

Header-size lower bounds for end-to-end communication in memoryless networks [☆]

Pierre Fraigniaud ^{a,*}, Cyril Gavoille ^b

^a CNRS, Lab. de Recherche en Informatique, University Paris-Sud, Batiment 490, 91405 Orsay Cedex, France

^b University of Bordeaux 1, LaBRI, 33405 Talence Cedex, France

Available online 25 October 2005

Abstract

The end-to-end communication problem is a protocol design problem, for sending a packet from a specified source-node s to a specified target-node t , through an unreliable asynchronous memoryless communication network. The protocol must insure reception and termination. In this paper, we measure the complexity of the protocol in term of header size, i.e., the quantity of information that must be attached to the packets to insure their delivery. We show that headers of $\Omega(\log \log \tau)$ bits are required in every network, where τ denotes the treewidth of the network. In planar networks, $\Omega(\log \tau)$ bits are required.

© 2005 Elsevier B.V. All rights reserved.

Keywords: End-to-end; Sequence transmission; Treewidth

1. Introduction

The *end-to-end communication* problem is the problem of sending a (sequence of) packet(s) from a specified source-node s to a specified target-node t , through an *unreliable* communication network G

(see, e.g., [13,29]). The *sequence transmission* [38] problem and the *reliable communication* [24] problem are other names for the end-to-end communication problem (cf. the survey [25]). By an unreliable network, it is generally meant that links can lose, reorder, and duplicate packets. Moreover, networks are assumed to be *asynchronous*, i.e., the time for a packet to traverse a link is finite but otherwise unbounded. In particular, a processor cannot distinguish between an inoperational link and an operational link which is just very slow. Hence, an instance of the end-to-end communication problem is described by an (unreliable and asynchronous) network, modeled by an undirected graph G , and two nodes s and t of G . Solving the problem consists in designing a distributed protocol which (1) allows s to send a packet, or a sequence of packets, to t

[☆] The two authors are supported by the Actions Spécifiques CNRS “Dynamo” and “Algorithmique des grands graphes”, and by the project “PairAPair” of the ACI Masses de Données. P. Fraigniaud receives additional support from the INRIA project “Grand Large”. A preliminary version of this paper appeared in the proceedings of the 17th International Symposium on Distributed Computing (DISC), Sorrento, Italy, October 1–3, 2003.

* Corresponding author.

E-mail addresses: pierre@lri.fr (P. Fraigniaud), gavoille@labri.fr (C. Gavoille).

through the network G , and (2) generates a finite amount of traffic for each packet. In other words, the end-to-end protocol must satisfy the two following requirements: (1) *reception*, i.e., the target must eventually receive at least one copy of each packet sent by the source; and (2) *termination*, i.e., after a finite time, no copy of the packet(s) remains in the network.

1.1. The oblivious single-packet end-to-end communication problem

This paper considers a static model, that is, we assume that each link is either operational or not. If there exists at least one operational path from s to t in G , it is the role of the protocol to find such a non-faulty path, which is of course a priori unknown. Note that the case of dynamic faults, that is when links can alternate between being operational and inoperational, can be treated similarly by assuming, as in [1,14,28], infinitely frequent path stability, i.e., infinitely often there is a path P from s to t such that a packet sent from s along P will arrive at t (see also [3]).

Performance of end-to-end communication protocols is commonly measured in terms of (1) the amount of communication performed over the links of the network, and (2) the amount of storage space used by intermediate nodes in the networks. *Oblivious* protocols, a.k.a. *memoryless* protocols, take their routing decision at every node x (i.e., on which link(s) x has to forward a packet) based solely on the content of the header of the packet. In particular, a node does not store any knowledge about the traffic that previously passed through it. It can however forward several copies of the received packet, and it can modify the header of this packet. As mentioned in [1], the practical advantage of oblivious protocols is their high tolerance to processor crashes. Indeed, as soon as a node recovers from a crash, the protocol is ready to proceed with no risk of corruption due to an altered writable memory (RAM). Moreover, the extremal behavior of oblivious protocols (in the sense that they consume no local memory at all) allows concentrating the analysis of end-to-end communication protocols on the amount of information transmitted over the links of the network. More specifically, this paper focuses on minimizing the *packet-header size*, i.e., the quantity of additional information that must be attached to every packet to insure its correct delivery.

The header content has two distinct functions: (1) to control the order in which packets arrive at the destination; (2) to find a route from the source s to the destination t (reception), and to insure that residual packets are eventually removed from the network (termination). In this paper, we are interested in the routing part of the problem, that is, in the problem of finding the non-faulty route from the source to the destination, while insuring termination. We will therefore concentrate our analysis on the process of sending a single packet from s to t . In other words, we consider the *single-packet* end-to-end communication problem, as opposed to the *stream-of-packets* end-to-end communication problem, the latter problem requiring the transmission of a sequence of packets from the source to the destination [22]. Hence, let us summarize our problem.

1.1.1. Our problem

We are given an unreliable and asynchronous network G , and two nodes s and t of G . We consider the design of an oblivious distributed protocol which allows the transmission of a packet from s to t (if there is a fault-free path between s and t in G), and which eventually leaves the network empty of packets. Such a protocol is required to use packet-headers of small size. The quality of the protocol is indeed measured by the maximum size of the headers involved during its execution.

1.1.2. Previous work

In the context of static link failures (i.e., every link is operational or not but its status does not change during the execution of the routing protocol), the *hop-count* protocol [32] uses headers of size $O(\log n)$ in n -node networks. It proceeds by flooding the network as follows. The source sends a copy of the original packet to all its neighbors, with header 1. A node receiving a packet whose header contains the hop count $i < n - 1$ updates the header by replacing i by $i + 1$, and forwards a copy of the packet to each of its neighbors. A node receiving a packet whose header contains the hop count $n - 1$ removes the packet from the network. If s and t are connected despite the faulty links, then a path of length at most $n - 1$ exists between s and t , and therefore at least one copy of the packet sent by s eventually arrives at t . Moreover, the remaining copies of the packet are removed from the network after a finite time since no packet can traverse n or more links.

In [1], Adler and Fich showed that if a network G contains H as a *minor*, i.e., if H can be obtained from G by edge contraction, node and edge deletion, then the oblivious single-packet end-to-end communication problem in G requires headers of size at least as large as for H . Hence the problem is closed under taking minors. They also showed that the complete graph of n nodes requires headers of size $\Omega(\log n)$. Therefore, an n -node network G which contains the complete graph of k nodes as a minor requires headers of size $\Omega(\log k)$. Adler and Fich also gave upper bounds on the header size based on the notion of *feedback vertex sets*. A feedback vertex set is a subset of nodes S such that every cycle in the network contains at least a node in S . They showed that if there exists a feedback vertex set of size f , then there exists an oblivious protocol using headers with size $O(\log f)$. They also pointed out that any minimum feedback vertex set of an $\sqrt{n} \times \sqrt{n}$ mesh is of size $\Omega(n)$, and thus the feedback vertex set protocol does not offer significant improvement in the mesh, compared to the hop-count protocol. Since the mesh is planar, it does not contain K_5 as minor, and thus it leaves a big gap between the best known upper and lower bounds for mesh networks. Adler et al. [2] recently closed this gap for “narrow” meshes, i.e., $p \times q$ meshes with $p = O(1)$. They have shown that there exists a protocol for n -node meshes using headers of size $O(p(\log p + \log \log n))$, and that any $p \times q$ mesh requires headers of size $\Omega(\log \log n)$ for $3 \leq p \leq q$. A constant size header protocol exists if $p \leq 2$.

1.2. Outline of our results

In this paper, we provide lower bounds of the header size of oblivious end-to-end protocols as a function of the treewidth of the input graph. The notion of treewidth is a very rich concept with many algorithmic implications. Many NP-complete problems are polynomial for graphs of bounded treewidth [10,16]. In addition, the notion of treewidth is in the kernel of the graph minor theory [20]. Recall that treewidth can be seen as a measure of “how far” a graph is from a tree. More precisely, let $G = (V, E)$ be a graph of n nodes, a *tree-decomposition* of G is a pair (\mathcal{S}, T) where \mathcal{S} is a collection $\{S_i | i \in I\}$ of subsets of nodes and $T = (I, F)$ is a tree, such that the three following conditions are satisfied:

1. $\cup_{i \in I} S_i = V$;
2. For every edge $e = \{u, v\}$ of G , there exists $i \in I$ such that both u and v belong to S_i ;
3. For every $u \in V$, the subgraph of T induced by the set of nodes $\{i \in I | u \in S_i\}$ is a tree.

For any graph G , there exists at least one tree-decomposition of G by choosing $\mathcal{S} = \{V\}$, i.e., T is reduced to a single node. The *width* of a tree-decomposition (\mathcal{S}, T) is defined as $\max_{i \in I} |S_i| - 1$. The *treewidth* $\text{tw}(G)$ of G is then defined as the minimum of the width of any tree-decomposition of G . Roughly speaking, $\text{tw}(G)$ denotes the minimum size c such that G has a recursive separator of at most c nodes. A c -decomposable graph has treewidth $O(c)$, a tree has treewidth 1, a cycle has treewidth 2, a square mesh has treewidth \sqrt{n} , and a complete graph has treewidth $n - 1$. Determining the treewidth of a graph is NP-hard [11]. However, Bodlaender [17] gave a linear time algorithm for recognizing graphs of bounded treewidth. There are also $O(\log n)$ -approximation algorithms for computing the treewidth of an arbitrary graph [18,33], and even $O(\log \tau)$ -approximation algorithms where τ is the treewidth [9,19].

In this paper, we provide a lower bound on the header size for end-to-end communication as a function of the treewidth of the network.

Theorem 1. *Any protocol designed for the instance (G, s, t) of the oblivious single-packet end-to-end communication problem uses headers of size at least $\Omega(\log \log \tau)$ bits where τ is the treewidth of the graph $G_{s,t}$ obtained from G , s , and t by deleting every edge e not on a simple path from s to t .*

As we will see in more details in Section 3, the bound of Theorem 1 derives from an optimal bound for square meshes, and from an upper bound of the “excluding grid” theorem of Robertson and Seymour [35]. This latter bound is likely far from best possible and, as mentioned in [21], Robertson et al. [36] think that the upper bound might be exponentially improved so that the lower bound of Theorem 1 would be $\Omega(\log \tau)$. In fact, independently from whether the bound of the excluding grid theorem can be exponentially improved, we believe that end-to-end protocols with headers size smaller than $\Omega(\log \tau)$ are unlikely to exist. This belief is supported by the fact that, in order to insure termination, any protocol seems to be required to compute how many times any given message crosses

a separator of the network. Hence we conjecture the following:

Conjecture. *Any protocol designed for the instance (G, s, t) of the oblivious single-packet end-to-end communication problem uses headers of size at least $\Omega(\log \tau)$ bits where τ is the treewidth of the graph $G_{s,t}$.*

If this conjecture holds, then the hop-count protocol will turn out to be asymptotically optimal in large classes of graphs (namely those graphs with treewidth $\Omega(n^\epsilon)$ for any $\epsilon > 0$). Actually, for planar graphs, we prove the lower bound $\Omega(\log \tau)$:

Theorem 2. *Any protocol designed for the instance (G, s, t) of the oblivious single-packet end-to-end communication problem, G planar, uses headers of size at least $\Omega(\log \tau)$ bits where τ is the treewidth of the graph $G_{s,t}$.*

In particular Theorem 2 solves the conjecture mentioned in [1] stating that headers of $\Omega(\log n)$ bits are required to insure packet-delivery in the two dimensional square mesh. All these results establish (as also suggested in [1]) connections between, on one hand, the number of header bits needed to send a single packet through an unreliable network G , and, on the other hand, graph-theoretic properties of G 's topology.

Finally, notice that our lower bounds obviously apply to the stream-of-packets problem too. Moreover, they also obviously apply to the case of dynamic faults, that is when links can alternate between being operational and inoperational. However, they do not necessarily apply to the models in [3,4] that do not assume memoryless networks, and for which protocols use the fact that links can be in different states.

1.3. Related work

The case of links with fixed or bounded traversal time has been considered in [8,41]. Probabilistic faults and delivery times have been considered in [27,30,31]. In the deterministic setting, i.e., the context of this paper, solutions for the stream-of-packets end-to-end problem differ according to the type of faults. Wang and Zuck [40] have shown that any protocol tolerating both packet reordering and duplication requires unbounded headers. Afek et al. [6] have shown that packet reordering and loss create the same effect, that is either unbounded headers or non-termination (i.e., the same packet

can be received an unbounded number of times). Fekete and Lynch [23] have shown that just packet loss implies that some header information must be attached to the packets for the stream of packets to be treated correctly. These three latter results hold even if G consists of a single edge $\{s, t\}$. Despite these impossibility results, efficient protocols have nevertheless been successfully designed. We refer to [12,37,39] if links are subject to packet duplication, reordering and loss, to [15,30] if links are subject to packet duplication and loss (but no reordering), and to [5,7,29] in case of static link failures. In [1], it is noticed that the latter protocols can be adapted to the case where links can lose packets—but otherwise transmit them in order and without duplication, by using the techniques in [4]. Although some of these protocols use very short headers, none of them is oblivious, i.e., they all require the local storage of information at intermediate nodes during the execution of the protocol. The stream-of-packet protocol of Dolev and Welch [22] is oblivious (i.e., intermediate processors do not change state) and apply to the static link failure (i.e., the model of this paper). It uses headers of $O(\log p)$ bits where p is the number of distinct simple paths between s and t in G . Although p can be quite large in general (e.g., $\Omega((n-2)!) in K_n), this protocol was proved optimal for many topologies, including complete graphs, meshes, and series-parallel graphs (see [1]). Fich and Jakoby [26] considered the same model for directed acyclic graphs. They proved that a single bit header suffices in DAGs, which contrasts with the case of arbitrary graphs.$

Several oblivious protocols have been proposed in the literature to solve the single-packet end-to-end problem. We already presented the *hop-count* protocol described by Postel [32], using headers of $O(\log n)$ bits in n -node graphs. We also mentioned the protocol of Dolev and Welch [22] which applies to the single-packet problem as well. Adler and Fich [1] derived lower and upper bounds for the header size in specific networks (e.g., meshes, hypercubes, butterflies, de Bruijn, etc.). Adler et al. [2] addressed the problem in $p \times q$ meshes, $p = O(1)$. They proved the somewhat surprising result stating that headers of size $\Theta(\log \log n)$ bits are necessary and sufficient in $3 \times n/3$ meshes (whereas headers of constant size are sufficient in $2 \times n/2$ meshes).

For a more detailed descriptions of the end-to-end protocols mentioned above, we refer to [25], and the references therein.

2. End-to-end routing in meshes

We start by proving the following lemma which is an extension of a result by Adler and Fich [1]. The result in [1] is valid for complete graphs. We extend it to arbitrary graphs. Then, in Theorem 3, we show how to apply this result to meshes.

Lemma 1. *Let $e_i = \{x_i, y_i\}$, $i = 1, \dots, k$, be k distinct edges of G . Let $\Pi \subseteq \Sigma_k$ be a subset of permutations of k symbols such that, for any $\pi \in \Pi$, there is a simple path P_π in G from s to t , traversing the e_i 's from the x_i 's to the y_i 's, in the order $e_{\pi(1)}, e_{\pi(2)}, \dots, e_{\pi(k)}$. Then any oblivious end-to-end communication protocol from s to t requires headers of at least $(\log_2 |\Pi|)/k$ bits.*

Proof. Given an oblivious end-to-end communication protocol A from s to t , let h be the total number of distinct headers involved in protocol A . We use the same terminology as in [1]. The *transcript* of a simple path P_π is defined as the word w_π of length k on the alphabet $\{0, \dots, h-1\}$ obtained by concatenating the headers of the packets transmitted through the e_i 's, from the x_i 's to the y_i 's, when only the edges of P_π are operational. More precisely, let us consider a packet b_π going from s to t along P_π . This packet generates a word w_π whose i th letter is the value of the header h_i of b_i when traversing the edge e_i from x_i to y_i . This word is the transcript of P_π . (Obviously, the transcript may depend on the choice of the packet b_π . This packet is chosen arbitrarily.) There are at most h^k transcripts, and therefore, if $h^k < |\Pi|$, then at least two paths have the same transcript. Thus assume, for the purpose of contradiction, that $h^k < |\Pi|$, and let P_π and $P_{\pi'}$ be two paths with the same transcript $h_1 h_2 \dots h_k$. Since $\pi \neq \pi'$, there is a pair (i, j) , $i \neq j$, such that e_i is traversed before e_j in P_π , and e_i is traversed after e_j in $P_{\pi'}$.

Assume now that both paths P_π and $P_{\pi'}$ are operational. Since the protocol has no way to distinguish an operational link from a link that is just very slow, and since the protocol is oblivious, the headers of packets b_π and $b_{\pi'}$, respectively following P_π and $P_{\pi'}$, will be the same as those used to traverse these two paths when only one of them is operational. In particular, when y_i receives a packet from x_i with header h_i , it cannot know whether P_π , or $P_{\pi'}$, or both paths, are operational. Therefore, y_i forwards the packet through both paths. The same holds for y_j .

As a consequence, the following situation occurs. A packet traverses e_i with header h_i , eventually

reaches x_j along P_π , traverses e_j with header h_j , eventually reaches x_i along $P_{\pi'}$, traverses e_i again, still with header h_i , and so on. This creates an infinite loop, in contradiction with the termination requirement. Therefore $h^k < |\Pi|$ cannot hold, and thus $h \geq |\Pi|^{1/k}$. \square

The main theorem is the following:

Theorem 3. *Any oblivious end-to-end communication protocol in the $p \times q$ mesh requires headers of at least $\Omega(\log \min\{p, q\})$ bits.*

Proof. Let $M_{p,q}$ be the $p \times q$ mesh. Assume, w.l.o.g. (without loss of generality), that $p \leq q$ (otherwise exchange the role of p and q). We see the mesh $M_{p,q}$ as with p rows and q columns. Rows are labeled from 0 to $p-1$, and columns from 0 to $q-1$. One can draw the mesh so that $(0,0)$ is the top-left corner, and $(p-1, q-1)$ is the bottom-right corner. Let s and t be the source-node and the target-node, respectively. The $p \times q$ mesh contains three node-disjoint $(p-2) \times \lfloor \frac{q-2}{3} \rfloor$ sub-meshes. At least one of these sub-meshes does not contain s nor t . Let M be this sub-mesh. One can construct two disjoint paths P_s and P_t in $M_{p,q} \setminus M$, respectively leading from s to the “top-left” corner of M , and from t to the “top-right” corner of M . Any oblivious end-to-end protocol must perform successfully even if all links in $M_{p,q} \setminus (M \cup P_s \cup P_t)$ are faulty. Hence one can assume, w.l.o.g., that s is node $(0,0)$ and t is node $(0, q-1)$.

Let $k = \lfloor \min\{\frac{p-1}{4}, \sqrt{\frac{q}{12}}\} \rfloor$. Hence $p \geq 4k+1$ and $q \geq 12k^2 \geq 4k^2 + 6k + 2$. Let us consider the edges $e_i = \{x_i, y_i\}$, $i = 1, \dots, k$, where $x_i = (2k, i(4k+2) - 1)$ and $y_i = (2k, i(4k+2))$. Fig. 1 illustrates the idea of the construction: eight edges are displayed on row $2k$ (the horizontal and vertical scales are different). From the setting of k , there are exactly $2k$ rows above row $2k$ (which contains all the e_i 's), and at least $2k$ rows below row $2k$. Also, there are exactly $4k$ columns separating column 0 from e_1 , $4k$ columns separating two consecutive e_i 's, and at least $4k$ columns separating column $q-1$ from e_k .

Let $\pi \in \Sigma_k$ be any permutation of k symbols. Let us show that there is a simple path from s to t , passing through e_i from x_i to y_i for every i , in the order $e_{\pi(1)}, e_{\pi(2)}, \dots, e_{\pi(k)}$. To each $e_{\pi(i)}$ are associated four rows and $4k$ columns. The rows associated to $e_{\pi(i)}$ are rows

$$U_i = 2i - 2, \quad U'_i = 2i - 1,$$

$$D_i = 2k + 2i - 2, \quad \text{and} \quad D'_i = 2k + 2i - 1.$$




Fig. 1. Associated rows and columns.

Fig. 1 shows the four rows associated to some $e_{\pi(i)}$. U stands for “up”, and D for “down”. The dashed rows are examples of rows associated to some $e_{\pi(j)}$, $j > i$. Note that, since $p \geq 4k + 1$, row $p - 1$ is not associated to any $e_{\pi(i)}$. Let $c_i = i(4k + 2) - 1$ be the column coordinate of x_i , $i = 1, \dots, k$. The columns associated to $e_{\pi(i)}$ are columns

$$I_{i,j} = c_j - 2i + 1, \quad O_{i,j} = c_j - 2i, \\ I'_{i,j} = c_j + 2i - 1, \quad \text{and} \quad O'_{i,j} = c_j + 2i, \quad j = 1, \dots, k.$$

Fig. 1 shows the $4k$ columns associated to some $e_{\pi(i)}$. I stands for “inside”, and O for “outside”. The dashed columns are examples of columns associated to some $e_{\pi(j)}$, $j > i$. There is no overlapping between associated columns. Moreover, columns $0, \dots, 2k$ and $4k^2 + 4k, \dots, q - 1$ are not associated columns.

Now, for $i = 1, \dots, k$, we define the path P_i as follows. (The construction is illustrated in Fig. 2

where the integer displayed above an edge indicates the order of the edge in the permutation π .)

P_1 starts from node $s = (0, 0)$ following the row U_1 until it reaches column $I_{1,j}$ where $j = \pi(1)$. Then P_1 follows $I_{1,j}$ downward until row $2k$. At this point P_1 follows row $2k$ rightward until it traverses edge $e_{\pi(1)}$. At node y_j , P_1 goes down to reach row $2k + 1$, and follows this row leftward until column $O_{1,j}$. Then it goes upward on $O_{1,j}$ until row U'_1 , and finally follows this row until node $(1, 0)$ where it ends.

P_i starts from node $(2(i - 1), 0)$. Let l_1, l_2, \dots, l_h be the indices $l < i$ such that $e_{\pi(l)}$ is on the left of $e_{\pi(i)}$. For instance, in Fig. 2, this sequence is $2, 1, 3$, for $e_{\pi(4)}$. Up to relabeling, one can assume, w.l.o.g., that $e_{\pi(l_j)}$ is on the left of $e_{\pi(l_{j+1})}$. Let $j = \pi(i)$. P_i goes rightward from node $(2(i - 1), 0)$ along row U_i until it reaches column I_{i, l_1} . Then it goes downward until it reaches row D_i , and goes right. P_i follows D_i until column I'_{i, l_1} , and then goes upward along that




Fig. 2. The path P_s .

column, until it reaches row U_i . This “detour” around the edges $e_{\pi(i),s}$ is repeated for each $1 \leq a \leq h$. (See for instance the path P_4 in Fig. 2.) More precisely, let $a \in \{1, \dots, h\}$. P_i remains on U_i until it reaches column I_{i,l_a} . Then it goes downward until it reaches row D_i , and goes right. P_i follows D_i until column I'_{i,l_a} , and then goes upward along that column, until it reaches row U_i again.

After the last detour around $e_{\pi(i),s}$, P_i goes rightward on U_i until it reaches $I_{i,j}$. Then it goes downward on $I_{i,j}$ until it reaches row $2k$, and turn rightward along that row. P_i follows row $2k$ until it traverses $e_j = e_{\pi(i)}$. At node y_j , P_i goes down to reach row $2k + 1$, and follows this row leftward until column $O_{i,j}$. Then it goes upward on $O_{i,j}$ until row U'_i . At this point, P_i starts its journey leftward on U'_i , back to column 0.

More precisely, P_i follows the same shape of path as when it went rightward from column 0 (see Fig. 2). That is, it proceeds along U'_i but takes detour around edges $e_{\pi(i),s}$. Let $a \in \{1, \dots, h\}$. A detour around $e_{\pi(i),s}$ consists in leaving U'_i to go downward along O'_{i,l_a} until row D'_i is reached. Then P_i follows D'_i leftward until column O_{i,l_a} , and follows this column upward until it reaches again U'_i . After the last detour (i.e., the detour around $e_{\pi(i),s}$), P_i follows row U'_i until node $(2i - 1, 0)$ where it ends.

The following two properties are satisfied:

- $\mathcal{P}1$ For every $1 \leq i \leq k$, P_i is a simple path.
- $\mathcal{P}2$ For every $i \neq j$, P_i and P_j have no node in common. More precisely, the paths fit one into another as displayed in Fig. 3. In this figure, $a < b < i < j$. P_a goes straight above e because there is no need of a detour as P_b is not set yet. P_b traverses e . Both P_i and P_j must go around e . By construction, since $i < j$, they follow different associated rows and columns and therefore do not intersect.




Fig. 3. Two paths P_i and P_j do not intersect.

Now, let P be the following path from s to t :

$$P = P_1, f_1, P_2, f_2, \dots, P_{k-1}, f_{k-1}, P_k, Q$$

where f_i is the edge between nodes $(2i - 1, 0)$ and $(2i, 0)$, and Q is the simple path from node $(2k - 1, 0)$ to t following column 0, then row $p - 1$, and finally column $q - 1$. From properties $\mathcal{P}1$ and $\mathcal{P}2$, P is a simple path from s to t which traverse edges e_i s in the order $e_{\pi(1)}, e_{\pi(2)}, \dots, e_{\pi(k)}$.

The construction applies for any permutation $\pi \in \Sigma_k$. Therefore, thanks to Lemma 1, any oblivious end-to-end communication protocol in the $p \times q$ mesh requires headers of at least $\Omega(\log(k!)/k) = \Omega(\log k)$ bits. Since $k = \lfloor \min\{\frac{p-1}{4}, \sqrt{\frac{q}{12}}\} \rfloor$, we get $\log k = \Omega(\log \min\{p, q\})$ which completes the proof. \square

3. Proofs of Theorems 1 and 2

3.1. Proof of Theorem 1

We use the “excluding grid” theorem of Robertson and Seymour, whose short proofs can be found in [20,21]:

Theorem 4 (Robertson and Seymour [35]). *For every integer r there is an integer k such that every graph of treewidth at least k has an $r \times r$ mesh as minor.*

So, let us define $f(r)$ as the smallest integer k satisfying Theorem 4. The constructive proof of the excluding grid theorem, given in [21] shows that $f(r) \leq 2^{5r \log_2 r}$. Let τ be the treewidth of $G_{s,t}$, i.e., the graph obtained from G by deleting every edge not on a simple path from s to t . Let $r = \lfloor (\frac{1}{5} \log_2 \tau)^{1/6} \rfloor$. By the excluding grid theorem, $G_{s,t}$ contains an $r \times r$ mesh as minor since $\tau \geq f(r)$.

Now, any oblivious end-to-end protocol on $G_{s,t}$ requires headers of size at least as large as those required for the $r \times r$ mesh (recall that the end-to-end problem is closed under minor taking, as shown by Adler and Fich [1]). By Theorem 3, the headers for $r \times r$ meshes are of size at least $\Omega(\log r) = \Omega(\log \log \tau)$, which completes the proof of Theorem 1.

Remark. Any polynomial upper bound on $f(r)$ in r would prove a lower bound of $\Omega(\log \tau)$ on the header size. So far, the best known upper bound on $f(r)$ is $f(r) \leq 2^{9r^5}$ (cf. [36]). It is conjectured that the correct order of magnitude for $f(r)$ is $O(r^2 \log r)$.




Fig. 4. An 8×3 cylinder.

3.2. Proof of Theorem 2

For $r \geq 0$, and $c > 0$, two integers, the $r \times c$ cylinder is the graph composed of r radial lines and c circles as shown in Fig. 4.

Theorem 5 (Robertson and Seymour [34], p. 62). *If G is planar and has no $r \times r$ cylinder as minor, then G has treewidth at most $\frac{3}{2}(r^2 + 2r) - 2$.*

It is clear that an $r \times r$ cylinder contains an $r \times r$ mesh as minor. Let $\tau = \text{tw}(G_{s,t})$, and let $r = \lfloor \sqrt{\tau/5} \rfloor$ so that $\tau > \frac{3}{2}(r^2 + 2r) - 2$. Thus, by Theorem 5, $G_{s,t}$ contains the $r \times r$ mesh as minor. Therefore the header size for end-to-end communication in $G_{s,t}$ is at least the header size for end-to-end communication in the $r \times r$ mesh, that is at least $\Omega(\log r) = \Omega(\log \tau)$ by Theorem 3. This completes the proof of Theorem 2.

4. Conclusion

This paper provides lower bounds on the header size for oblivious end-to-end communication in both arbitrary and planar graphs. The design of matching upper bounds remains open.

The treewidth of a graph G can be alternatively defined as the minimum k for which G is a subgraph of a triangulated graph H that has maximum clique size $k + 1$. (Recall that a graph is triangulated if it contains no chordless cycle of length greater than three.) Solving the end-to-end communication problem in triangulated graphs of given maximum clique size may be a good way to solve the problem for arbitrary graphs.

References

- [1] M. Adler, F. Fich, The complexity of end-to-end communication in memoryless networks, in: 18th ACM Symposium on Principles of Distributed Computing (PODC), 1999, pp. 239–248.
- [2] M. Adler, F. Fich, L. Goldberg, M. Paterson, Tight size bounds for packet headers in narrow meshes, in: 27th Int. Colloquium on Automata, Languages and Programming (ICALP) LNCS 1853, Springer, 2000, pp. 756–767.
- [3] Y. Afek, B. Awerbuch, E. Gafni, Applying static network protocols to dynamic networks, in: 28th IEEE Annual Symposium on Foundations of Computer Science (FOCS), 1987, pp. 358–370.
- [4] Y. Afek, E. Gafni, End-to-end communication in unreliable networks, in: 7th ACM Symposium on Principles of Distributed Computing (PODC), 1988, pp. 131–148.
- [5] Y. Afek, E. Gafni, Bootstrap network resynchronization, in: 10th ACM Symposium on Principles of Distributed Computing (PODC), 1991, pp. 295–307.
- [6] Y. Afek, H. Attiya, A. Fekete, M. Fischer, N. Lynch, Y. Mansour, D.-W. Wang, L. Zuck, Reliable communication over unreliable channels, *Journal of the ACM* 41 (6) (1994) 1267–1297.
- [7] Y. Afek, B. Awerbuch, E. Gafni, Y. Mansour, A. Rosén, N. Shavit, Slide—The key to polynomial end-to-end communication, *Journal of Algorithms* 22 (1) (1997) 158–186.
- [8] A. Aho, J. Ullman, A. Wyler, M. Yannakakis, Bounds on the size and transmission rate of communication protocols, *Computers and Math with Applications* 8 (3) (1982) 205–214.
- [9] E. Amir, Efficient approximation for triangulation of minimum treewidth, in: 17th Conference on Uncertainty in Artificial Intelligence (UAI), Morgan Kaufmann, 2001, pp. 7–15.
- [10] S. Arnborg, Efficient algorithms for combinatorial problems on graphs with bounded decomposability—a survey, *BIT* 25 (1985) 2–23.
- [11] S. Arnborg, D. Corneil, A. Proskurowski, Complexity of finding embeddings in a k -tree, *SIAM Journal of Algebraic and Discrete Methods* 8 (1987) 277–284.
- [12] B. Awerbuch, S. Even, Reliable broadcast protocols in unreliable networks, *Networks* 16 (1986) 381–396.
- [13] B. Awerbuch, Y. Mansour, N. Shavit, Polynomial end-to-end communication, in: 30th IEEE Annual Symposium on Foundations of Computer Science (FOCS), 1989, pp. 358–363.
- [14] B. Awerbuch, O. Goldreich, A. Herzberg, A quantitative approach to dynamic networks, in: 9th ACM Symposium on Principles of Distributed Computing (PODC), 1990, pp. 189–203.
- [15] K. Bartlett, R. Scantleburg, P. Wilkinson, A note on reliable, full-duplex transmission over half-duplex links, *Communication of the ACM* 12 (1969) 260–261.
- [16] H. Bodlaender, A tourist guide through treewidth, *Acta Cybernetica* 11 (1993) 1–21.
- [17] H. Bodlaender, A linear-time algorithm for finding tree-decompositions of small treewidth, *SIAM Journal on Computing* 25 (1996) 1305–1317.
- [18] H. Bodlaender, J. Gilbert, H. Hafsteinsson, T. Kloks, Approximating treewidth, pathwidth, frontsize, and shortest elimination tree, *Journal of Algorithms* 18 (1995) 238–255.

- [19] V. Bouchitté, D. Kratsch, H. Müller, I. Todinca, On treewidth approximations, in: 1st Cologne-Twente Workshop on Graphs and Combinatorial Optimization (CTW), 2001.
- [20] R. Diestel, *Graph Theory*, second ed., Springer-Verlag, New York, 2000.
- [21] R. Diestel, T.R. Jensen, K.Y. Gorbunov, C. Thomassen, Highly connected sets and the excluded grid theorem, *Journal of Combinatorial Theory Series B* 75 (1999) 61–73.
- [22] S. Dolev, J. Welch, Crash resilient communication in dynamic networks, *IEEE Transactions on Computers* 46 (1) (1997) 14–26.
- [23] A. Fekete, N. Lynch, The need for headers: an impossibility result for communication over unreliable channels, in: 1st Int. Conf. on Theories of Concurrency: Unification and Extension (CONCUR) LNCS 458, Springer, 1990, pp. 199–215.
- [24] A. Fekete, N. Lynch, Y. Mansour, J. Spinelli, The impossibility of implementing reliable communication in the face of crashes, *Journal of the ACM* 40 (3) (1993) 1087–1107.
- [25] F. Fich, End-to-end communication protocols, in: 2nd Int. Conference on Principles of Distributed Systems (OPODIS), Hermes, 1998, pp. 37–43.
- [26] F. Fich, A. Jakoby, Short headers suffice for communication in a DAG with link failures, in: 14th International Symposium on Distributed Computing (DISC) LNCS, Springer, 2000, pp. 360–373.
- [27] O. Goldreich, A. Herzberg, Y. Mansour, Source to destination communication in the presence of faults, in: 8th ACM Symposium on Principles of Distributed Computing (PODC), 1989, pp. 85–101.
- [28] A. Herzberg, Connection-based communication in dynamic networks, in: 11th ACM Symposium on Principles of Distributed Computing (PODC), 1992, pp. 13–24.
- [29] E. Kushilevitz, R. Ostrovsky, A. Rosén, Log-space polynomial end-to-end communication, in: 28th ACM Symposium on Theory of Computing (STOC), 1995, pp. 559–568.
- [30] R. Ladner, A. LaMarca, E. Tempero, Counting protocol for reliable end-to-end transmission, *Journal of Computer and System Sciences* 56 (1998) 96–111.
- [31] Y. Mansour, B. Schieber, The intractability of bounded protocols for on-line sequence transmission over non-fifo channels, *Journal of the ACM* 39 (4) (1992) 783–799.
- [32] J. Postel, Internet protocol, Internet Engineering Task Force Request For Comments 791, 1981.
- [33] B. Reed, Finding approximate separators and computing treewidth quickly, in: 24th ACM Symposium on Theory of Computing (STOC), 1992, pp. 221–228.
- [34] N. Robertson, P.D. Seymour, Graph minors. III. Planar treewidth, *Journal of Combinatorial Theory Series B* 36 (1984) 49–64.
- [35] N. Robertson, P.D. Seymour, Graph minors. V. Excluding a planar graph, *Journal of Combinatorial Theory Series B* 41 (1986) 92–114.
- [36] N. Robertson, P.D. Seymour, R. Thomas, Quickly excluding a planar graph, *Journal of Combinatorial Theory Series B* 62 (1994) 323–348.
- [37] N. Stenning, A data transfer protocol, *Computer Networks* 1 (2) (1976) 99–110.
- [38] E. Tempero, R. Ladner, Recoverable sequence transmission protocols, *Journal of the ACM* 42 (5) (1995) 1059–1090.
- [39] U. Vishkin, A distributed orientation algorithm, *IEEE Transaction on Information Theory* 29 (1983) 624–629.
- [40] D.-W. Wang, L. Zuck, Tight bounds for the sequence transmission problem, in: 8th ACM Symposium on Principles of Distributed Computing (PODC), 1989, pp. 73–83.
- [41] D.-W. Wang, L. Zuck, Real-time sequence transmission problem, in: 10th ACM Symposium on Principles of Distributed Computing (PODC), 1991, pp. 111–124.



Pierre Fraigniaud received the Ph.D. degree in Computer Science from the ENS Lyon, France, in 1990. He is currently Directeur de Recherches at the French CNRS, and located at the Computer Science Department of Univ. Paris-Sud (LRI) where he is leading the research group “Graph Theory and Fundamental Aspects of Communications” (GraFComm). He is an associated member of the INRIA project Grand

Large on global parallel and distributed computing. His main research interests include several aspects of communication networks, parallel and distributed systems, and telecommunication systems. He is member of the Editorial Board of *Theory of Computing Systems (TOCS)* and *Journal of Interconnection Networks (JOIN)*, and recently acted as Program Chair for the 19th International Symposium on Distributed Computing (DISC '05), and for the 13th ACM Symposium on Parallel Algorithms and Architectures (SPAA '01). He is member of the Steering Committee of the series of International Symposium on Theoretical Aspects of Computer Science (STACS). He is Chair of the COST Action 295 DYNAMO on “Dynamic Communication Networks: Foundations and Algorithms”.



Cyril Gavoille received the Ph.D. from Ecole Normale Supérieure de Lyon in 1996. He is currently a full professor at Bordeaux University in the Computer Science Laboratory LaBRI. He is the author of 20 international journal papers and of 40 international conference papers. His main research interests are distributed computing and data-structures, specially routing algorithms. He was chairman of SIROCCO in 1999

(Colloquium on Structural Information and Communication Complexity), co-chair of AlgoTel (Algorithmic Aspects of Telecommunications) Conference in 2003, co-chair of LOCALITY workshop at DISC (Distributed Conference) in 2005, and has participated in 18 international program committee conferences (PODC, PDCN, HiPC, IWDC for 2005).