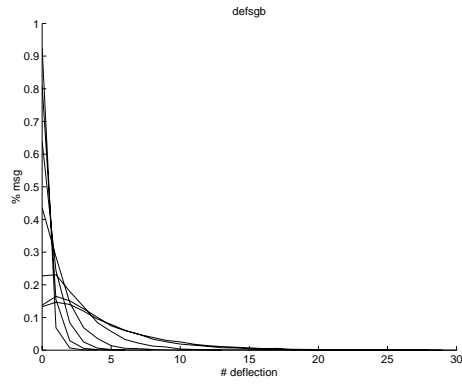
Figure 12: Distribution of the number of deflections for synchronous routing under the sporadic traffic

4.1.2 Asynchronous routing

As for the synchronous case, we noted that there is no big difference between Poissonian and sporadic traffic when looking at the dispersion of the number of deflections. However synchronous and asynchronous routing present totally different behavior. Figure 13 shows the important degradation of the performances when asynchronous routing is used, although we will see that this phenomenon is not really du to the asynchronism!
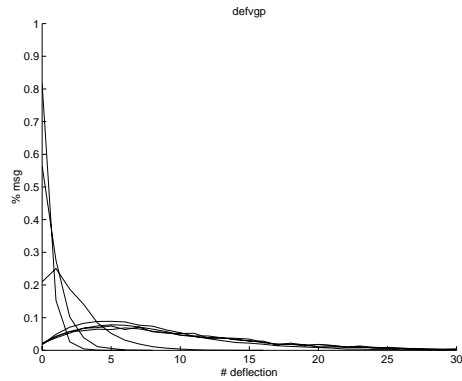

Figure 13: Distribution of the number of deflections for asynchronous routing under the Poissonian traffic

Figure 13 and 14 offer roughly the same shape (same median, same standard deviation, etc.). This shows that the performance degradation of asynchronous routing is mainly due to the difficulty of minimizing the number of deflections locally. This is more an algorithmic problem than an intrinsic problem of the asynchronism.
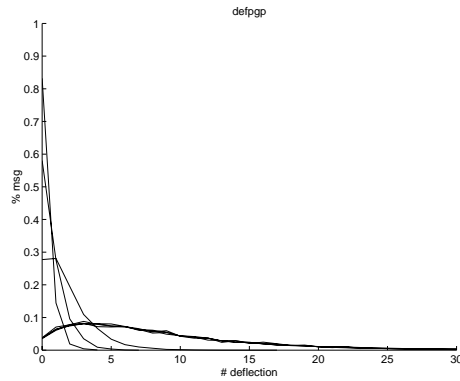
12

Figure 14: Distribution of the number of deflections for partially synchronous routing under the Poissonian traffic

4.2 Behavior of the queues

Two parameters characterize the behavior of the queues: (1) the number of packets currently in the queue when a given packet in introduced in that queue, and (2) the time a given packet waits in the queue before being introduced in the network. As opposed to the network internal behavior, these two parameters will clearly show the difference between a Poissonian traffic and a sporadic traffic.

4.2.1 Synchronous mode

As Figure 15 clearly shows, the number of packets that a given packet finds in the queue of a slightly loaded network is quite small in the average. On the contrary, when the network is saturated, the queues are almost always completely full. Even more, Figure 16 shows that the average time every packet waits in a queue of a saturated network is about 350 slots, that is 3.5 times the size of the queues. On the contrary, when a packet enters a queue of a slightly loaded network, it waits a time roughly equal to the number of packets waiting in the queue.

The behavior of the queues under a sporadic traffic is completely different when the network is saturated (see Figure 17). Indeed, the repartition of the number of packets in the queue is almost uniform. This is induced by the large standard deviation of a sporadic traffic. This is not in contradiction with the fact that, when the network get saturated, both Poissonian and sporadic traffics lost about the same number of packets. Indeed, remember that the spy traffic is Poissonian. Although it does not appear clearly on the left hand side of Figure 17, a large number of spy messages find the queue filled up by an arbitrary large number of packets even when the network is far to be saturated.

4.2.2 Asynchronous routing

Figure 18 shows the waiting time in a queue under a Poissonian traffic for asynchronous routing. We point out one major difference between this figure and Figure 16. For the maximum load, the waiting time in the asynchronous mode is about three time the waiting time in the synchronous mode.
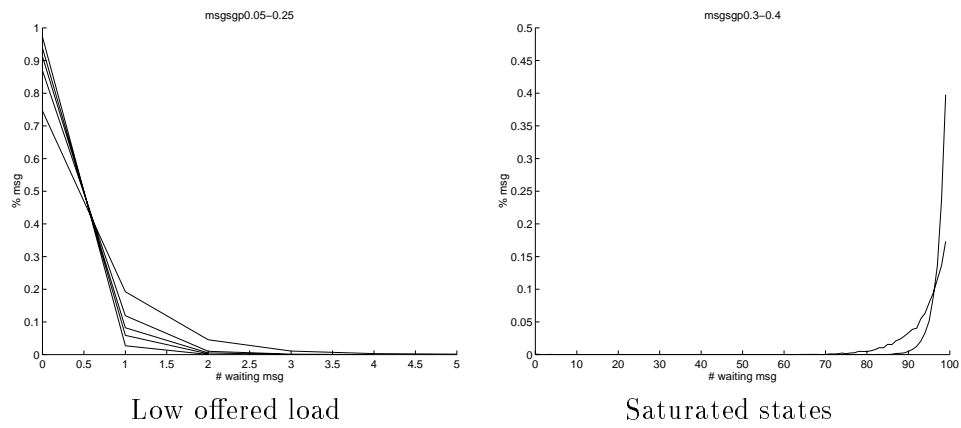
13

Low offered load   Saturated states

Figure 15: Number of packets when entering the queue under the Poissonian traffic for synchronous routing

Low offered load   Saturated states
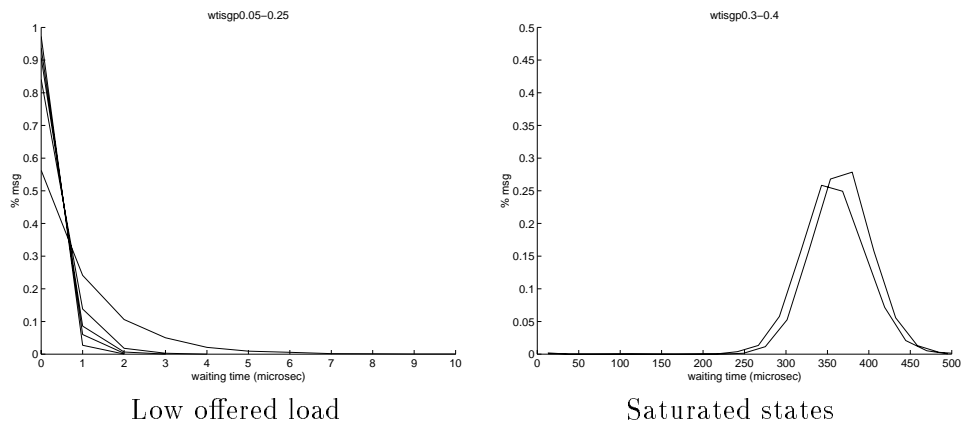
Figure 16: Waiting time in the queue under the Poissonian traffic for synchronous routing

Low offered load   Saturated states
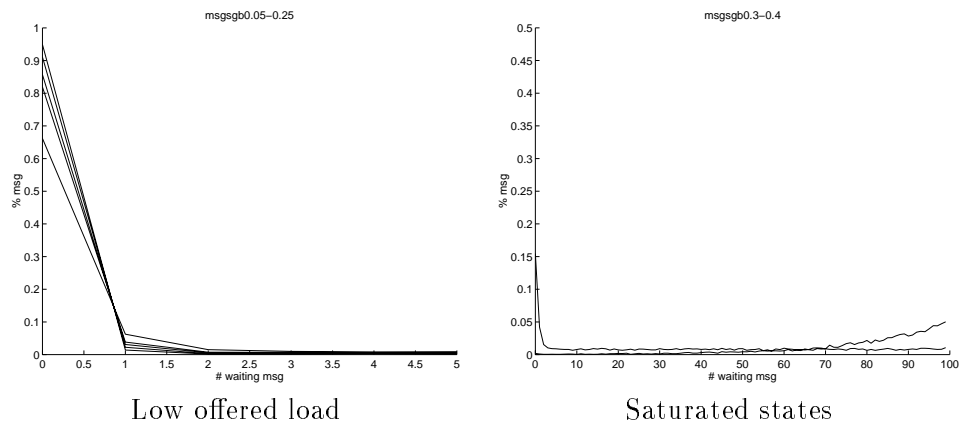
Figure 17: Number of packets when entering the queue under the sporadic traffic for synchronous routing
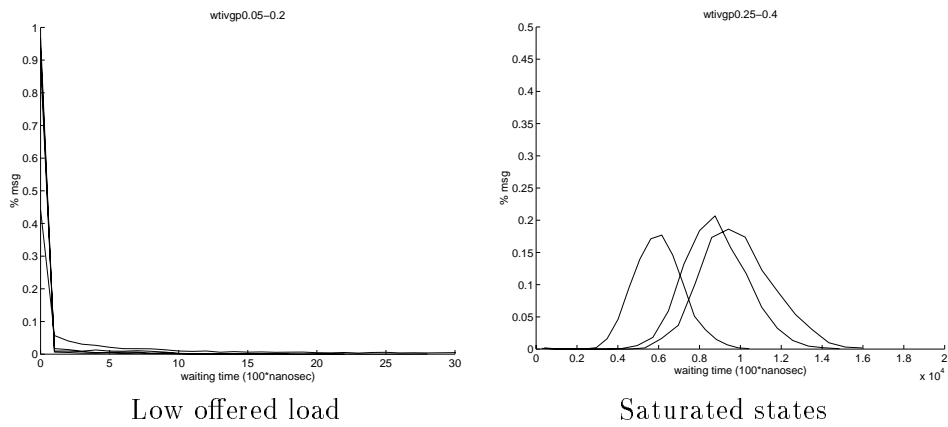

Low offered load   Saturated states

Figure 18: Waiting time in the queue under the Poissonian traffic for asynchronous routing

5 Conclusion

In this paper, we have experimented the performances of deflection routing in synchronous and asynchronous networks. We have shown that, as it could have been easily guessed, synchronous routing performs better that asynchronous routing (the former gets saturated for an higher offered load than the latter, the number of deflection of messages is smaller in synchronous networks than in asynchronous networks, etc.). However, we have also shown that the main reason of this difference of behavior is the fact that synchronism helps to minimize locally the number of deflections (for instance using a locally–matched deflection routing). Therefore, the question that naturally arises is to figure out whether it is possible to replace the greedy deflection routing in asynchronous networks by some kind of locally–matched algorithm.

It is technically possible to implement locally–matched deflection routing in asynchronous networks. Indeed, one can use the fact that, as explained in the introduction (see Figure 3), one can know a bit in advance the possible arrivals in a router. However, the underlying algorithmic problem is a bit more tricky since it turns out to be an on-line scheduling problem. We think that it is of a major interest to investigate deeper this question since synchronism is technically difficult to implement, and expensive to realize.

References

[1] ACAMPORA, A. S., AND KAROL, M. J. An Overview of Lightwave Packet Networks. *IEEE Network* (Jan. 1989), 29–41.

[2] ACAMPORA, A. S., AND SHAH, S. I. Multihop Lightwave Networks : A Comparison of Store–and–Forward and Hot–Popato Routing. *IEEE Trans. on Communication 40*, 6 (June 1992), 1082–1089.

[3] BARANSEL, C., DOBOSIEWICZ, W., AND GBURZYNSKI, P. Routing in Multihop Packet Switching : Gb/s Challenge. *IEEE Network* (May 1995), 38–61.

[4] BLUMENTHAL, D. J., FEUERSTEIN, R., AND SAUER, J. R. First Experimentation of Multihop All–Optical Packet Switching. *IEEE Photonics technology letters 6*, 3 (Mar. 1994), 457–460.

[5] BLUMENTHAL, D. J., PRUCNAL, P. R., AND SAUER, J. R. Photonic Packet Switches : Architectures and Experimental Implementations. In *Proc. of the IEEE* (Nov. 1994), vol. 82(11), IEEE, pp. 1650–1667.

[6] BORNOGOVO, F., FRATTA, L., AND BANNISTER, J. Unslotted Deflexion Routing in All–Optical Networks. *IEEE* (1993), 119–125.

[7] CHOUDHURY, A. K., AND LI, V. O. Performance Analysis of Deflexion Routing in the msn. In *Int. Conf. on Communication* (1991), IEEE, pp. 1659–1665.

[8] CLÉROT, F. Private communication.

[9] COLLANGE, D. Reclassification des modèles de trafic. Tech. rep., CNET, France, Jan. 1995.

[10] LELAND, W., TAQQU, M., WILLINGER, W., AND WILSON, D. On the Self–Similar Nature of Ethernet Traffic. *IEEE/ACM Transactions on Networking 2*, 1 (Feb. 1994), 1–15.

[11] MARSAN, M. A., BIANCO, A., LEONARDI, E., AND NERI, F. Topologies for Wavelength–Routing All–Optical Networks. *IEEE/ACM Transaction on Networking 1*, 5 (1993), 534–546.

[12] MAXENCHUK, N. F. Routing in the MSN. *IEEE Trans. on Communication 35*, 5 (1987).

[13] MOLLENAUER, J. F. Standards for Metropolitan Area Networks. *IEEE Communications Magazine 26*, 4 (1988), 15–19.

[14] PAXSON, V., AND FLOYD, S. Wide Area Traffic : The Failure of Poisson Modelling. *IEEE/ACM Transactions on Networking 3*, 3 (June 1995), 226–244.

[15] RAMASWANI, R., AND SIVARAJAN, K. N. Routing and Wavelentgh Assignement in All–Optical Networks. *IEEE/ACM Transaction on Networking 3*, 3 (Oct. 1995), 489–500.

[16] SIVARAJAN, KUMAR N.AND RAMASWANI, R. Lightwave Networks based on de Bruijn Graphs. *IEEE/ACM Transaction on Networking 2*, 1 (Feb. 1994), 70–79.