3

Service Guarantee in Deflection Networks

M. Ajmone Marsan, E. Leonardi, F. Neri, and C. Pistritto
Dipartimento di Elettronica, Politecnico di Torino
Corso Duca degli Abruzzi 24, 10129 Torino, Italy
Fax: +39-11-5644099, e-mail: neri@polito.it

Abstract

An approach is described for the provision of guaranteed-quality connection-oriented services to isochronous traffic (as well as connectionless services to bursty traffic) in mesh networks with deflection routing. The approach is based on the mapping of the mesh topology on a multiring topology, whose linearity simplifies the management of slots devoted to isochronous services. Numerical results obtained by simulation show that, when the network load due to bursty traffic is low, the introduction of isochronous connection-oriented services has a marginal impact on the performance of connectionless services for bursty traffic. On the contrary, when the network load due to bursty traffic is high, the performance of connectionless services for bursty traffic suffers significantly for the introduction of isochronous connection-oriented services; however, the performance of the latter is insensitive of the network load level.

Keywords

Mesh topology, service guarantee, deflection routing, Manhattan Street Network.

1 Introduction

Topologies are of paramount importance in the determination of the performances of telecommunication networks in terms of many different quality parameters (throughput, delay, reliability, etc.). This is why so many research efforts have always been devoted to the identification of topologies with desirable characteristics, and to the topological optimization of networks.

Actually, at least two different topological levels can normally be found in a telecommunications network: at a higher level we find the *virtual* topology (also often called the *logical* topology), and at a lower level we find the *physical* topology. Normally, the optimization is performed at the level of the virtual topology, but when this virtual topology has to be implemented, it must be *mapped* onto a convenient physical topology that accounts for a number of logistic constraints. Of course, a virtual topology can be implemented on top of a physical topology provided that a mapping $\mathcal{M} : T_v \rightarrow T_p$ is possible from the virtual topology T_v onto the physical topology T_p.

A typical example is provided by the field of optical networks, where a number of topological optimizations and analyses have been performed at the virtual topology level, assuming a very simple physical topology, such as an optical star or ring, and exploiting wavelength and space division to perform the mapping of the virtual topology onto the physical topology [1].

In general, the mapping of a topology T_1 onto another topology T_2 requires the possibility of emulating some characteristics of T_1 (the desirable ones) within T_2 (possibly losing some

T. Hasegawa et al. (eds.), *Local and Metropolitan Communication Systems*

marginal characteristics of T_1). This mapping is independent of the fact that T_1 is a virtual topology and T_2 is a physical topology, and is possible also for two virtual topologies T_{v1} and T_{v2}.

In this paper we exploit this mapping between virtual topologies in order to provide guaranteed-quality connection-oriented services to isochronous traffic (as well as connectionless services to bursty traffic) in square-grid mesh networks with deflection routing.

The original proposal of the use of deflection routing in networks with square-grid topologies is due to Maxemchuk [2], who proposed and studied the Manhattan Street Network (MSN). Many other studies later followed, considering unidirectional as well as bidirectional MSN (BMSN) [3, 4, 5, 6].

Several authors have recently shown that the performances achievable with MSN and BMSN are by far superior to those that can be obtained with linear topologies, such as rings and busses, from the point of view of numerous quality indexes [7, 8]. This suggests that square-grid topologies (or mesh topologies in general) are possible candidates for the design of future very high-speed local and metropolitan area networks (LANs and MANs).

With respect to ATM high-speed networks, that are receiving today most of the attentions of the networking community, MSN and BMSN have similar space diversity advantages and scalability properties, due to the meshed (non linear) topology. While ATM only provides for a connection oriented information transfer, and special attention must be paid when connectionless services are required, MSN and BMSN were conceived for the datagram service exploited by most of the data applications in use today. Moreover, the regularity of the topology (same number of input and output links), together with the datagram paradigm and deflection routing algorithms, allows interesting approaches to the problem of congestion control, for which no final solutions have been found until now in the ATM world. MSN and BMSN can thus be considered as an interesting arena were possible answers to the questions raised by high-performance networking for a large population of users can be experimented.

However, a disadvantage of MSN (and in general of deflection routing networks) with respect to typical MAN architectures is in the fact that it is difficult to provide guaranteed-quality conection-oriented services to isochronous traffic, and to insure fairness in the access to the network of the different users. Both of these characteristics are much easier to obtain from linear topology networks.

The proposal that is presented in this paper achieves both the high performance of regular mesh topologies and the service guarantee characteristics of linear networks by mapping square-grid virtual topologies onto multiring virtual topologies, i.e., topologies comprising a set of parallel rings.

In the sequel of this paper we first describe the considered square-grid topologies, and their mapping onto the multiring; then we illustrate how connection-oriented services can be implemented in the resulting virtual topology. Numerical results are given to statically characterize the topologies, and to illustrate the performances that can be obtained for bursty and isochronous traffic.

Previous proposals for the introduction of circuit services in deflection networks with square grid topologies were presented in [9, 10], but they are not meant to provide the same service quality as our approach, that guarantees truly periodic access to slots that suffer no deflection, so that the delivery delay is constant, and no reassembly procedures are necessary for isochronous traffic, since slots always arrive in sequence.

2 Three Square-Grid Regular Mesh Topologies

Three different square-grid regular mesh topologies are considered in this paper; each topology comprises $M = N^2$ nodes arranged in N rows and N columns with N nodes each. A node can be identified through the pair (i, j) whose elements are the row ($i = 0, \ldots, N-1$) and column ($j = 0, \ldots, N-1$) ordinal numbers.

The three topologies are such that exactly two point-to-point unidirectional links arrive at and depart from each node. This means that the total number of links in the three topologies is $2M$, and that the macroscopic node structure is identical for the three topologies, and is also identical to the one in standard MAN configurations, such as FDDI [11], and DQDB [12]: two transmitters and two receivers within each node process the signals on four point-to-point channels.

Each topology comprises $2N$ rings, each made of a sequence of N point-to-point links. Each ring reaches exactly N nodes, that lie in the same column or row. Thus, N horizontal rings cover the rows, and N vertical rings cover the columns in the topology, and node (i, j) is connected to the i-th horizontal ring and to the j-th vertical ring.

The three topologies are characterized by the directions in which the different rings run.

The first topology that we take into consideration is the well-known unidirectional MSN. In this case, N must be even, and all (horizontal and vertical) odd-numbered rings run clockwise, while all even-numbered[1] rings run counterclockwise (or viceversa).

The second topology that we take into consideration is the unidirectional square-grid (USG). In this case, N can take any value; all horizontal rings run in the same direction, and all vertical rings run in the same direction; such directions can be either clockwise or counterclockwise.

The third topology that we take into consideration is the alternating square-grid (ASG). In this case, N must be a multiple of 4, and four sets of adjacent rings alternate their direction. The horizontal and vertical rings for which $\left\lfloor \frac{4i}{N} \right\rfloor$ is even run clockwise, while all horizontal and vertical rings for which $\left\lfloor \frac{4i}{N} \right\rfloor$ is odd run counterclockwise (or viceversa). When $N = 4$, the ASG topology coincides with the MSN topology.

It is interesting to investigate the distance properties of the three topologies, in order to obtain a first insight in their possible behavior. In Fig. 1 we show the distributions of the minimum distances in the case of $M = 64$ nodes ($N = 8$); the horizontal axis measures the minimum distance[2] and the vertical axis measures the average number of nodes at the different values of minimum distance from one node in the topology. It can be observed that the minimum distance distributions of the MSN and ASG topologies are quite similar, whereas the minimum distance distribution for the USG topology is significantly worse, comprising a large number of long distances, as expected due to the unidirectionality of the topology.

The distribution of the minimum distance is not sufficient to characterize the quality of a topology; another very important aspect is the number of alternate paths with minimum distance. This figure is presented in Fig. 2 for the three topologies, where the horizontal axis measures the minimum distance, and the vertical axis measures the average number of alternate paths to reach a node at a given minimum distance from one node in the topology. From these results we see that the USG topology is the one providing by far the most spatial diversity, i.e., the largest number of minimum distance paths, for a given minimum distance value, followed by the MSN, and finally by the ASG (actually, for low distances, the ASG is slightly better than

[1] The number zero is considered as even.

[2] The minimum distance between nodes i and j is the minimum number of links to be traversed to reach node j from node i.

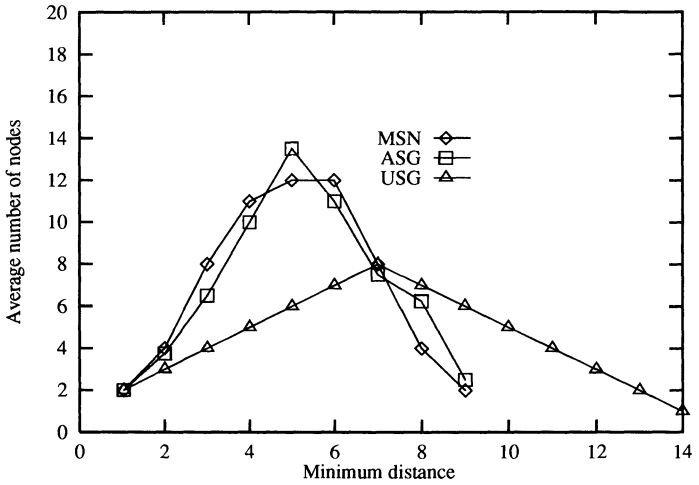

Figure 1: Distribution of the minimum distance for the three topologies

the MSN).

Yet another interesting parameter for the determination of the quality of a topology for a deflection network is the cost of a deflection, i.e., the cost to be paid whenever the routing algorithm deviates a cell from one of the shortest paths between the source and the destination. This cost is the difference in length between a shortest length path and a path comprising one deflection. The cost of a deflection is different in the three topologies: whereas in the MSN a deflection always has a cost 4, the cost becomes N for the USG topology, and can be seen to be variable between 2 and $2 + N/2$ in the ASG.

As far as distances are concerned, we have seen that, in the case of 64 nodes, the three topologies have strong and weak points; we can thus expect the USG and ASG topologies to exhibit performances comparable to those of the MSN, while providing the mapping advantages that we shall see in later sections. Indeed, simulation results obtained assuming a physical topology identical to the virtual meshed topology confirm this expectation, as we shall see in the next section.

For larger numbers of nodes we may expect the ASG topology to perform more similarly to the USG topology, due to the larger number of adjacent rings running in the same direction.

3 Results for Bursty Traffic with Square Grid Physical Topologies

The three topologies were simulated in the case of $M = 64$ nodes located on a square lattice with 35 km side, so that nodes are spaced on each row and column by 5 km. The physical topology is assumed to coincide with the square grid virtual topology; thus all point-to-point unidirectional links have length 5 km, with the exception of the return links that close the rings, whose length is 35 km. On each unidirectional link the data rate is assumed to be 100 Mbit/s. The internal node processing time is assumed to equal one slot.

Nodes generate bursty (data) traffic in the form of user *messages* whose length (in bytes)

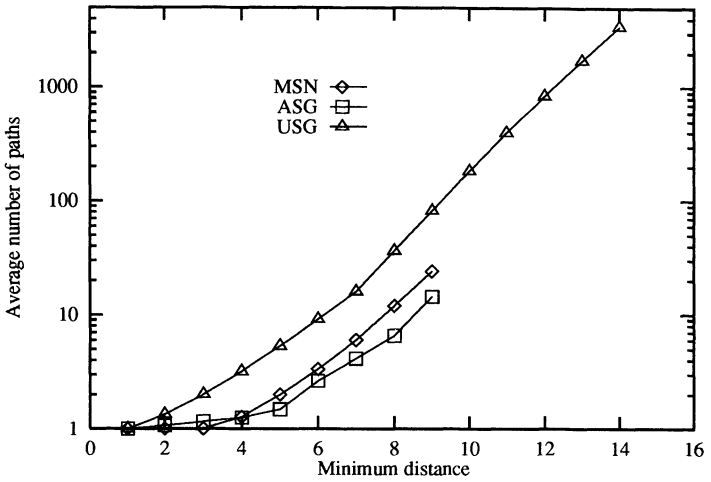

Figure 2: Number of alternate paths at minimum distance for the three topologies

obeys a geometric distribution truncated to a maximum value of 9188 bytes, with an average equal to 512 bytes. Messages are segmented into cells by first adding a fixed segmentation overhead equal to 8 bytes, and then filling the DQDB-like 44-byte payload of the required number of cells. The additional overhead (for addressing and protocol control information) inside each cell is supposed to be equal to 20 bytes, so that the total cell length equals 64 bytes. Simple minimum distance is taken as the metric for deflection routing: deflection occurs only when two cells have the same output port as the only port leading to their destinations along a shortest path. The deflected cell is randomly chosen.

Messages are generated according to a Poisson process of given rate, and the cells resulting from the segmentation are inserted into the node transmission buffer, that can contain at most 30 kbytes of data. If the transmission buffer cannot store the whole message, all its cells are lost. This means that our simulation model adopts a higher level of abstraction than is normally used when slotted protocols are simulated. We do not focus our attention onto slots, but on the service data units (called messages) at the interface between the access protocol and the layers above. As a consequence, our results account for the bandwidth losses due to slotting overheads and partial filling of slots. Measuring performance at the slot level, as is usually done, provides better results, but may not be representative of the service quality seen by a user of the network.

On the reception side, nodes keep track of how many messages they are currently receiving, i.e., how many instances of the message reassembly process are active. Enough buffer space is assumed to be available at the receiver so as not to consider losses of messages being received.

The execution of simulation runs provides estimates for several performance parameters, together with their confidence intervals. Since a rigorous statistical analysis of the simulation output is difficult, owing to the complexity of the model, an approximate second-order estimator was implemented following the batch means [13] approach, so that confidence intervals could be estimated, and a stopping rule for the simulation experiment could be automatically found. Also the initial transient time of the simulation experiment, when samples of the meaningful performance indices are not collected, is automatically determined according to the evolution of the average message access delay estimate. All simulation experiments were run for the time

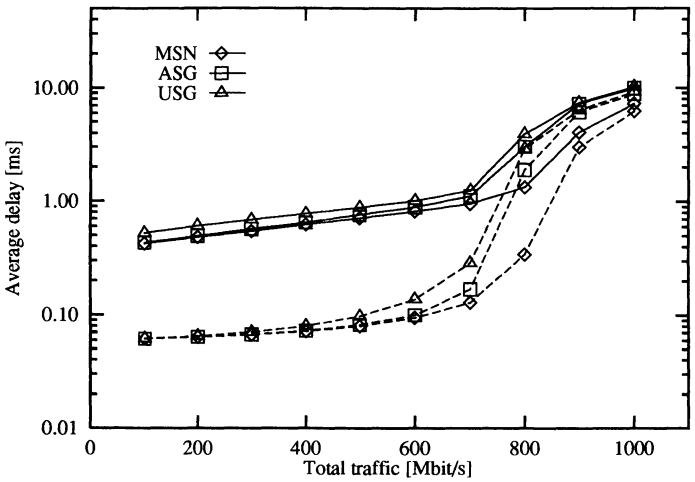

Figure 3: Average access delay (dashed lines) and average delivery delay (solid lines), versus total traffic

necessary for the average message access delay estimator to fall within a 5% confidence semi-interval with a 90% confidence level. It should be noted that confidence intervals for other performance parameters (e.g., the message access delay variance) may be much wider than for the average message access delay.

The following performance indices were measured in the simulation experiments.

- Average and variance of the user message *waiting delay*: the time a message spends in the transmission queue, from its arrival at the node until the instant when its first bit is transmitted.
- Average and variance of the user message *service delay*: the time from the instant when the first bit of the message is transmitted until the transmission of its last information bit.
- Average and variance of the user message *access delay*: the sum of the waiting and service delays.
- Average and variance of the user message *delivery delay*: the time from the arrival of the message at the source node until the reception of its last information bit at the destination, i.e., the sum of the access delay and the transit delay (which is variable due to deflections).
- Average and 99-th percentile of the number of deflections suffered by a cell.
- Average and 99-th percentile of the number of reassembly machines used at a node to reassemble messages; this is the number of messages whose reception is simultaneously in progress at a node due to the interleaving of cell arrivals.
- Message loss probability: the fraction of user messages lost due to buffer overflow at the source node.

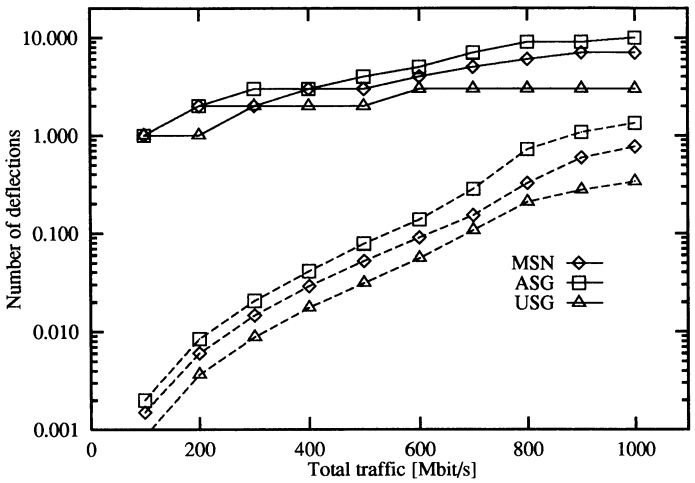

Figure 4: Average (dashed lines) and 99-th percentile (solid lines) of the number of deflections, versus total traffic

- Network utilization, or throughput: the average amount of user information – including message header bits, but excluding fields used for access and segmentation control – successfully transmitted on the network.
- Average number of channels traversed by a cell travelling from the source node to the destination node.

A sample of the numerical results obtained for the three virtual topologies in the case of a square grid physical topology are presented in Figs. 3–5.

Fig. 3 shows curves of the average access delay (dashed lines) and of the average delivery delay (solid lines), measured in ms, versus the total traffic generated by the 64 users, measured in Mbit/s of useful data (bytes in the user-generated messages). Both sets of curves indicate that the performances of the three topologies are similar, the MSN being the one that provides the lowest delays, the ASG exhibiting intermediate performance, and the USG suffering the longest average delays. Furthermore, we can observe that, as expected, the relative difference between the access and the delivery delay (the transit delay) is larger for lower traffic, where the access delay is quite low, and the transit delay is close to the propagation delay. Instead, for higher traffic, the access delay becomes large and dominates the delivery delay. The other results collected for averages and variances of the different types of delay show similar behaviors, and the same ranking of topologies.

Fig. 4 shows curves of the average (dashed lines) and 99-th percentile (solid lines) of the number of deflections suffered by a cell in the three topologies versus the total traffic generated by the 64 users, measured in Mbit/s of useful data. In this case we see that the USG topology suffers the lowest amount of deflections, the MSN behaves significantly worse, and the ASG has the largest number of deflections. However, it must be remembered that the cost of a deflection is different for the three topologies. It is thus interesting to observe the numerical results for the average number of channels traversed by a cell travelling from the source to the destination node. These results are presented in Fig. 5. The hierarchy among the three topologies is again the one

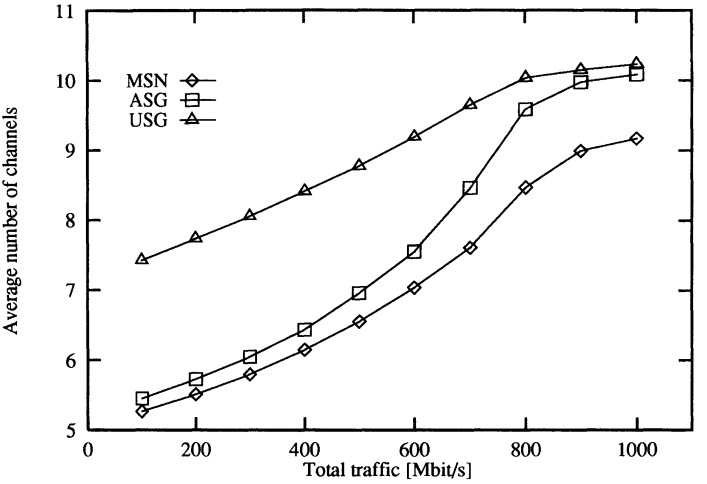

Figure 5: Average number of channels traversed by a cell versus total traffic

seen in Fig. 3: MSN is best, ASG is medium, and USG is worst, but once again differences are not that large.

4 Mapping Square Grid Meshes onto Multirings

A mapping $\mathcal{M} : T_{v1} \rightarrow T_{v2}$ is a functional relation establishing a one-to-one correspondence between the nodes in T_{v1} and the nodes in T_{v2}, as well as between the point-to-point unidirectional links in T_{v1} and those in T_{v2}.

Each one of the three regular mesh topologies that we are considering comprises N horizontal and N vertical rings, each made of N point-to-point links. The mapping $\mathcal{M}$ we propose is from a regular mesh virtual topology T_{v1} to a multiring virtual topology T_{v2} comprising a set of $2N$ parallel rings.

This mapping defines a one-to-one correspondence between one ring of the regular mesh and one ring of the multiring; this means that the rings in the multiring can be logically identified as row ring R_i^r, $i = 0, \ldots, N-1$, and column ring R_j^c, $j = 0, \ldots, N-1$. Naturally, the mapping $\mathcal{M}$ is such that node (i, j) is connected to R_i^r, and R_j^c, so that exactly N nodes are connected to each ring in the multiring, and that each node is connected to exactly two rings, like in the regular mesh topology. Furthermore, the node ordering of the mesh is maintained in the multiring: nodes are traversed by one ring of the multiring in the same order as they are by the corresponding column or row ring in the mesh.

Given the above setup, the definition of the mapping $\mathcal{M}$ only requires the determination of the node sequence along the multiring, so as to produce *the same node ordering along each ring*, both in the square grid and in the multiring. The mapping $\mathcal{M}$ thus reduces to an ordering function $\psi(i, j)$ that, given the two planar node indexes (row and column), produces the linear node index, i.e., gives the appropriate ordering of nodes along the multiring.

We can visualize a mapping through an $N \times N$ matrix Ψ, whose entries are the $\psi(i, j)$, or with a drawing in which 2 groups of N horizontal lines depict the rings R_i^r, and R_j^c, and M vertical segments indicate the ordering of nodes and their connections to the multiring.

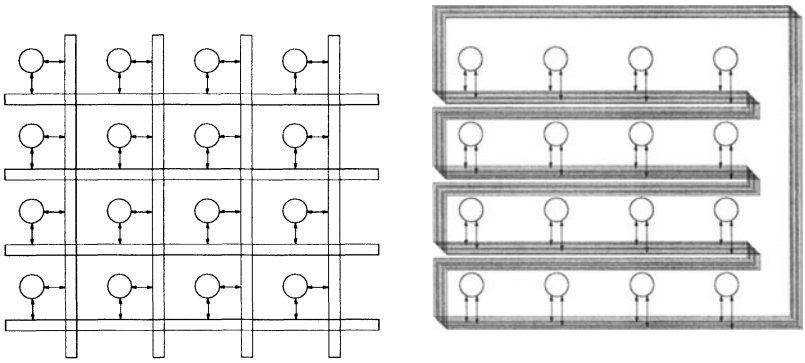

Figure 6: Square grid (left) and multiring (right) for a 16-node USG

$$\Psi = \begin{bmatrix} 0 & 1 & 2 & 3 & 4 & 5 & 6 & 7 \\ 8 & 9 & 10 & 11 & 12 & 13 & 14 & 15 \\ 16 & 17 & 18 & 19 & 20 & 21 & 22 & 23 \\ 24 & 25 & 26 & 27 & 28 & 29 & 30 & 31 \\ 32 & 33 & 34 & 35 & 36 & 37 & 38 & 39 \\ 40 & 41 & 42 & 43 & 44 & 45 & 46 & 47 \\ 48 & 49 & 50 & 51 & 52 & 53 & 54 & 55 \\ 56 & 57 & 58 & 59 & 60 & 61 & 62 & 63 \end{bmatrix}$$

Figure 7: Mapping from a 64-node USG to a multiring; matrix representation

Note that mapping the square grid onto the multiring reduces the topological stiffness for which MSN and BMSN are often criticized: the multiring more naturally adapts to the geographical locations of network users, and operations like adding or removing nodes from the network can be more easily implemented in the multiring. At the same time, the fault tolerance of the network is reduced, since all rings run in parallel.

USG

The simplest mapping is obtained in the case of the USG topology, for which one acceptable mapping is defined by the ordering function $\psi(i,j) = iN + j$. Note that other mappings exist in this case. A detailed study of the possible aspects relating to the mapping optimization is left for further study.

Fig. 6 shows the square grid and multiring-mapped 4×4 USG topologies. The matrix representation of the mapping is shown in Fig. 7 in the case of a 64-node USG, with the 8×8 matrix Ψ; the pictorial representation of the mapping is shown in Fig. 8 in the case of a 16-node USG.

MSN

The MSN topology can be mapped on a multiring only when $N \leq 4$, and, under this constraint, the mapping is not unique. The matrix and the pictorial representations of a possible mapping

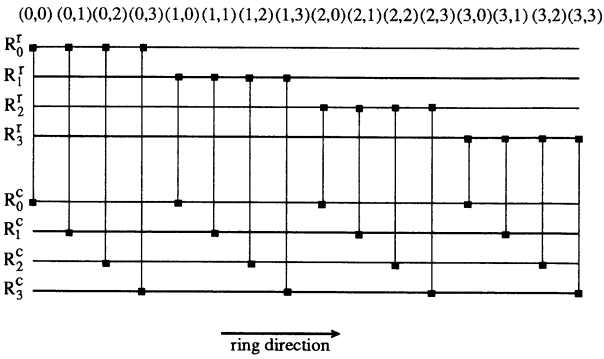

Figure 8: Mapping from a 16-node USG to a multiring; pictorial representation

$$\Psi = \begin{bmatrix} 0 & 4 & 8 & 12 \\ 5 & 1 & 13 & 9 \\ 10 & 14 & 2 & 6 \\ 15 & 11 & 7 & 3 \end{bmatrix}$$

Figure 9: Mapping from a 16-node MSN to a multiring; matrix representation

for a MSN topology with $N = 4$ ($M = 16$) are shown in Figs. 9 and 10. The proof of the fact that the mapping is impossible for MSN topologies with larger number of nodes can be found in [14].

ASG

The ASG topology can be viewed as a generalization of a 4×4 MSN topolgy, or as a combination of a 4×4 MSN topology and a USG topology. Indeed, in the ASG topology we can identify four groups of (horizontal and vertical) parallel rings running in the same direction (as in a USG), but the directions of the different groups alternate (as in a MSN).

This fact can be exploited to obtain the mapping from an ASG topology to a multiring through a combination of the mappings found for the USG and for the 4×4 MSN. The matrix representation of the resulting mapping in the case of a 64-node USG is shown in Fig. 11, where the 8×8 matrix Ψ is given. Mappings for larger ASG topologies are constructed in a similar manner. An algorithm for the construction of the mapping in this case is presented in [14].

5 Results for Bursty Traffic with Multiring Mapped Topologies

The second set of simulation results refers to USG and ASG topologies identical to those considered in Section 3, except for the fact that now the topologies are assumed to be mapped onto a multiring, using the mappings described in the previous section.

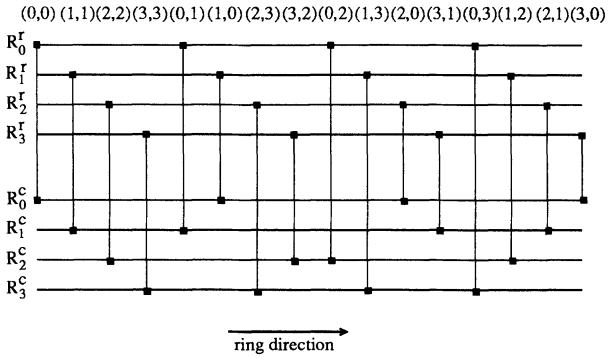

Figure 10: Mapping from a 16-node MSN to a multiring; pictorial representation

$$\Psi = \begin{bmatrix} 0 & 1 & 18 & 19 & 32 & 33 & 50 & 51 \\ 2 & 3 & 16 & 17 & 34 & 35 & 48 & 49 \\ 21 & 20 & 7 & 6 & 53 & 52 & 39 & 38 \\ 23 & 22 & 5 & 4 & 55 & 54 & 37 & 36 \\ 40 & 41 & 58 & 59 & 8 & 9 & 26 & 27 \\ 42 & 43 & 56 & 57 & 10 & 11 & 24 & 25 \\ 61 & 60 & 47 & 46 & 29 & 28 & 15 & 14 \\ 63 & 62 & 45 & 44 & 31 & 30 & 13 & 12 \end{bmatrix}$$

Figure 11: Mapping from a 64-node ASG to a multiring; matrix representation

It must be emphasized that the simulation results were obtained under assumptions that greatly penalize the multiring mapped topologies: indeed, we assume that users are physically located on the vertices of a square lattice graph (like in Section 3), and that the multiring channels can only run along the edges of the graph, so that the implementation of the multirings requires very long links. In other words, a virtual mesh topology T_{v1} (either USG or ASG) is mapped onto a second virtual topology T_{v2} (the multiring), which in turn is mapped onto a square grid physical topology T_p. As a result of this double mapping, each ring of the multiring becomes about ten times longer than was necessary when the mesh virtual topology T_{v1} was directly mapped onto the (identical) square grid physical topology T_p.

In spite of this increase in the length of the undirectional point-to-point links, the results obtained with the two topologies after the mapping onto the multiring are for the majority of cases very close to those obtained with a physical square grid topology where rings are much shorter (i.e., without the mapping onto the multiring). Often, the confidence intervals of the two point estimates obtained in the two cases overlap, so that the results are practically not distinguishable. For example, Fig. 12 shows curves of the average and variance of the access delay, measured in ms, and ms^2, respectively, versus the total traffic generated by the 64 users, measured in Mbit/s of useful data, before and after the mapping of the USG and ASG topologies onto the multiring. Clearly, the access delays are not influenced by the mapping.

Instead, the delivery delay is adversely affected by the longer propagation delays due to the

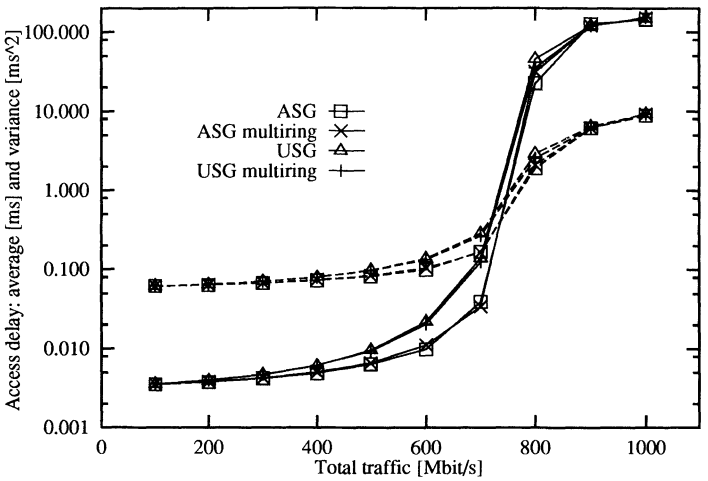

Figure 12: Average (solid lines) and variance (dashed lines) of the access delay versus total traffic before and after the mapping onto the multiring

greater distances in the multiring topologies. Fig. 13 shows curves of the average delivery delay measured in ms, versus the total traffic generated by the 64 users, measured in Mbit/s of useful data, before and after the mapping of the USG and ASG topologies onto the multiring. The fact that the delivery delays increase with the link length is unavoidable, but may not be a source of major problems, since propagation delays are constant.

The performance index that suffers most from the mapping is the number of reassembly machines used at destination nodes, whose average and 99-th percentile are shown in Fig. 14 for the USG and ASG topologies in the two cases with and without the mapping onto the multiring. It can be observed that the increase in the length of the links induces quite a significant increase in the number of reassembly machines needed at destination nodes, particularly at high traffic values. Ten-fold increases in the average and five-fold increases in the 99-th percentile are shown in the figure. An interesting observation is that the USG topology suffers significantly less than the ASG topology with respect to this parameter.

Before concluding this section we wish to stress once more the fact that the simulations were performed assuming a physical location of nodes that greatly favors the (non-mapped) square grid topologies, thus penalizing the multirings resulting from the mapping. The performance penalty to be paid is essentially due to the longer propagation delays, that induce longer delivery delays and a greater interleaving of slots in the network, so that a larger number of reassembly machines becomes necessary at the destinations. This may not be true if different physical distributions of users are assumed, and should be weighted against the possibility of introducing a true guaranteed-quality isochronous service on the multiring-mapped square grids, as we shall explain in the next section.

6 Providing Isochronous Services

The key advantage that is gained by mapping regular meshes onto multirings is in the fact that the linearity of the multiring topology allows the circulation of groups of slots devoted to the

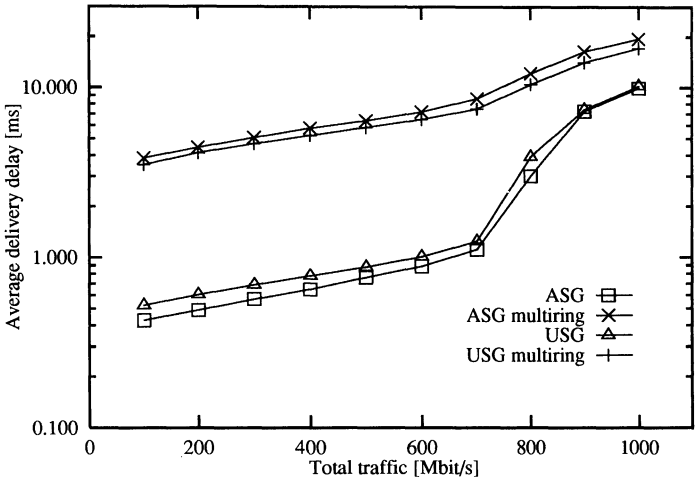

Figure 13: Average delivery delay versus total traffic before and after the mapping onto the multiring

transmission of isochronous traffic. Such slots will be called isochronous slots (I-slots), whereas the "normal" slots for bursty data traffic will be called bursty slots (B-slots). The slot type must be recognized by nodes; thus, slots are marked in their header with a *slot-type* bit.

In this section we describe one possible approach for the provision of isochronous services in the multiring-mapped topologies, but others are surely feasible.

We assume that a special *master* station is added in the multiring topology for the management of connection-oriented services for isochronous traffic. This master station is connected to all the $2N$ rings, contrary to normal nodes, that are only connected to 2 rings. We also assume that the isochronous services are provided with just one periodicity (e.g., one slot every 125 μs per connection), or at most with integer multiples or fractions of one periodicity (i.e., k slots every 125 μs or 1 slot every k 125 μs).

For a first description of the master station operations, assume that the isochronous traffic source and destination stations are not connected to a common ring (they are in different rows and columns of the square grid topology).

When a connection request is received (as a bursty message) and accepted, the master station periodically issues on two rings, one connected to the source and one to the destination, one I-slot. The periodic generation of the I-slots requires an insertion buffer in the master station, to decouple the generation of I-slots from the traffic in B-slots. The insertion buffer is also useful to equalize the latencies along all rings. Observe that the latencies along the rings are similar[3], since the number of nodes and the link lengths are equal on each ring; differences are only due to the slot misalignments produced by the node clocks, and can be compensated at the master station, or inside nodes. However, the generation of I-slots increases the insertion buffer lengths on the two rings, whose latencies become one slot longer than those of the other rings.

I-slots circulate around the rings on which they were issued, with no deflection since they are

[3]This is a crucial assumption in our proposal. Note however that a proper use of insertion buffers at the master station can compensate for different ring latencies. This observation leads to interesting topological alternatives to this paper's proposals.

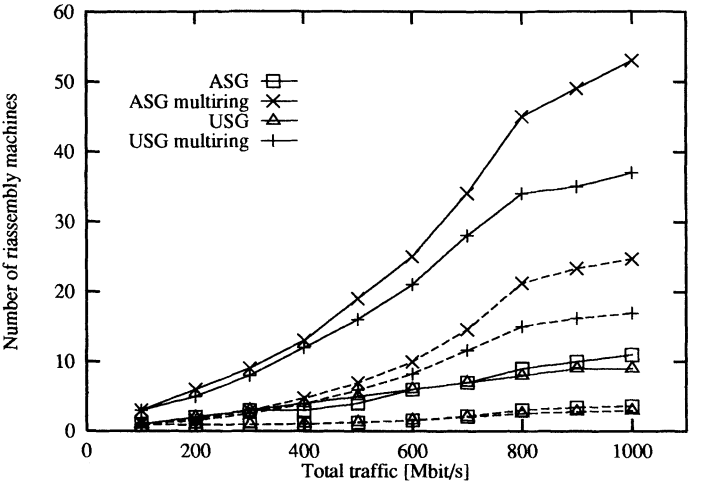

Figure 14: Average (dashed lines) and 99-th percentile (solid lines) of the number of reassembly machines used at destination nodes

given priority in the routing calculation over B-slots, and return to the scheduler at the same time, since their two ring latencies are equal. The source of the connection-oriented isochronous traffic loads one of these I-slots when it passes by. When this slot reaches the master station, its content is transfered to the second I-slot, that propagates to the destination station, where data are received. When the two slots return at the master station after completing their second tour of the multiring, they are removed from the insertion buffer, thus returning the latencies along the two rings to their original values. The same two I-slots can also be used for the data traveling in the opposite direction. Thus, two I-slots need to complete two tours of the multiring in order to establish one bidirectional connection, in this case. Of course, the allocation of I-slot pairs is performed periodically, and I-slots are removed from the ring after they complete their second tour.

The deflection routing algorithm in each network node must be integrated with a policy for handling I-slots. Whenever a node different from the master station receives one or two I-slots, it must let them continue their journey along the same ring on which they arrived. If one I-slot and one B-slot arrive at the same time, the B-slot has a lower priority in the output link selection, and may thus be deflected.

When the two users of the isochronous service are connected to a common ring (they share either one row or one column of the virtual square grid), only one I-slot is necessary to establish a bidirectional isochronous service, but two tours of the multiring are always required for bidirectional connections: the I-slot must be generated by the master station, reach the first user where it is loaded, then reach the second user where it is read, and loaded with the return data, then propagate to the first user where the return data are read, and finally reach the master station, where it is extracted from the insertion buffer.

A subtle point refers to the allocation of I-slots for multiple connections and to their influence on the delays along the rings. If at a given point in time a different number of I-slots is allocated on different rings, the property of the multiring topology of providing constant delays along all rings is destroyed, and the proposed technique for the provision of isochronous services fails. For

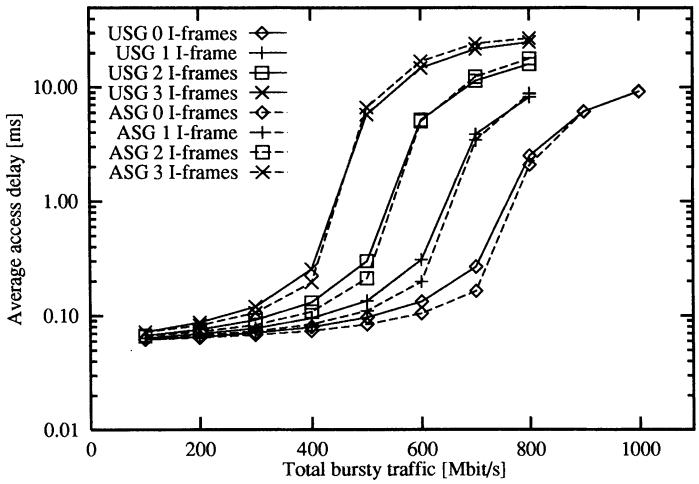

Figure 15: Average access delay versus total bursty traffic for a variable number of I-frames, and for the USG (solid lines) and the ASG (dashed lines) topologies

this reason, the allocation of I-slots must take place in *I-frames*. That is, whenever the master station needs to allocate one I-slot on some rings, then one I-slot is generated on all the $2N$ rings in the multiring. The set of $2N$ I-slots simultaneously generated by the master station is called an I-frame. By generating whole I-frames, the propagation delays along all the $2N$ rings are increased by an identical amount (one slot), so that the delays along all rings remain the same.

When new connection requests arrive at the master station, the allocation of new I-slots is performed within the available I-frames, and a new I-frame is generated only when the available ones cannot satisfy the new request.

Broadcast isochronous services can be implemented by allocating one half of an I-frame, comprising the N I-slots on all the row (or column) rings, and by copying at the master station the data written by the source on one I-slot to all other I-slots, that reach all possible destinations in their second tour of the multiring.

It should be observed that the proposed technique for the implementation of connection-oriented isochronous services is conceived so that such services are completely independent of the bursty traffic, and exactly synchronous. Of course, the presence of I-frames has a negative impact on the performance of bursty services. The quantification of this impact is provided by the numerical results in the next section.

7 Results for Mixed Traffic with Multiring Mapped Topologies

The third set of simulation results refers to USG and ASG topologies mapped onto a multiring, identical to those considered in Section 5, with the addition of isochronous services.

Without loss of generality, we assume that all connection requests must be satisfied with one slot every 125 μs. Rather than considering the effect of individual connections, which varies depending on the source and destination locations, we report results for a variable number of

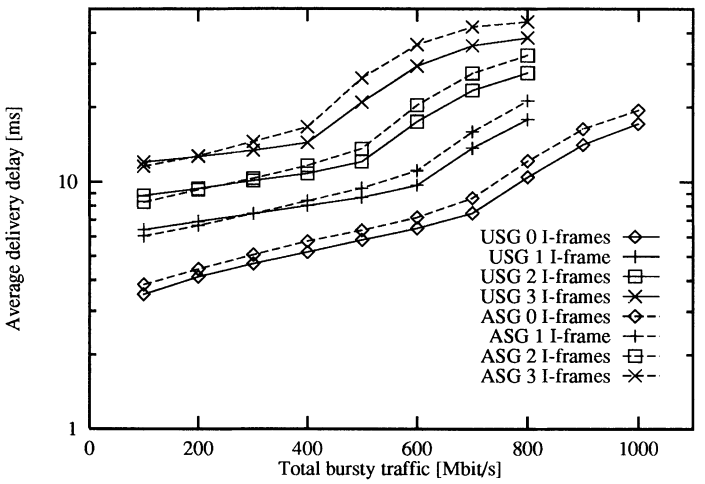

Figure 16: Average delivery delay versus total bursty traffic for a variable number of I-frames, and for the USG (solid lines) and the ASG (dashed lines) topologies

I-frames within a 125 μs period. With the chosen parameters, each I-frame with a 125 μs period consumes about 10% of the network raw data rate.

In general, the numerical results show a progressive degradation of the performance of bursty traffic with the increase in the number of I-frames. For example, Fig. 15 shows curves of the average access delay, measured in ms, versus the total bursty traffic generated by the 64 users, measured in Mbit/s of useful data, for 0, 1, 2, and 3 I-frames. The performance degradation is negligible for low loads, but the network capacity (in terms of bursty traffic) is progressively reduced by the allocation of I-frames. Similar results were obtained for the variance of the access delay. Note that these performance indices for the ASG topology mapped onto a multiring are similar to, but consistently better than those of the USG topology mapped onto a multiring.

Similar results are visible in Fig. 16, that shows curves of the average delivery delay, versus the total bursty traffic generated by the 64 users, for 0, 1, 2, and 3 I-frames. The performance degradation, however, is now remarkable even for low loads. With respect to the average delivery delay, the USG topology is slightly better than the ASG topology, except for very low loads.

The reduction in the network capacity due to the allocation of I-frames is very clearly seen from the curves of the network utilization presented in Fig. 17, versus the total bursty traffic generated by the 64 users, for 0, 1, 2, and 3 I-frames. The curves indicate that the cost of each I-frame in terms of capacity for the bursty traffic is close to 100 Mbit/s. Indeed, a simple computation shows that 16 slots with 44-byte payload that complete two tours of the multiring every 125 μs consume about 90 Mbit/s.

Finally, Fig. 18 shows curves of the average number of deflections incurred by a bursty data slot, and Fig. 19 presents results for the 99-th percentile of the number of reassembly machines active at a destination, versus the total bursty traffic generated by the 64 users, for 0, 1, 2, and 3 I-frames. Also for these performance indices we observe a significant degradation when the number of I-frames is increased; in particular, for high values of the bursty traffic, the number of reassembly machines necessary at destinations grows to values that are quite large. It is interesting to observe that the average number of deflections incurred by a bursty data slot is

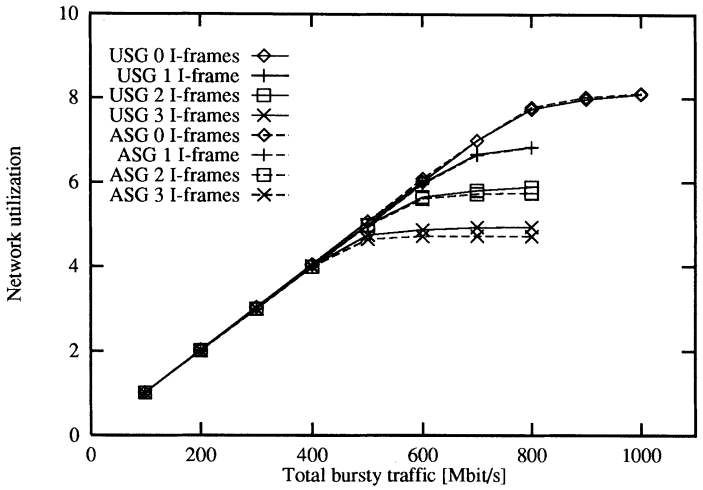

Figure 17: Network utilization versus total bursty traffic for a variable number of I-frames, and for the USG (solid lines) and the ASG (dashed lines) topologies

much larger for the ASG topology than for the USG topology. Moreover, also for what regards the number of reassembly machines, the USG topology exhibits lower values for both the average and the 99-th percentile.

These numerical results indicate that the price that has to be paid to establish guaranteed isochronous connection-oriented services is quite high if the bursty traffic load is high. Instead, if the load due to bursty data services is low, so that much extra bandwidth is available, the proposed technique for the provision of connection-oriented isochronous services appears to be promising.

8 Conclusions

We have presented an approach for the provision of guaranteed-quality connection-oriented services to isochronous traffic (as well as connectionless services to bursty traffic) in mesh networks with deflection routing.

The key element of the proposed approach is in the mapping of the mesh topology on a multiring topology. The linearity of the multiring simplifies the management of slots devoted to isochronous services. Several techniques can be envisioned to provide isochronous services in the network obtained with the mapping; our proposal is based on the availability of a master station equipped with an insertion buffer, and connected to all the rings in the multiring. The master station performs a centralized switching of connection oriented traffic, while bursty traffic is switched in a distributed manner according to deflection routing.

The mapping from a mesh to a multiring also simplifies the design of schemes to guarantee the fairness in the sharing of the network resources by the different users. This topic was however not covered in this paper.

The main advantages and drawbacks of the proposed approach can be summarized as follows.

- The complexity of the master station may be significant.

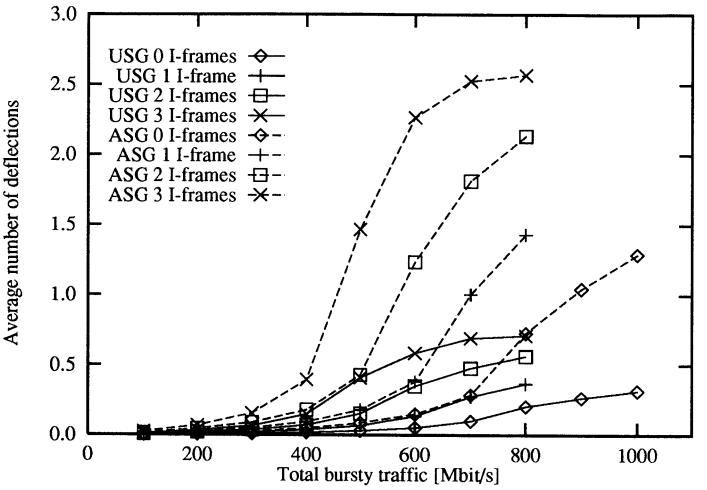

Figure 18: Average number of deflections versus total bursty traffic for a variable number of I-frames, and for the USG (solid lines) and the ASG (dashed lines) topologies

- The minimum amount of bandwith that can be allocated to isochronous services corresponds to one I-frame; it is not possible to allocate bandwidth for just one connection.
- The allocation of I-frames increases the probability of deflection of bursty traffic.
- The allocation of I-frames entails an increase of the ring latency; this has a negative effect on the bursty traffic.
- The mapping from a mesh topology onto a multiring may reduce the fault-tolerance characteristics of the network, making them closer to those typical of linear topologies.
- The synchronous arrival of slots to isochronous sources is guaranteed, and independent of the bursty traffic load.
- The routing algorithm in nodes remains simple.
- The allocation of whole I-frames avoids any sort of conflict among isochronous slots.
- Broadcast isochronous traffic can be handled in a very simple manner.

Numerical results obtained by simulation show that, when the network load due to bursty traffic is low, the introduction of isochronous connection-oriented services has a marginal impact on the performance of connectionless services for bursty traffic. When the network load due to bursty traffic is high, connectionless services suffer significantly for the introduction of isochronous connection-oriented services; however, the performance of the latter is insensitive of the network load level.

The two mesh topologies for which a mapping on a multiring was found provide similar performances for most of the considered quality indexes. Only from the point of view of the number of reassembly machines necessary at each receiving node, the USG topology seems to exhibit some advantages for the network size considered in this paper, but these results may be sensitive to the number of nodes in the network.

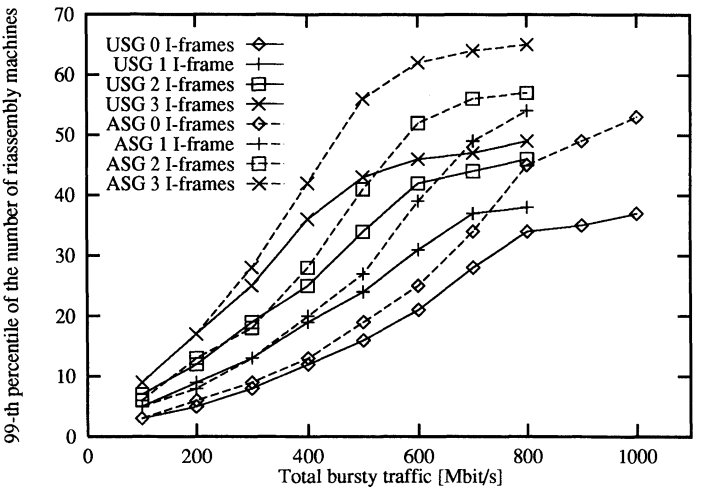

Figure 19: 99-th percentile of the number of reassembly machines used at a node versus total bursty traffic for a variable number of I-frames, and for the USG (solid lines) and the ASG (dashed lines) topologies

References

[1] B. Mukherjee, "WDM-Based Local Lightwave Networks Part II: Multihop Systems," *IEEE Network,* vol. 6, n. 4, pp. 20-32, July 1992

[2] N. F. Maxemchuck, "The Manhattan Street Network," *GLOBECOM'85,* New Orleans, LA, USA, December 1985, pp. 255-261

[3] N.F. Maxemchuck, "Routing in The Manhattan Street Network," *IEEE Transactions on Communications,* vol. 35, n. 5, pp. 503-512, May 1987

[4] F. Borgonovo, E. Cadorin, "Routing in The Bidirectional Manhattan Network," *Third International Conference on Data Communication Systems and their Performance,* Rio de Janeiro, Brasil, June 1987

[5] F. Borgonovo, E. Cadorin, "Locally-Optimal Deflection Routing in the Bidirectional Manhattan Network," *IEEE INFOCOM '90,* San Francisco, CA, USA, June 1990, pp. 458-464

[6] G. Albertengo, R. Lo Cigno, G. Panizzardi, "Optimal Routing Algorithms for the Bidirectional Manhattan Street Network," *ICC'91,* Denver, CO, USA, June 1991, pp. 1676-1680

[7] M. Ajmone Marsan, G. Albertengo, C. Casetti, F. Neri, G. Panizzardi, "On the Performance of Topologies and Access Protocols for High-Speed LANs and MANs," *Computer Networks and ISDN Systems,* vol. 26, n. 6-8, March 1994, pp. 873-893

[8] R. Krishnan, N.F. Maxemchuck, "Is There Life Beyond Linear Topologies? A Comparison of DQDB and the Manhattan Street Network," *INFOCOM '93,* San Francisco, CA, USA, April 1993

[9] F. Borgonovo, L. Fratta, F. Tonelli, "Circuit Service in Deflection Networks," *INFOCOM '91,* Bal Harbour, FL, USA, April 1991

[10] F. Borgonovo, L. Fratta, "Deflection Networks: Architectures for Metropolitan and Wide Area Networks," *Computer Networks and ISDN Systems,* vol. 24, pp. 171-183, 1992

[11] International Organization for Standardization, *Information Processing Systems - Fiber Distributed Data Interface (FDDI),* ISO DIS 9314, 1987

[12] IEEE Std 802.6-1990, *Distributed Queue Dual Bus (DQDB) Subnetwork of a Metropolitan Area Network (MAN),* 1991

[13] K. Pawlikowsky, "Steady-State Simulation of Queueing Processes: A Survey of Problems and Solutions," *ACM Computing Surveys,* vol. 22, n. 2, June 1990, pp. 123-170

[14] C. Pistritto, "Simulazione del Traffico Isocrono in Reti Metropolitane con Topologia a Maglia Quadrata," *Tesi di Laurea,* October 1993, in Italian

---

Marco Ajmone Marsan received the Dr.Ing. degree in electronic engineering from Politecnico di Torino, Italy, in 1974, and the Master of Science from University of California, Los Angeles, in 1978. He currently is a full professor at the Electronics Department of Politecnico di Torino. His research interests are in the fields of performance evaluation of data communication and computer systems, communication networks, and queueing theory. M. Ajmone Marsan is a Senior Member of IEEE.

Emilio Leonardi received the Dr.Ing. degree from Politecnico di Torino, Italy, in 1991. He is currently a Ph.D. student at the Electronics Department of Politecnico di Torino. His research interests are in the fields of optical networks, and performance evaluation of communication systems.

Fabio Neri received the Dr.Ing. and Ph.D. degrees in electronic engineering from Politecnico di Torino, Italy, in 1981 and 1987. He currently is an associate professor at the Electronics Department of Politecnico di Torino. His research interests are in the fields of discrete event simulation, queueing theory, and performance evaluation of communication systems.

Claudio Pistritto received the Dr.Ing. degree in electronic engineering from Politecnico di Torino, Italy, in 1993. The topic of his thesis is reported in this paper.