

Scalable Cycle-Breaking Algorithms for Gigabit Ethernet Backbones

Francesco De Pellegrini, David Starobinski, Mark G. Karpovsky, and Lev B. Levitin

Abstract—Ethernet networks rely on the so-called spanning tree protocol (IEEE 802.1d) in order to break cycles, thereby avoiding the possibility of infinitely circulating packets and deadlocks. This protocol imposes a severe penalty on the performance and scalability of large Gigabit Ethernet backbones, since it makes inefficient use of expensive fibers and may lead to bottlenecks. In this paper, we propose a significantly more scalable cycle-breaking approach, based on the novel theory of turn-prohibition. Specifically, we introduce, analyze and evaluate a new algorithm, called Tree-Based Turn-Prohibition (TBTP). We show that this polynomial-time algorithm maintains backward-compatibility with the IEEE 802.1d standard and never prohibits more than $1/2$ of the turns in the network, for any given graph and any given spanning tree. Through extensive simulations on a variety of graph topologies, we show that it can lead to an order of magnitude improvement over the spanning tree protocol with respect to throughput and end-of-end delay metrics. In addition, we propose and evaluate heuristics to determine the replacement order of legacy switches that results in the fastest performance improvement.

Index Terms—Graph Theory, System Design, Simulations.

I. INTRODUCTION

For many years, Ethernet has been the prevalent local area network (LAN) technology, offering a wide-range of services in a simple and cost-effective manner. With the standardization of Gigabit Ethernet protocols, the scope of Ethernet has widened even further [1]. A large number of corporations and service providers are now adopting Gigabit Ethernet as their backbone technology. Gigabit Ethernet backbones cope with the increasing traffic demand resulting from the deployment of high-speed LANs, home networks, Voice over IP, and high-bandwidth applications. The key advantage of Gigabit Ethernet over alternate technologies, such as ATM, is to maintain backward-compatibility with the over one hundred millions Ethernet nodes deployed world-wide and the large number of applications running on these nodes [2, 3].

Gigabit Ethernet has the same plug-and-play functionalities as its Ethernet (10 Mb/s) and Fast Ethernet (100 Mb/s) precursors, requiring minimal manual intervention for connecting hosts to the network. In addition, Gigabit Ethernet relies on full-duplex technologies and on a flow control (backpressure) mechanism that significantly reduce the amount of congestion and packet loss in the network [1, 4]. More specifically, the flow control mechanism (IEEE 802.3x) prevents switches from

losing packets due to buffer overflow. This protocol makes use of *Pause* messages, whereby a congested receiver can ask the transmitter to suspend (pause) its transmissions. Each Pause message includes a timer value that specifies how long the transmitter needs to remain quiet.

Currently, the network topology for Gigabit Ethernet follows the traditional rules of Ethernet. The spanning tree protocol (IEEE 802.1d) is used to avoid the occurrence of any cycle in the networks, thus pruning the network into a tree topology [5].

The reasons for breaking cycles are two-fold. The first is to avoid broadcast packets (or packets with unknown destination) from circulating forever in the network. Unlike IP, Ethernet packets do not have a Time-to-Live (TTL) field. Moreover, Ethernet switches must be transparent, which means that they are not allowed to modify headers of Ethernet packets.

The second reason is to prevent the occurrence of deadlocks as a result of the IEEE 802.3x flow control mechanism [6]. Such deadlocks may occur when Pause messages are sent from one switch to another along a circular path, leading to a situation where no switch is allowed to transmit. The use of a spanning tree precludes this problem, since deadlocks cannot arise in an acyclic network [7].

The spanning tree protocol works well in LAN networks, which are often organized hierarchically and under-utilized [8]. However, it imposes a severe penalty on the performance and scalability of large Gigabit Ethernet backbones, since a spanning tree allows the use of only one cycle-free path in the entire network. As pointed out by the Metro Ethernet Forum, an industry-wide initiative promoting the use of optical Ethernet in metropolitan area networks, this leads to inefficient utilization of expensive fiber links and may result in uneven load distribution and bottlenecks, especially close to the root [9, 10].

One of the current approaches to address this issue is to overlay the physical network with logical networks, referred to as virtual LANs [5]. A spanning tree instance is then run separately for each virtual LAN (or group of virtual LANs). This approach of maintaining multiple spanning trees can add significant complexity to network management and be very CPU-intensive [9].

In this paper, we propose a significantly more scalable approach, based on the novel theory of *turn-prohibition* [12, 13], in order to solve the cycle-breaking problem in Gigabit Ethernet backbones. Turn-prohibition is much less restrictive than link-prohibition, the approach employed to construct a spanning tree. The main idea is to consider pairs of links around nodes, referred to as turns [14], and show that all the cycles in a network can be broken through the prohibition of

F. De Pellegrini is with the Department of Information Engineering at the University of Padova. E-mail: depe@dei.unipd.it.

D. Starobinski, M. Karpovsky, and L. Levitin are with the Department of Electrical and Computer Engineering at Boston University. E-mail: {staro,markkar,levitin}@bu.edu.

carefully selected turns in the network (a turn (a, b, c) around node b is prohibited if no packet can be forwarded from link (a, b) to link (b, c)).

One of the main challenges in making use of the turn-prohibition approach is to maintain backward-compatibility with the IEEE 802.1d standard. Our main contribution in this paper is to propose and analyze a novel algorithm, called Tree-Based Turn-Prohibition (TBTP), to address this issue. This algorithm receives a graph along with a spanning tree, as its input, and generates a set of prohibited turns, as its output.

The TBTP algorithm possesses several key theoretical properties. First, it breaks all the cycles in the networks and preserves connectivity. Second, it *never* prohibits turns along the given spanning tree. In particular, if the tree is generated by the IEEE 802.1d protocol, then legacy switches can gradually be replaced by switches capable of running turn-prohibition. Third, the algorithm prohibits at most $1/2$ of the turns in the network. Thus, the total number of permitted turns in the network always exceeds the total number of prohibited turns. This result is valid for *any* given graph and spanning tree on it. Furthermore, it is generalizable to weighted graph topologies. In this latter case, the algorithm guarantees that the total weight of permitted turns always exceeds the total weight of prohibited turns in the network. We note that the constraint of permitting all the turns along a given spanning tree is critical for backward-compatibility. Without this constraint, a tighter bound on the fraction of prohibited turns can be achieved, namely $1/3$ [12, 13].

The rest of this paper is organized as follows. In Section II, we give a brief overview of the IEEE 802.1d protocol and discuss related work. In Section III, we introduce the TBTP algorithm, prove its main properties, and analyze its worst-case time complexity. In Section IV, we provide a general framework for maintaining backward-compatibility in an heterogeneous network composed of both “intelligent” switches, capable of running turn-prohibition, and legacy switches. We also propose heuristics to determine the order in which legacy switches should be replaced in order to achieve the fastest performance improvement. In Section V, we present numerical results, comparing the TBTP algorithm with the standard spanning tree algorithm and an earlier turn-prohibition approach called Up/Down [15]. Through extensive simulations on a variety of graph topologies, we show that the TBTP algorithm significantly outperforms the two other schemes with respect to throughput and end-to-end delay metrics. The last section is devoted to concluding remarks.

II. BACKGROUND AND RELATED WORK

A. The Spanning Tree Algorithm (IEEE 802.1d)

The spanning tree algorithm, first proposed in the seminal paper of [16], is the standard for interconnecting LANs, according to the IEEE 802.1d protocol.

This algorithm requires a unique identifier (ID) for every switch and every port within a switch. Using a distributed procedure, it elects the switch with the smallest ID as the root. A spanning tree is then constructed, based on the shortest path from each switch to the root (switch and port IDs are used to break ties).

Every switch brings the ports connected to its parent and children into a *forwarding* state. All remaining ports are placed in a *blocking* state. Ports in the forwarding state are used to forward data frames, while ports in the blocking state can only be used for forwarding signaling messages between switches.

Packet forwarding is based on a *backward-learning* process. When a switch receives a packet from a certain host S , via one of its active ports P , it assumes that the same port can be used in the reverse direction to forward packets to host S . This way, each switch progressively constructs a table, maintained in a local cache, that maps each destination with its corresponding port in the switch. Note that cache entries are valid only for a limited amount of time, determined by an associated timer. If the destination of a packet is unknown, then the packet is forwarded over all active ports except the incoming one. This action is commonly referred to as *flooding* or *broadcast*.

B. Improvements

Several enhancements have been proposed in the literature in order to mitigate congestion near the root.

The Distributed Load Sharing (DLS) technique enables use of some of the links not belonging to the tree [17, 18]. Under specific topological constraints, this technique can provide alternate paths to the spanning tree, thus helping to relieve local congestion. In general, however, no guarantee is provided on the overall performance improvement of the network.

A recent improvement over the DLS idea is proposed in [19]. This work devises a backward-compatible scheme, called STAR (Spanning Tree Alternate Routing), that finds alternate paths shorter than paths on the spanning tree, for any additive metric. It also introduces optimization techniques to determine which legacy switches should be candidates for replacement by new switches. As with previous DLS solutions, the STAR scheme does not provide bounds on the amount of prohibited resources in the network and does not address the problem of deadlocks caused by the IEEE 802.3x protocol. Readers are referred to [19] for other previous work related to the STAR protocol.

In [20], a solution, based on the technique of diffusing computation, is proposed in order to avoid infinite packet loops. Unfortunately, this solution is not backward-compatible with the IEEE 802.1d protocol and does not address the issue of deadlocks.

A scheduling approach, suitable for a lossless Gigabit Ethernet LAN, is proposed in [6] in order to avoid deadlocks. Although this solution avoids changes in the headers of Ethernet frames, it requires the replacement of all the switches in the network and is inherently incompatible with the spanning tree approach.

Our contribution substantially differs from previous work by guaranteeing a provable bound on the amount of prohibited resources (turns) in the network. Moreover, the TBTP algorithm that we propose is a general graph-theoretic approach for breaking cycles in networks. This algorithm is, thus, application-independent and can be used to avoid both infinite packet loops and deadlocks. Finally, our proposed algorithm is purposefully designed to be backward-compatible, as it relies

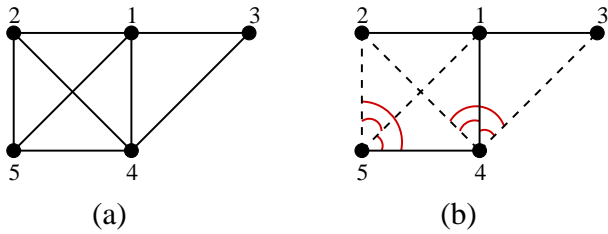


Fig. 1. (a) Example of a graph G . (b) Solution of the Up/Down algorithm for graph G . Tree-links are represented by solid lines and cross-links by dashed lines. Arcs represent prohibited turns (six turns are prohibited in total).

on the spanning tree generated by the IEEE 802.1d protocol. Preliminary ideas leading to this work were presented in [11].

III. SCALABLE CYCLE-BREAKING ALGORITHMS

In this section, we present our main contribution, the Tree-Based Turn-Prohibition (TBTP) algorithm, for breaking cycles in a scalable manner in Gigabit Ethernet Networks. We also briefly review an earlier turn-prohibition algorithm called *Up/Down* [15] and discuss the theoretical advantages of the TBTP algorithm over that algorithm. Beforehand, we introduce our network model and notations, and provide a formal statement of the problem.

A. Model

We model a Gigabit Ethernet network by a directed graph $G(V, E)$ where V is a set of nodes (vertices) representing switches and E is a set of links (edges). We do not consider end-hosts, since they would just be leaves on the graph. We restrict our attention to *bi-directional* network topologies, that is, networks where nodes are connected by full-duplex links. It is worth noting that essentially all modern Gigabit Ethernet networks make use of full-duplex links (in contrast to the original Ethernet where nodes were communicating over a shared medium link).

We define a *cycle* to be a path whose first and last links are the same, for instance, $(n_1, n_2, n_3, \dots, n_{\ell-1}, n_{\ell}, n_1, n_2)$. Our goal is to break all such cycles in the underlying graph in order to avoid deadlocks and infinitely circulating packets.

Note that the literature in graph-theory typically defines a cycle as a path such that the initial and final nodes in the path are the same [21]. We refer to this latter definition as a *cycle of nodes*. A spanning tree breaks all cycles of nodes in a graph.

Breaking all cycles of nodes is, however, unnecessary in general. For instance, referring to Figure 1(a), the path $(5, 1, 4, 3, 1, 4)$ contains a cycle, while the path $(5, 1, 4, 3, 1, 2)$ does not (although it contains a cycle of nodes). A cycle is, thus, created only when the same buffer or port is traversed twice, in the same direction. In particular, a path may traverse several times the same node without creating a cycle.

A pair of input-output links around a node is called a *turn*. The three-tuple (i, j, k) will be used to represent a turn from link (i, j) to link (j, k) , with $i \neq k$. For instance, in Fig. 1(a), the three-tuple $(2, 1, 3)$ represent the turn from link $(2, 1)$ to link $(1, 3)$ around node 1.

Due to the symmetrical nature of bi-directional graphs, we will consider the turns (i, j, k) and (k, j, i) to be the same turn. As shown in the sequel, an efficient approach for breaking cycles in a network is based on the *prohibition* of turns. For example, in Fig. 1(a), prohibiting the turn $(2, 1, 3)$ means that no packet can be forwarded from link $(2, 1)$ to link $(1, 3)$ (and from link $(3, 1)$ to link $(1, 2)$). Note that turns involving leaves in the graph can always be permitted, since they do not belong to any cycle. We will denote by $R(G)$ the set of all turns in the network. If the degree of each node i is d_i , then $|R(G)| = \sum_{i=1}^{|V|} d_i(d_i - 1)/2$.

B. Statement of the Problem

Suppose a spanning tree $T(G) = (V_T, E_T)$ providing connectivity for a graph G is given *a-priori*. We refer to links belonging to this tree as *tree-links*. Other links are referred to as *cross-links*.

Denote a set of prohibited turns by $S_T(G)$. Our goal is to determine a set $S_T(G)$, such that

- 1) Every cycle in the network contains at least one turn from $S_T(G)$, i.e., all the cycles in the network are broken.
- 2) $S_T(G)$ contains a minimum number of elements.
- 3) All the turns between tree-links in $T(G)$ are permitted.

Our problem, therefore, is to determine a minimal cycle-breaking set $S_T(G)$ that does not include turns from $T(G)$. If the tree is generated by the IEEE 802.1d protocol, then the solution to this problem would allow us to gradually replace legacy switches by switches capable of running turn-prohibition. We present, now, two algorithms to address this problem. For simplicity of exposition, we first assume that all nodes in the network are *intelligent* nodes, that is, nodes capable of implementing turn-prohibition. We address backward-compatibility issues with legacy nodes in Section IV.

C. Up/Down

A possible approach for the construction of a cycle-breaking set $S_T(G)$, is the so-called *up/down routing* algorithm [15]. In this approach, a spanning tree $T(G)$ is first constructed. Then, nodes are ordered according to the level at which they are located on the tree. The level of a node is defined as its distance (in number of hops) from the root. Nodes at the same level are ordered arbitrarily.

Once the nodes are ordered, a link (i, j) is considered to go “up” if $i > j$. Otherwise, it is said to go “down”. A turn (i, j, k) is referred to as an up/down turn if node $i > j$ and $j < k$. Respectively, a down/up turn is a turn such that $i < j$ and $j > k$. Clearly, any cycle must involve at least one up/down turn and one down/up turn. Therefore, it is possible to break all the cycles in a graph by prohibiting all the down/up turns. This means that packets can be transmitted over cross-links as long as they are not forwarded over down/up turns. An illustration of the Up/Down algorithm is provided in Fig. 1(b).

The Up/Down routing algorithm achieves much higher performance than the simple spanning tree approach implemented by the IEEE 802.1d protocol, as shown by our

simulations in Section V. However, the performance of this algorithm depends critically on the selection of the spanning tree and its root node. In particular, it has been shown that the fraction of turns prohibited by this scheme, that is, the ratio $|S_T(G)|/|R(G)|$, can be arbitrarily close to 1 in some networks [12].

D. The Tree-Based Turn-Prohibition (TBTP) Algorithm

We now introduce a novel cycle-breaking algorithm, called Tree-Based Turn-Prohibition (TBTP). This algorithm prohibits at most 1/2 of the turns for *any* given graph G and spanning tree $T(G)$. This is in contrast with the Up/Down algorithm that does not provide any guarantees on the fraction of prohibited turns in the network.

1) *The Algorithm:* The TBTP algorithm is a greedy procedure that receives a graph G and an associated spanning tree $T(G)$ as its arguments. At each iteration, the algorithm breaks all the cycles involving some selected node i^* . Towards this end, the algorithm prohibits all the turns around node i^* (except for turns between tree-links) and permits all turns of the type (i^*, j, k) , where (i^*, j) are cross-links originating from node i^* . All the cross-links connected to node i^* are then deleted and the procedure is repeated until all the nodes have been selected and no more cross-links remain. Following is a pseudo-code for the algorithm:

Algorithm TBTP($G, T(G)$):

- 1) For each node i , that has adjacent cross-link(s), construct two sets of turns

$$A(i) = \{(i, j, k) | (i, j) \in E \setminus E_T, (j, k) \in E\}$$

$$P(i) = \{(j, i, k) | (j, i) \in E \setminus E_T, (i, k) \in E\}$$

- 2) Select node i^* maximizing $|A(i)| - |P(i)|$ (if there are multiple such nodes, choose one at random).
- 3) Add all the turns from $P(i^*)$ into $S_T(G)$.
- 4) Delete all cross-links $(i^*, j) \in E \setminus E_T$, and repeat the procedure until all the cross-links have been deleted.

The intuition behind the node selection is as follows. Each node i , is associated with a potential set of permitted turns $A(i)$ and a potential set of prohibited turns $P(i)$. At each iteration, the selected node is the one that maximizes the difference between the cardinality of the sets, namely $|A(i)| - |P(i)|$. We will show in the sequel that there always exists at least one node for which this difference is greater or equal to zero. Thus, this selection strategy guarantees that the algorithm prohibits at most 1/2 of the turns in the graph.

Figure 2 illustrates the algorithm. At the first step of the algorithm, the sets $A(i)$ and $P(i)$ are computed for each node i . For instance, for node 1, these sets are $A(1) = \{(1, 5, 2), (1, 5, 4)\}$ and $P(1) = \{(2, 1, 5), (3, 1, 5), (4, 1, 5)\}$. The selected node is node 3 for which $|A(3)| - |P(3)| = 2$ (either node 2 or 5 could have been selected as well). As a result, turn $(1, 3, 4)$ is prohibited. The procedure is then repeated, but without cross-link $(3, 4)$. As shown in the figure, the procedure continues until no more cross-links remain. The final set of prohibited turns is $S_T(G) = \{(1, 3, 4), (1, 4, 2), (1, 5, 2), (1, 5, 4), (2, 5, 4)\}$ (other solutions

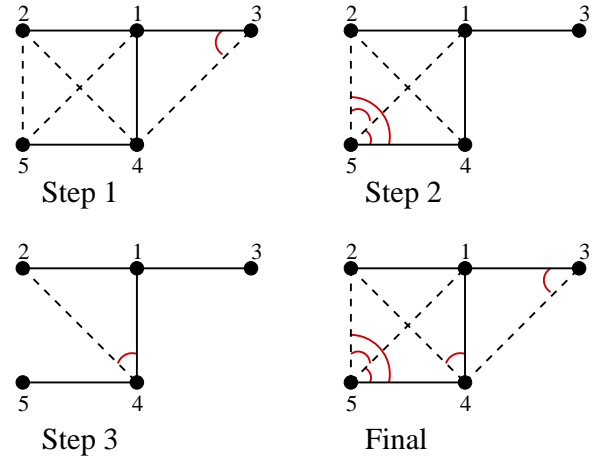


Fig. 2. Successive steps of the TBTP algorithm and final result (five turns are prohibited in total).

are also possible). Thus, the algorithm ends up in prohibiting five turns which is one fewer than Up/Down.

Note that the TBTP algorithm requires each switch to have full knowledge of the topology. Nevertheless, the algorithm can still be implemented in a decentralized fashion as a link-state algorithm, such as OSPF [5], where every switch knows the global topology.

2) *Analysis:* We now prove the main properties of the TBTP algorithm.

Theorem 1 *The TBTP algorithm preserves network connectivity.*

Proof: The algorithm never prohibits turns between tree-links. Thus, network connectivity is provided by the spanning tree $T(G)$. ■

Theorem 2 *The TBTP algorithm breaks all the cycles.*

Proof: If any cycle exists, then it must involve at least one cross-link. Thus, in order to prove the theorem, we need to show that a turn containing a cross-link cannot belong to any cycle.

The proof is by induction. The induction hypothesis is that each iteration of the algorithm, none of the turns around nodes previously selected, and containing a cross-link, belong to any cycle. This hypothesis is clearly true for the first node selected, say node i_1 . This is because the TBTP algorithm prohibits all the the turns of the type (j, i_1, k) around node i_1 , where (j, i_1) are cross-links.

Now suppose that $n - 1$ nodes have already been selected and none of the turns, containing a cross-link, around these nodes belong to any cycle. The next chosen node is, say, node i_n . We distinguish between two types of turns around node i_n . First, we consider turns which contain a cross-link adjacent to one of the previously selected nodes. These turns have been permitted in one of the former steps of the algorithm, but can not lead to a cycle, by the induction hypothesis. Second, we consider the turns around node i_n , containing a cross-link, and that do not involve a previously

selected node. The TBTP algorithm prohibits all these turns. As a consequence, the algorithm breaks all the remaining cycles that could have involved node i_n with one of its adjacent cross-links. The induction hypothesis is thus justified, and the proof of the theorem is complete. ■

We now show that the TBTP algorithm prohibits at most 1/2 of the turns in the network. We first prove the following lemma:

Lemma 1 *At each step of the algorithm, the following inequality holds*

$$\sum_{i \in G} (|A(i)| - |P(i)|) \geq 0$$

Proof: Each turn (i, j, k) , $\{(i, j) \in E \setminus E_T, (j, k) \in E\}$, can appear only once as a prohibited turn, namely in the set $P(j)$. On the other hand, the same turn will appear as a permitted turn in the set $A(i)$, and possibly also in the set $A(k)$ if $(j, k) \in E \setminus E_T$. Thus, each turn is counted once in the set of prohibited turns and, at least, once in the set of permitted turns, thereby proving the lemma. ■

Theorem 3 *The TBTP algorithm prohibits at most 1/2 of the turns in the network.*

Proof: From Lemma 1, there must exist at least one node i for which the difference $|A(i)| - |P(i)|$ is greater or equal to zero. Since, at each step, the algorithm selects the node i^* that maximizes this difference, it is guaranteed that the number of permitted turns is greater or equal to the number of prohibited turns. ■

It is worth noting that the number of turns permitted by the TBTP algorithm is actually strictly greater than the number of prohibited turns. This is because turns between tree-links are all permitted as well.

3) *Algorithm complexity:* We next show that the TBTP algorithm has practical computational complexity.

Theorem 4 *The computational complexity of the TBTP algorithm is $\mathcal{O}(|V|^2 d^2)$, where d represents the degree of graph G (i.e., the maximum degree of any node in the graph).*

Proof: The analysis of the computational complexity can be broken down into the following components:

- 1) The construction of the spanning tree $T(G)$. This component has complexity $\mathcal{O}(|V|d)$.
- 2) Computations of the sets $A(i)$ and $P(i)$. Computing the number of permitted/prohibited turns around any node i is of the order $\mathcal{O}(d^2)$. At each iteration of the algorithm, this computation is performed for all $|V|$ nodes in the graph. Since the number of iterations is at most $|V|$, we obtain that the overall complexity of this component is $\mathcal{O}(|V|^2 d^2)$.
- 3) The selection of node i^* . At each iteration, the node that maximizes the difference $|A(i)| - |P(i)|$ is selected, which requires $\mathcal{O}(|V|)$ computations. Since there are at most $|V|$ iterations, the complexity of this component is $\mathcal{O}(|V|^2)$.

We therefore obtain that the algorithm complexity is $\mathcal{O}(|V|^2 d^2)$. ■

Note that the above analysis assumes a straight-forward implementation of the algorithm. The computational complexity could be reduced by using advanced data structures, such as priority queues. Likewise, at each step of the algorithm, we do not have to repeat the computations of the set $A(i)$ and $P(i)$ for each node i , but only for a subset of the nodes.

4) *Extension for Weighted Graphs:* So far, we have only considered the case of unweighted graphs. However, in switched Ethernet networks, different links may have different capacity, e.g., 100 Mb/s versus 1 Gb/s. Consequently, the relative importance of different turns vary as well.

In order to address this issue, we can extend our results to weighted graphs. Each turn (i, j, k) in the graph is associated with a weight $w(i, j, k)$. This weight can be set according to any metric of interest [13].

The TBTP algorithm for weighted graphs remains the same as for unweighted graphs. We just have to replace $|A(i)|$ and $|P(i)|$ by the sum of the weight of turns in the corresponding sets. Moreover, since the proof of Theorem 3 holds unchanged, we obtain the following result for weighted graph topologies:

Corollary 1 *The sum of the weights of prohibited turns by the TBTP algorithm is at most 1/2 of the sum of the weights of all turns in G .*

As an illustration, consider the example of Fig. 2, but with the following distribution of weights for the turns: $w(1, 3, 4) = 10$ and the weights of all other turns set to 1. In such a case, the algorithm would prohibit the following set of turns $S_T(G) = \{(2, 4, 1), (2, 4, 3), (2, 4, 5), (3, 4, 1), (3, 4, 5), (2, 5, 1), (2, 5, 4)\}$. The overall weight of the prohibited turns is 7 which represents 25% of the total weight of the turns in the network.

IV. BACKWARD-COMPATIBILITY

In the description of the Up/Down and TBTP algorithms in the previous section, we have assumed that all switches in the network are capable to perform turn-prohibition. In large networks, however, a massive replacement of all switches is a major issue and very unlikely to happen. In this section, we suggest a strategy for a smooth transition towards a complete replacement.

A. Approach

We consider an heterogeneous network consisting of both intelligent and legacy nodes. In order to ensure backward-compatibility, we require intelligent switches to be able of running both the spanning tree algorithm (IEEE 802.1d) and turn prohibition.

At the very beginning, all the switches run the distributed spanning tree algorithm. As a result, every node belongs to a spanning tree $T(G)$ that is generated by the IEEE 802.1d protocol.

Next, intelligent nodes send “neighbor discovery” messages to their directly connected neighbors. As a result, an intelligent node knows if it is connected to another intelligent node or to a legacy node.

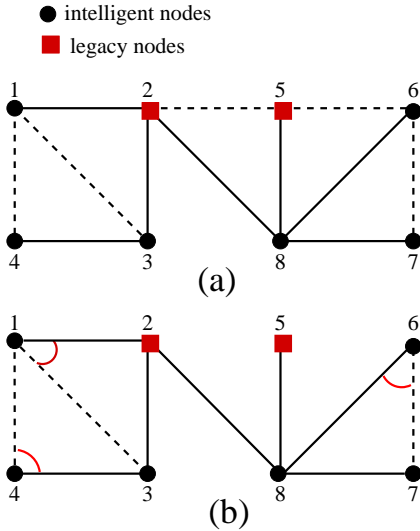


Fig. 3. (a) An heterogeneous graph, G , composed of legacy and intelligent nodes. (b) Corresponding modified topology G' and prohibited turns.

Now suppose that two intelligent nodes are connected by a cross-link. Then use of this link may be possible, depending on the turn-prohibition algorithm being employed. However, if a cross-link connects between an intelligent node and a legacy node, then use of this link is not possible. This is because the legacy node would not accept to forward or receive packets over that link.

Therefore, a backward-compatible solution has to examine each turn (i, j, k) , where (i, j) is a cross-link and both nodes i and j are intelligent, and decide if this turn can be permitted or not. We now distinguish between the cases of Up/Down and TBTP.

1) *Up/Down*: We remind that the Up/Down algorithm requires nodes to be ordered. This information can readily be obtained from the IEEE 802.1d protocol, since each node n knows its distance $d(n)$ from the root. Thus, each turn (i, j, k) for which $d(i) \neq d(j)$ and $d(j) \neq d(k)$ can be resolved. It will be prohibited only if node j is farther away from the root than the two other nodes. Also, turns such that $d(i) = d(j)$ can be resolved using the IDs of the switches, that is, if $ID(i) > ID(j)$ then $i > j$ and vice-versa.

Note that this approach works as well when node k in the turn (i, j, k) is a legacy node, because node j knows whether node k is its parent or one of its children (remind that if node k is a legacy node, then link (j, k) must be a tree-link).

2) *TBTP*: As mentioned in Section III, the TBTP algorithm requires each node to have global knowledge of the network topology. This is of course not possible in an heterogeneous network, since legacy nodes would not participate in the process of collecting the topology (e.g., by generating or forwarding link-state packets).

Our solution is to implement the TBTP algorithm independently for each connected group of intelligent nodes. We refer to such groups as *components of connectivity*. Figure 3(a) gives an example of an heterogeneous network with two components of connectivity, namely nodes 1,3,4 and nodes 6,7,8.

In each component connectivity, the intelligent nodes collect

the internal topology of that component and run the TBTP algorithm to decide which turns should be permitted or prohibited.

In order to show why this approach breaks all the cycles, consider the following modified topology G' . This topology consists of the original topology G , but without all the cross-links connected to legacy nodes (on one or both ends). We now observe that our backward-compatible solution is equivalent to having run the centralized version of the TBTP algorithm, described in Section III, on the topology G' . Therefore, it follows from Theorem 2 that all the cycles in the network are broken. Figure 3 depicts an example of an heterogeneous graph G and its corresponding modified topology G' .

B. Packet Forwarding (Routing)

In a turn prohibition-based algorithm, such as Up/Down or TBTP, several paths may exist between a source and a destination, as opposed to a spanning tree where only a single path exists. Thus, there is a need to deploy a routing algorithm that can determine the “best” path, according to some appropriate metric, from each source to each destination.

In [12], we have described one such routing algorithm that is a generalization of the Bellman-Ford routing algorithm. In the case of TBTP, the implementation of this routing algorithm adds little overhead, since every node knows the full topology.

Similarly to the original Bellman-Ford algorithm, the generalized version can be implemented in a fully distributed way. Therefore, when used with Up/Down, nodes do not need to have global knowledge of the topology. Furthermore, since the network contains no cycles, the count-to-infinity problem will not arise.

In the case of heterogeneous networks, the generalized Bellman-Ford algorithm is implemented separately in each component connectivity. In order to maintain backward-compatibility, intelligent nodes which are directly connected to legacy nodes must also implement the backward-learning process of the IEEE 802.1d protocol. In such a case, when an intelligent node learns about a certain destination outside its component of connectivity, it must inform all the other nodes within its component.

Finally, each routing entry has an associated time-out, e.g. 15 minutes. If the entry is not refreshed within this period of time, then it is removed from the routing table. As with the IEEE 802.1d protocol, a switch that receives a packet for an unknown destination will broadcast the packet along links of the spanning tree $T(G)$.

C. Heuristics for Node Replacement

By implementing turn-prohibition instead of link-prohibition, the TBTP and Up/Down algorithms clearly outperform the simple spanning tree algorithm. Therefore, the gradual replacement of legacy nodes by intelligent node will improve network performance.

An interesting question within this context is as to which legacy nodes should be replaced first in order to achieve maximum improvement. One simple strategy is the *random*

one in which a legacy node is picked at random and replaced by an intelligent node.

We propose here another heuristic. This heuristic relies on the fact that a random replacement could lead to sets of intelligent nodes completely isolated, i.e., no cross-link can be enabled for forwarding. Instead, we propose a *Top-Down* approach whereas the replacement starts from the root, i.e., replacing the level 1 node first, then proceed with level 2 nodes and so on. This replacement strategy will guarantee the existence of at most one component of connectivity in the network. Moreover, it will relieve congestion from the root. Therefore, we expect that this heuristic will lead to a faster performance improvement than the random one. Our numerical results in the next section confirm this rationale.

D. Reconfiguration

In the case of node/link failures, the spanning tree $T(G)$ is reconfigured according to the specifications of the IEEE 802.1d standard or its recently improved version, IEEE 802.1w, which supports faster spanning tree reconfiguration. Once the spanning tree is reconfigured, a new set of permitted/prohibited turn is determined using the procedure described above.

V. NUMERICAL RESULTS

In this section, we compare the performance of the various cycle-breaking approaches described in the previous sections. In order to obtain comprehensive sets of results, we make use of both *flow level* and *packet level* simulations.

Flow level simulations are based on fluid models. They provide fast, qualitative results and can be run over large number of graphs with different topology properties.

Packet level simulations, on the other hand, model network dynamics at the granularity of single packets. They can provide more accurate estimates on quality of service metrics, such as end-to-end delay and throughput. Unfortunately, they are computationally intensive, especially for large graphs.

A. Flow Level Simulations

The goal of our flow level simulations is to compare the performance of the algorithms as related to different network topology properties, such as the size of the network or the degree of the nodes. Another important objective is to evaluate the heuristics for node replacement discussed in Section IV.

Our simulations are based on random, connected graphs. Each node has fixed degree d , and the network size (number of nodes in the network) is $|V|$. We assume that all links in the network have the same capacity C and we set $C = 1$ Gb/s. All the results presented correspond to average over 100 graphs with identical parameters.

Once a random graph is generated, each of the cycle-breaking algorithms (spanning tree, Up/Down and TBTP) is run on top of it in order to determine a set of prohibited links/turns. Routing matrices over a turn-prohibited graph are computed using the generalized Bellman-Ford algorithm of [12].

As a reference, we also simulate an ideal scheme when no turn in the graph is prohibited. We refer to this scheme as *shortest path*, since each flow is established along the shortest path from the source to the destination.

We consider two metrics. The first is the fraction of prohibited turns in the network, that is, the ratio $|S_T(G)|/|R(G)|$. A lower value for this quantity is considered to be better since it implies less unutilized network resources.

The second metric is throughput. We compute this metric as follows. We assume a fluid model where flows are transmitted at a fixed rate. Each node establishes a session (flow) with k other nodes in the network, picked at random. In all of our simulations, we set $k = 4$. Each flow is routed along the shortest-path over the turn-prohibited graph (if multiple routes exist, then one is selected at random). Next, we determine the bottleneck link, which is the link shared by the maximum number of flows. The throughput is then defined as the capacity of the bottleneck link divided by the number of flows sharing it. In other words, the throughput is the maximum rate at which each flow can be transmitted without saturating the network.

1) *Complete Replacement of the Switches*: In our first set of simulations, we compare the performance of the algorithms assuming that all switches in the network are intelligent, that is, capable of performing turn-prohibition.

We first evaluate the performance and scalability properties of each algorithm as a function of the number of nodes in the network $|V|$. We assume the degree of each node to be $d = 8$.

Table I shows that the TBTP algorithm prohibits about 10% fewer turns than Up/Down. As expected, the simple spanning tree approach performs far worse, prohibiting about 90% of the turns in the network¹. These results are almost insensitive to the network size.

Figure 4 depicts the throughput performance with 99% confidence intervals. The results presented are in agreement with those obtained for the fraction of prohibited turns. We observe that, for 32 nodes, the TBTP achieves a throughput approximately 10% higher than that of Up/Down. Moreover, the relative difference in the performance between these two algorithms increases with the network size. Another important observation is that the throughput of TBTP is within a factor of at most 1.5 from that of the “shortest-path” scheme. This means that the cost of breaking all the cycles in the network may not be too significant in terms of network performance. This is in clear contract with the spanning tree approach which achieves an order of magnitude lower throughput.

¹Although the spanning tree approach prohibits links, it is still possible to count the fraction of prohibited turns in the network

$ V $	TBTP	Up/Down	Spanning Tree
32	0.29	0.31	0.91
56	0.28	0.30	0.91
88	0.28	0.30	0.91
120	0.28	0.30	0.91
152	0.27	0.30	0.91

TABLE I
FRACTION OF PROHIBITED TURNS FOR NETWORKS OF VARYING SIZE AND
FIXED DEGREE $d = 8$.

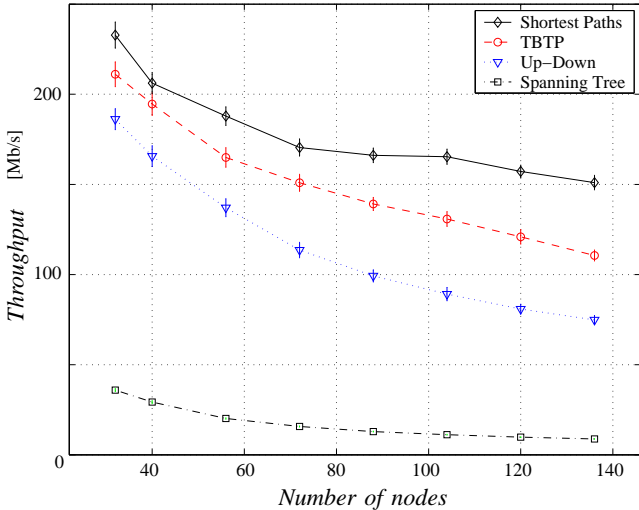


Fig. 4. Throughput versus increasing number of nodes $|V|$. The degree of each node is $d = 8$.

Next, we evaluate the effect of the graph density on the performance of the algorithms. We vary the degree of the nodes d , but keep the number of nodes fixed to $|V| = 120$. Table II provides evidence on the scalability of the turn-prohibition techniques, as compared to the spanning tree. While the fraction of prohibited turns prohibited by the TBTP and Up/Down algorithms increase slightly with the degree d , the spanning tree approach experiences a much sharper increase.

Figure 5 shows the throughput performance of the various algorithms as a function of the node degree. We note that increasing the degree of nodes implies an addition of links in the network. The spanning tree algorithm is the only algorithm that does not benefit from this increase. The reason is that a spanning tree permits the use of only $|V| - 1$ links in the network, independently of the topology.

Figure 5 also depicts the performance of the Turn-Prohibition (TP) algorithm developed in [12]. This algorithm prohibits at most 1/3 of the turns in any network, but is not backward-compatible with the IEEE 802.1d protocol. We observe that the performance of the new TBTP algorithm lies in-between those of the TP and Up/Down algorithms. Moreover, the performance of TBTP gets closer to TP as graphs become denser. Thus, the TBTP algorithm shows a good trade-off between the performance of the TP approach and the constraint imposed by permitting all the turns along a given spanning tree.

d	TBTP	Up-Down	Spanning Tree
4	0.23	0.25	0.74
6	0.27	0.29	0.86
8	0.28	0.30	0.91
10	0.28	0.31	0.93
12	0.29	0.31	0.95

TABLE II

FRACTION OF PROHIBITED TURNS FOR NETWORKS OF FIXED SIZE, $|V| = 120$, AND NODES OF VARYING DEGREE d .

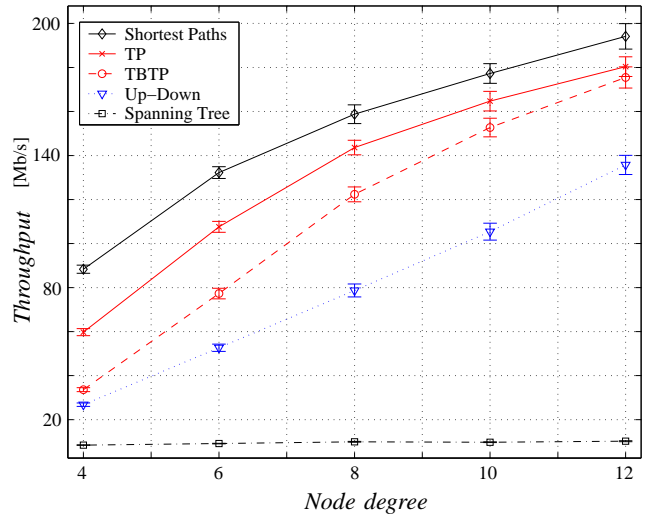


Fig. 5. Throughput versus increasing degree d , for $|V| = 120$ nodes.

2) *Partial Replacement of the Switches*: We now evaluate and compare the heuristics for node replacement described in Section IV. Simulations are run for random graphs of $|V| = 120$ nodes and degree $d = 8$. For each graph G , a spanning tree $T(G)$ is first constructed. Then legacy nodes are replaced according to one of the heuristics, that is, either the random or the Top-Down approach.

Figure 6(a) plots the throughput as a function of the number of intelligent nodes in the network, for each of the two heuristics. The results depicted in that figure are for the case where intelligent nodes implement the TBTP algorithm. As expected, the figures shows that the Top-Down heuristic leads to a faster improvement in performance than the random heuristic. For instance, when 60 out of the 120 network nodes are intelligent, Top-Down achieves a throughput roughly 50% higher than that obtained by the random approach. For both heuristics, the marginal gain in the throughput increases with the number of intelligent nodes.

Figure 6(b) shows similar results when intelligent nodes implement the Up/Down algorithms, though the difference between the two heuristics is less significant in this case.

B. Packet Level Simulations

In this section, we present packet level simulation results obtained with the NS2 simulator [22]. These simulations allow us to estimate the average end-to-end delay of packets as a function of the traffic load, for each of the cycle-breaking methods.

We consider two sample topologies. The first one, referred to as graph G_1 , is similar to that adopted for the flow level simulation and consists of a randomly generated graph with 32 nodes of degree 8. The second graph, G_2 , has also 32 nodes. It is generated using the BRITe topology generator [23], based on the so-called Barabási-Albert model [24]. This model captures the power-law node degree distribution that has been observed in a variety of networks.

Our traffic model is as follows. Each node establishes an UDP session (flow) with four other nodes, picked at random

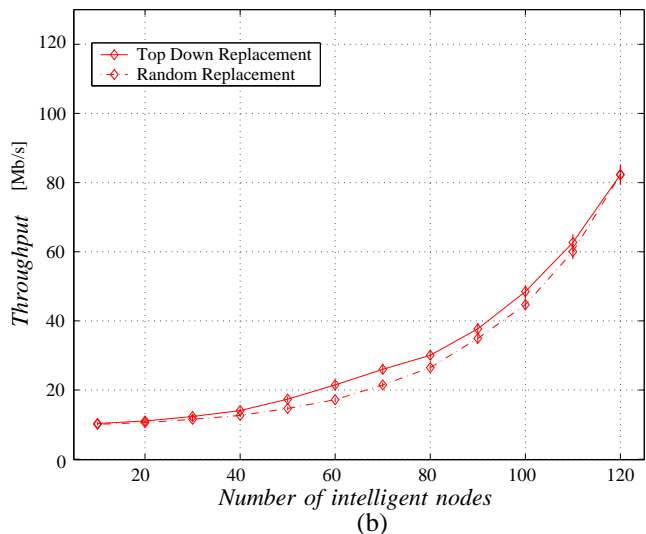
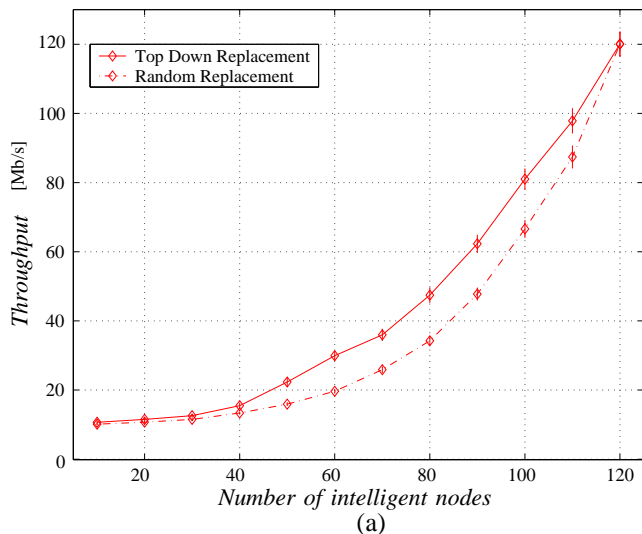


Fig. 6. Throughput versus number of intelligent nodes (total number of nodes is fixed at $|V| = 120$): (a) TBTP (b) Up-Down.

in the network. Each UDP session generates traffic according to an exponential ON-OFF source model. The average ON and OFF periods are 460 ms and 540 ms, respectively. The average traffic rate (in bit/s) is a simulation parameter and denoted by λ . The size of each packet is 600 bytes long. Each link has a propagation delay of 0.5 ms and capacity of 1 Mb/s. Note that due to memory and computation constraints, we could not simulate gigabit links. We conjecture, however, that the simulation results would have been qualitatively similar to those presented here.

The results for the end-to-end delay are presented in Figure 7. The results obtained for the two sample graphs are similar. We observe that the average end-to-end delay incurred with the spanning tree algorithm is always higher than with the two turn-prohibition approaches. Moreover, the maximum sustainable throughput, i.e., the traffic rate value λ at which the end-to-end delay starts to diverge, is increased by a factor of about five when the turn-prohibition techniques are employed. The TBTP algorithm achieves a higher throughput than Up/Down, as the latter prohibits a larger number of turns.

Finally, it is interesting to compare Figures 4 and 7, for the case of $|V| = 32$ nodes. We observe that flow level simulations predict well how each scheme performs one with respect to the other (we remind that the flow-level simulations use 1 Gb/s links). This result validates the use of flow level simulations as a fast and reliable method to predict network performance.

VI. CONCLUSION

In this paper, we have addressed the problem of breaking cycles in a scalable manner in Gigabit Ethernet networks. This problem has gained particular importance recently with the emergence of large metropolitan area networks (MANs), based on the Gigabit Ethernet technology. Toward this goal, we have proposed, analyzed and evaluated a novel cycle-breaking algorithm, called Tree-Based Turn-Prohibition (TBTP). This polynomial-time algorithm guarantees the prohibition of at most $1/2$ of the turns in the network, while permitting all

the turns belonging to a pre-determined spanning tree. We have shown that the algorithm can be generalized to the case where turns are assigned non-uniform weights (these weights can be set according to any metric of choice). In that case, the TBTP algorithm prohibits at most $1/2$ of the total weight of turns in the network. This generalization is especially useful for networks with links of non-uniform capacity.

We have also presented a general framework for incrementally deploying TBTP-capable switches in a way that is backward-compatible with the existing IEEE 802.1d standard. Since several paths may exist between any source-destination pair, we have described a method to perform packet forwarding (routing) in a heterogeneous network consisting of intelligent and legacy switches.

The performance of the proposed algorithm was thoroughly evaluated using both flow-level and packet-level simulations. The simulation showed that the TBTP algorithm achieves a throughput that is about an order magnitude higher than that obtained with the spanning tree standard. Furthermore, for a wide range of topology parameters, the performance of the TBTP-algorithm differs only by a small margin from that of shortest path routing (which achieves the highest possible throughput in theory, but does not break cycles).

Finally, we have proposed a heuristic, called Top/Down, to determine the order in which legacy switches should be replaced. This scheme proceeds by first replacing the root node, then nodes at level 1, and so forth. We have shown that Top/Down replacement outperforms a scheme where legacy nodes are replaced at random, achieving a throughput 50% higher in some cases. An interesting open research area is to investigate whether this strategy might be further improved and devise other possible replacement schemes.

REFERENCES

- [1] H. Frazier and H. Johnson, "Gigabit Ethernet: from 100 to 1,000 Mbps," *IEEE Internet Computing*, Vol. 3, No. 1, pp. 24–31, Jan./Feb. 1999.
- [2] H. Frazier, "The 802.3z Gigabit Ethernet Standard," *IEEE Network*, Vol. 12, No. 3, pp. 6–7, May/June 1998.

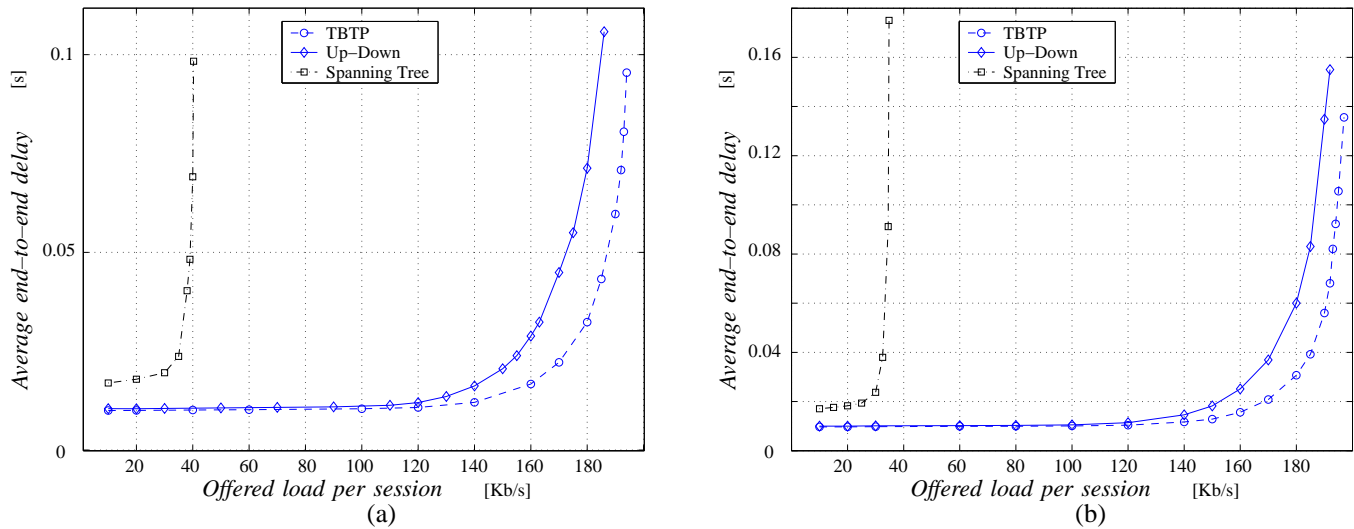


Fig. 7. Average end-to-end delay versus offered load λ , (a) Sample graph G_1 (b) Sample graph G_2 .

- [3] D. Clark, "Are ATM, Gigabit Ethernet Ready for Prime Time?," *IEEE Computer*, Vol. 31, No. 5, pp. 11–13, May 1998.
- [4] W. Nourreddine and F. Tobagi, "Selective Backpressure in Switched Ethernet LANs," in the proceedings of *Globecom '99*, Vol. 2, pp. 1256–1263, Dec. 1999.
- [5] R. Perlman, *Interconnections*, Addison-Wesley, 2000.
- [6] M. Karol, S. Golestani, and D. Lee, "Prevention of Deadlocks and Live-locks in Lossless, Backpressured Packet Networks," in the proceedings of *INFOCOM 2000*, Vol.3, pp. 1333–1342, March 2000, Tel Aviv, Israel.
- [7] J. Duato, "A Necessary and Sufficient Condition for Deadlock-Free Routing in Cut-Through and Store-and-Forward Networks," *IEEE Transactions on Parallel and Distributed Systems*, Vol. 7, No. 8, pp. 841–854, August 1996.
- [8] M. Soha and R. Perlman, "Comparison of Two LAN Bridge Approaches," *IEEE Network*, Vol.2 No 1, pp. 37–48, January 1988.
- [9] "Metro Ethernet Networks—A Technical Overview," Metro Ethernet Forum White Paper, July 2002. Available online at: <http://www.metroethernetforum.org/>
- [10] L. Bosack and C. Hedrick, "Problems in Large LANs," *IEEE Network*, Vol.2 No 1, pp. 52–56, January 1988.
- [11] D. Starobinski and M. Karpovsky, "Turn Prohibition-based Packet Forwarding in Gigabit Ethernets," in the proceeding of the *Gigabit Networking Workshop (GBN 2001)*, Anchorage, Alaska, April 2001. <http://www.comsoc.org/tcgn/conference/gbn2001/staro-presentation.pdf>
- [12] D. Starobinski, M.G. Karpovsky and L. Zakrevski "Application of Network Calculus to General Topologies using Turn-Prohibition," *IEEE/ACM Transactions on Networking*, to appear.
- [13] L.B. Levitin, M.G. Karpovsky, "Deadlock Prevention in Networks Modeled as Weighted Graphs," in the proceedings of *ICINSAT-2002*, pp. 42–47, 2002.
- [14] C. Glass and L. Ni, "The Turn Model for Adaptive Routing," *Journal of ACM*, Vol. 41, No. 5, pp. 874–902, September 1994.
- [15] M. Shoreder et al., "Autonet: A High-Speed, Self-Configuring Local Area Network Using Point-to-Point Links," *IEEE Journal on Selected Areas in Communications*, Vol. 9, No. 8, pp. 1318–1335, October 1991.
- [16] R. Perlman, "An Algorithm for Distributed Computation of a Spanning Tree in an Extended LAN", in the proceedings of Ninth ACM Data Communications Symposium, Vol. 20, No. 7, pp. 44–52, September 1985, New York, USA.
- [17] J. Hart, "Extending the IEEE 802.1 MAC Bridge Standard to Remote Bridges," *IEEE Network*, Vol.2 No 1, pp. 10–15, January 1988.
- [18] Perlman et al., "Utilization of redundant links in bridged networks," United States Patent Number 5,150,360, September 22, 1992.
- [19] K. Lui, W. C. Lee, K. Nahrstedt, "STAR: a transparent spanning tree bridge protocol with alternate routing," *ACM SIGCOMM Computer Communication Review*, vol. 32, no. 3, pp. 33–46, July 2002.
- [20] T.L. Rodeheffer, C.A. Thekkath, D.C. Anderson, "SmartBridge: A Scalable Bridge Architecture," in the proceedings of *ACM SIGCOMM 2000*, September 2000, Stockholm, Sweden.
- [21] T. Cormen, C. Leiserson, and R. Rivest, *Introduction to algorithms*, McGraw-Hill, 1990.
- [22] The Network Simulator - ns-2, available online at <http://www.isi.edu/nsnam/ns/>.
- [23] Boston University Representative Topology Generator - BRITE, available online at <http://www.cs.bu.edu/brite/>.
- [24] A.-L. Barabási and R. Albert, "Emergence of Scaling in Random Networks," *Science*, No. 286, pp. 509–512, 1999.