# Virtual path layouts in simple ATM networks

Jean-Claude Bermond, Nausica Marlin, David Peleg, Stéphane Pérennes

# Virtual Path Layouts in Simple ATM Networks*

Jean-Claude Bermond† Nausica Marlin† David Peleg ‡ Stéphane Perennes†

**Abstract**

In an ATM network, Virtual Paths (VPs) are built on top of the physical network and form a virtual network. Routing is done using these VPs. The sum of the bandwidths of the VPs that share a physical link constitutes the *load* of this link, which is limited by its *capacity*. The number of VPs used along a connection is defined as the *hop count* of that connection. Naturally, the hop count should be kept as low as possible so as to guarantee the efficiency of the virtual channels.

In its most general formulation, the *Virtual Path Layout* problem is an optimization problem in which, given certain communication demands between ordered pairs of nodes and constraints on the capacity and hop count, it is first required to design a system of VPs satisfying the constraints, and then minimizing some function of the load and hop count.

As the general problem is NP-hard, hence difficult to solve, we consider here a more restrictive version, in which all VPs have equal bandwidths, communication demands are uniform (namely, all pairs of nodes are equally required to communicate), and all physical links have equal capacity.

In practice, physical links based on optical fibers and cables are directed, and the amounts of information flowing in the two directions of a connection are not necessarily identical (consider, for example, a video-on-demand application). However, traditional research on this problem has focused on the undirected model. In this paper, we investigate the problem of designing directed VPs in a directed ATM model. Our aim is to minimize the maximum hop-count for a given capacity. We present good VP Layouts for some of the most popular types of simple networks, namely cycles, tori, and general trees. Our constructions are the best possible, up to a constant. In general, it turns out that the directed model admits better VP Layouts than the undirected one.

**Keywords :** ATM, Virtual Path Layout, Hop Count, Route Optimization, Network Design.

*This work has been supported by the French-Israeli cooperation "AFIRST"

†I3S Université de Nice, SLOOP joint project CNRS-UNSA-INRIA, 2004 route des Lucioles, BP93, F-06902 Sophia-Antipolis, France. {bermond, nmarlin, speren} @sophia.inria.fr

‡Department of Applied Mathematics and Computer Science, The Weizmann Institute, Rehovot 76100, Israel. peleg@wisdom.weizmann.ac.il

# 1 Introduction

Recent developments in the field of communication, such as the appearance of fiber optic based communication and the ATM network architecture [De 95, KG98], have thoroughly changed the way messages are routed in today's new communication networks. In ATM, data is transferred in fixed length packets, each routed independently. The routes are composed of two types of simple paths in the network, namely, *virtual paths* (or VPs), and *virtual channels* (or VCs).

The VPs can be thought of as predetermined path segments, built out of ordinary links, and used as a collection of fixed "expressways" for passing messages over large distances in the network. The VCs, on the other hand, are ad-hoc routes established in response to specific requests for connecting network users (e.g., telephone calls). The purpose of defining a fixed collection of VPs in the network is to simplify network management and make it easier to select the VC routes; the route of a VC is obtained as a concatenation of a number of complete VPs (see Fig. 1).

For example, let us consider setting up two Virtual Circuits, one connecting $A$ to $B$ and one connecting $B$ to $C$. The first of these connections can be obtained by concatenating the VPs $(Ada)$, $(abf)$ and $(fgB)$, yielding the route $(AdabfgB)$. Similarly, the second connection can be built along the VPs $(Bgc)$, $(cfbe)$ and $(edC)$, yielding the route $(BgcfbedC)$.

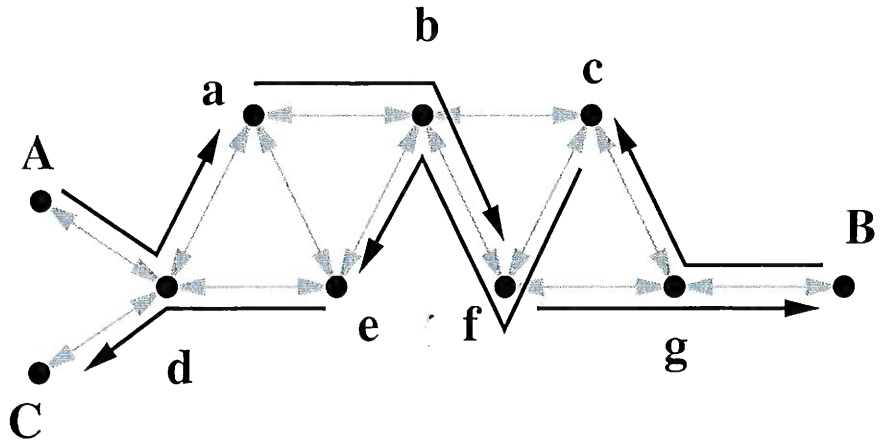


Figure 1: A network and a collection of VPs on it. Note that the physical links are symmetric (bi-directed), yet the VPs are directed.

A major problem in this framework is the one of defining the set of VPs in such a way that some good properties are achieved. The following considerations are involved. A capacity (or bandwidth) is assigned to each VP. The sum of the capacities of the VPs that share a physical link constitutes the *load* of this link. (In the above example, each directed link carries a load of 1 unit, although some links carry 1 load unit in both directions.) Naturally, this load must not exceed the link's capacity, namely, the amount of data it can carry. The sum of the capacities of all the physical links is a major component in the cost of the network, and should be kept as low as possible.

At the same time, a second important parameter is the maximum number of VPs in a virtual

channel, termed *hop count* in the literature. (In the above example, both VCs have hop count 3.) This parameter should also be kept as low as possible so as to guarantee low set up times for the virtual channels.

In its most general formulation, the *Virtual Path Layout (VPL)* problem is an optimization problem in which, given a certain communication demand between pairs of nodes and constraints on the maximum load and hop count, it is first required to design a system of VPs satisfying the constraints and then minimizing some given function of the load and hop count.

We employ a restricted model similar to the one presented by Cidon, Gerstel and Zaks in [GZ94]. In particular, we assume that all VPs have equal capacities, normalized to 1. Hence the load of a physical link is simply the number of VPs that share this link.

Although links based on optical fibers and cables are directed, traditional research uses an undirected model. Indeed, this model imposes the requirement that if there exists a VP from $u$ to $v$ then there exists also a VP from $v$ to $u$. In fact, that is the way ATM networks are implemented at the present time. However, the two VPs (the one from $u$ to $v$ and the one in the other direction) do not need to have the same capacity. Indeed, in many applications the flows on the VPs are not equal. For example, in a video application where $u$ is a server and $v$ a consumer there is a VP from $u$ to $v$ using a large capacity (transmission of video data) and a VP from $v$ to $u$ used only for control or acknowledgments with a very small capacity which can be considered as negligible. Therefore, it seems more reasonable to use a directed model like the one introduced by Chanas and Goldschmidt in [CG98]. This would allow us to model the situation described above by a single VP of capacity 1 in the main direction.

It is convenient to represent the structure of VPs constructed for the communication network $G$ by a virtual directed graph $H$ defined on the top of the physical network over the nodes of $G$, whose arcs represent the VPs defined in the network. Specifically, $H$ contains an arc $(a, b)$ precisely if there exists a VP leading from $a$ to $b$. This virtual digraph $H$ provides a *directed VP layout (DVPL)* for the physical graph. The virtual graph $H$, along with the accompanying VP layout, is called an *admissible* virtual digraph for the physical network $G$, if it doesn't load any arc more than its capacity.

In this framework, each VC in $G$ can be viewed as a simple dipath in the virtual digraph $H$. Therefore, for a given source-destination pair $A, B$, the hop count of any VC connecting $A$ to $B$ must be at least the distance from $A$ to $B$ in $H$.

In this article, we focus on the all-to-all problem (i.e., all pairs of nodes are equally likely to communicate). Thus, the resulting maximum hop count can be viewed as the *diameter* of the virtual graph $H$ induced by the VPs. Consequently, we consider the central problem of finding a tradeoff between the maximum load and the virtual diameter. In particular, we address the following problem: Given a capacity on each physical arc, minimize the diameter of an admissible virtual graph $H$.

**Related Work** The problem has been considered in the undirected case, for example, in [GZ94, GWZ95, SV96, GCZ96, KKP97, EFZ97]. The problem of minimizing the maximum load over all VPL with bounded hop-count is studied in [FZ97, BBGG97], and minimizing also the average load is considered in [GWZ95]. The one-to-many problem is handled in [FZ97, GWZ95],

where the focus is on minimizing the eccentricity of the virtual graph from a special node called the root (this problem is the *rooted VP layout problem*) rather than minimizing the diameter of the virtual graph. A duality in the chain network between the problem of minimizing the hop-count knowing the maximum load, and the one of minimizing the load, knowing the maximum hop-count, is established in [FZ97]. The reader can find an excellent survey of the results in the undirected model in [Zak97].

**Our Results** The following table summarizes the results, giving upper bounds on the virtual diameter (the minimum diameter of an admissible virtual digraph) as a function of the number of vertices $n$ in the physical graph, its diameter $D_G$, its maximum in- and out-degree $d$, and the capacity $c$ considered as a constant. For lack of space, we only present upper bounds on the virtual diameter for the cycle, the grid and the arbitrary tree. All the proofs can be found in the full version [BMPP98]. As shown therein, all the layouts proposed here are near-optimal. The results mentioned in the table for the path in the special case of $c = 1$ are due to [Cha98, BCG98]. We use their optimal layout in the path to build our layout in the arbitrary tree.

| Graph $G$ | Capacity | Upper Bound |
|---|---|---|
| Path $P_n$ | $c = 1$ | $\lceil \frac{n}{2} \rceil + \lceil \log n \rceil$ |
| | $c = \mathcal{O}(1)$ | $\mathcal{O}(c \cdot n^{\frac{1}{2c-1}})$ |
| Cycle $C_n$ | $c = 1$ | $2\sqrt{2n} + \mathcal{O}(1)$ |
| | $c = \mathcal{O}(1)$ | $\mathcal{O}(c \cdot n^{\frac{1}{2c}})$ |
| Torus $TM(a,b), a \leq b$ | $c = \mathcal{O}(1)$ | $\mathcal{O}(a \cdot b^{1/2ac})$ |
| Mesh $M(a,b)$, $\log b \leq a \leq b$ | $c = \mathcal{O}(1)$ | $\mathcal{O}(\log n)$ |
| Arbitrary Tree $T$ | $c = 1$ | $\tilde{D}(P_{D_T}, c)$<br>$\frac{n}{=}\theta(log^2(n))$ |
| | $c == 2$ | $32 \cdot n^{1/3}$<br>${D_G}^{1/3}.\log n$ |
| | $c = \mathcal{O}(1)$ | $8c \cdot n^{1/(2c-1)}$<br>${D_G}^{1/(2c-1)} \cdot \log n$ |
| Complete binary Tree $T$ | $c = 2$ | $D_G/2 + 2$ |

## 2 The Cycle

../FiguresIFIP 2 and 3 show VP layouts on the symmetric directed cycle with capacity 1 and 2 on each physical link. The label on a VP is its *dilation*, i.e., the number of physical links used in this VP.

Let us describe the construction for capacity 1 in more detail. We refer to the "long" (dilation $m$) VPs as *highways*, and call the starting nodes of the highways *main nodes*. To route from a vertex $x$ to a vertex $y$, one can either use only counterclockwise "short" (dilation 1) VPs, or go from the node $x$ counterclockwise to the nearest main node [in at most $m - 1$ steps], then clockwise using highways till the first main node after $y$ [in at most $n/m - 1$ steps], and then counterclockwise again to $y$ [in at most $m - 1$ additional steps]. The maximum hop count is

$2m + n/m - 3$. Choosing $m = (\frac{n}{2})^{1/2}$, we obtain an upper bound of $2\sqrt{2n} - 3$ on the maximum hop count.

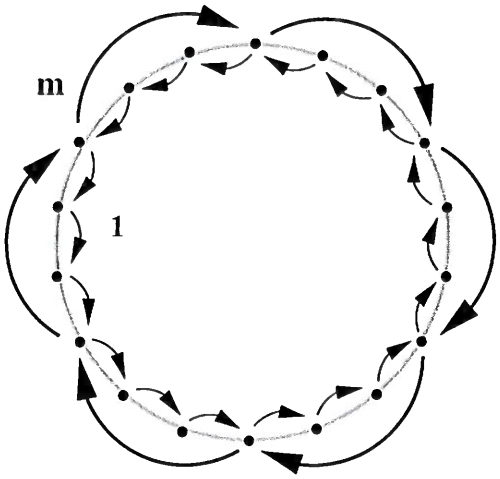


Figure 2: $C_n$, $c = 1$, $m = (\frac{n}{2})^{1/2}$

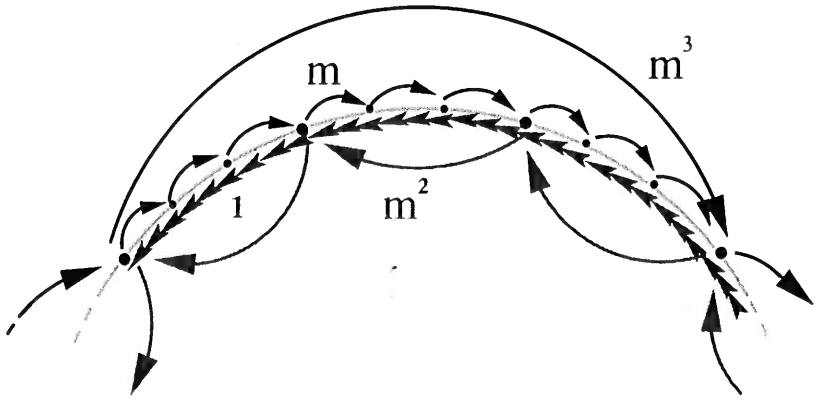


Figure 3: $C_n$, $c = 2$, $m = (\frac{n}{2})^{1/4}$

Similarly, for capacity 2, let us classify the VPs into classes by their length, referring to VPs of dilation $m^i$ (for $i = 1, 2, 3$) as *type-$i$ highways*, and call the starting nodes of a type-$i$ highway *main nodes of type $i$*. Note that a main node of type $i$ is also of type $i - 1$. Now, to route from a node $x$ to a node $y$, proceed as follows. Using counterclockwise VP of dilation 1, one reaches the first main node counterclockwise from $x$ of type 1 [in at most $m - 1$ steps]. Then using clockwise VPs of dilation $m$, reach the first main node of type 2 [in at most $m^2/m - 1 = m - 1$ steps]. Next, using counterclockwise VPs of dilation $m^2$, reach the first main node of type 3 [in $m - 1$ steps]. From this node of type 3, one reaches the first node of type 3 just after $y$ [in $n/m - 1 = 2m - 1$ steps]. Then one symmetrically reached $y$ going down [$3 \times (m - 1)$ steps]. The upper bound established on the maximum hop count is $8m - 7 \leq 4 \times 2 \times (\frac{n}{2})^{1/4}$.

This construction can be generalized to obtain an upper bound of $4c(\frac{n}{2})^{\frac{1}{2c}}$ for any constant capacity $c$ on the physical links.

## 3 The Toroidal Mesh

The physical digraph $G$ is the toroidal mesh of dimensions $a \times b$, $TM(a,b)$. Recall that $TM(a,b) = C_a \square C_b$, the Cartesian sum (also called product) of two cycles.

Figure 4 depicts the entire torus $TM(2,b)$ with a quasi-optimal layout for capacity 1; compare with the layout given earlier for the cycle with capacity 2. The internal cycle of length $b$ is used for setting up VPs of dilation 1 and VPs of dilation $m = (\frac{n}{2})^{1/4}$. The external cycle of length $b$ is used for VPs of dilation $m^2$ and $m^3$. To route from a node $x$ to a node $y$, we use the same path as on the cycle $C_b$ with capacity 2, except that we might have to use a transversal arc four times: at the beginning - if $x$ is on the external cycle; then from the type 2 node of the internal cycle to the external cycle in order to use VPs of dilation $m^2$ and $m^3$; and in the opposite direction to reach $y$.

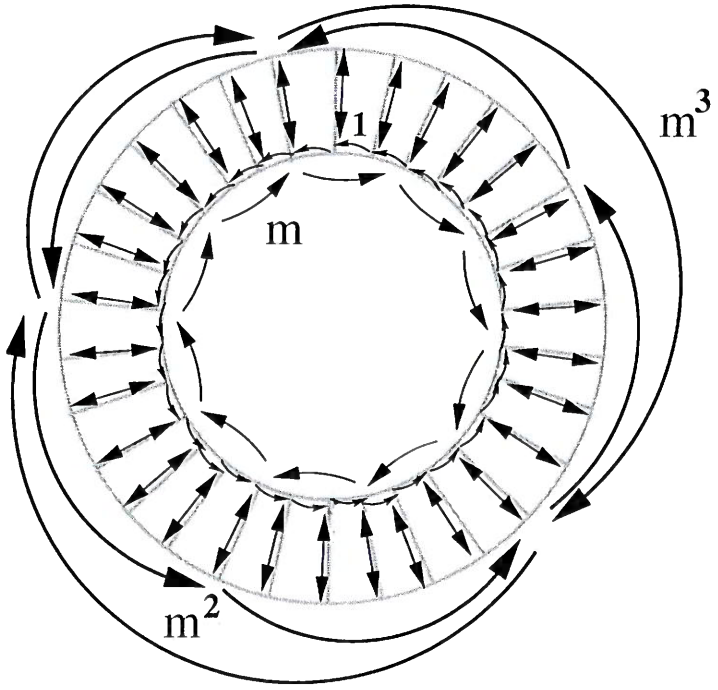


Figure 4: $TM(2,b)$, $c = 1$, $m = (\frac{b}{2})^{1/4}$

Again, this idea can be generalized to obtain a layout for capacity $c$ with hop-count

$$8c \cdot \left(\frac{a}{2}\right)^{1/2c-1} + 2a + 1 + (4ac-1)\left(\frac{b}{2}\right)^{1/2ac}.$$

On the cycles of length $b$, we apply the optimal layout for the cycle with capacity $a \times c$. On the cycles of length $a$, we apply the optimal layout for the cycle with capacity $c-1$, and we also put VPs of dilation 1 in both directions. In one direction, these VPs come from the optimal layout for the cycle with capacity $c-1$; in the other direction, we build VPs of dilation 1 using the remaining capacity unit.

# 4 The Arbitrary Tree

To construct good VP layouts for an arbitrary tree with capacity $c$, we use induction.

As shown on figure 5, a subtree hanging from a node is a tree rooted at this node such that the degree of the root in the subtree is one. A middle node $m$ of a tree $T$ is a node such that the subtrees hanging from $m$ contain at most $\lceil \frac{n}{2} \rceil$ nodes.

Any tree admits at least one middle node. For proof, take $x$ a node of $T$. Let $M(x)$ be the size of the larger subtree hanging from $x$. If $M(x) \geq \lceil \frac{n}{2} \rceil + 1$ then, all together, other subtrees hanging from $x$ contain at most $\lceil \frac{n}{2} \rceil - 1$ nodes. We replace $x$ by its neighbor $x'$ in the largest subtree. Since $M(x') \leq M(x) - 1$, the process stops at a middle node.

For a given rooted tree $(T, r)$, we define our VPL inductively.

Let $m$ be a middle node for $T$. We use the best known VPL on the path $P$ between $r$ and $m$ (maximum hop count $O(c.n^{\frac{1}{2c-1}})$). We consider $T_i$, the subtrees hanging from a node of the path $P$ and we apply the construction to theses trees $T_i$ rooted at the node they are hanging from. Note that, since $m$ is a middle node, if $T$ has $n$ nodes then $T_i$ has at most $n/2$ nodes.

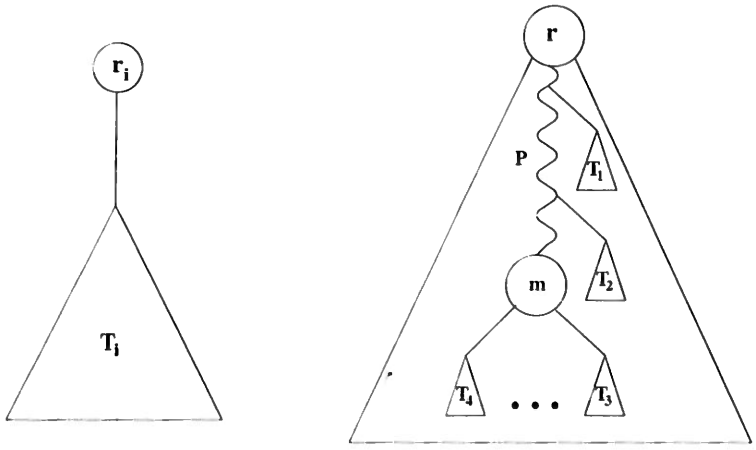


Figure 5: The tree $T_i$ hanging from $r_i$ ($r_i$ is the root of $T_i$). The VP Layout on $T$.

For this layout, let $h(T, r)$ be the maximum hop count from any node to the root or from the root to any node. Let $f(n)$ be the maximum of $h(T, r)$ taken over all the rooted trees with $n$ nodes.

Consider a vertex $x \in T$ located in one of the subtrees $T_i$. To go from $x$ to the root $r$, one first reaches some node of the path $P$, so doing we are moving inside the subtree $T_i$, then we have to move on the path $P$ to reach $r$. Hence $h(T, r)$ is at most $f(n/2)$ plus the maximum hop count in a path of length $n/2$. Let us recall that for any $c$ this last value is $\mathcal{O}(c.n^{\frac{1}{2c-1}})$ and for $c = 1$ it is improved to $\lceil \frac{n}{4} \rceil + \lceil \log(n/2) \rceil$.

For $c = 1$, we get

$$f(n) \leq f(\lceil n/2 \rceil) + \lceil n/4 \rceil + \lceil \log(n/2) \rceil .$$

Thus by summing $f(n) \leq n/2 + \log^2(n)$

The last trick is to remark that for a given tree $T$, we can choose the first root as we want. We choose $r = m$ and, if $T$ has $n_0$ nodes, we get: $h(T, m) \leq f(n_0/2) \leq n_0/4 + \log^2(n_0/2)$. The

maximum hop count is at most $2.h(T,m)$ that is at most $n/2 + 2.log^2(n)$.

For capacity $c$, the same computation leads to $\mathcal{O}(c.n^{1/(2c-1)})$.

## 5 Conclusions

The main idea behind our construction is to prudently partition the capacity available to us on the links, and use it for setting up a hierarchy of "expressways" of various "speeds." This allows us to significantly reduce the resulting hop counts, by enabling messages routed to large distances to perform large "jumps" efficiently covering significant portions of their route, and use "slower roads" only at the initial and final segments of their trip. We expect this idea to be applicable and useful also for other kinds of topologies, and in principle also in networks of arbitrary topology, and we intend to explore this idea further in the future.

## References


[BBGG97] L. Becchetti, P. Bertolazzi, C. Gaibisso, and G. Gambosi. On the design of efficient ATM routing schemes. manuscript, 1997.

[BCG98] M. Burlet, P. Chanas, and O. Goldschmidt. Optimization of VP layout in ATM networks. In preparation, 1998.

[BMPP98] J.-C. Bermond, N. Marlin, D. Peleg, and S. Perennes. Directed virtual path layouts in atm networks. in preparation, 1998.

[CG98] P. Chanas and O. Goldschmidt. Conception de réseau de VP de diamètre minimum pour les réseaux ATM. In *Road-f'98*, pages 38–40, 1998.

[Cha98] Pascal Chanas. *Dimensionnement de réseaux ATM.* PhD thesis, CNET Sophia, Sept. 1998. In preparation.

[De 95] Martin De Pricker. *Asynchronous Transfer Mode, Solution for Broadband ISDN.* Prentice Hall, August 1995. 3rd edition 332p.

[EFZ97] T. Eilam, M. Flammini, and S. Zaks. A complete characterization of the path layout construction problem for ATM networks with given hop count and load. In *24th International Colloquium on Automata, Languages and Programming (ICALP)*, volume 1256 of *Lecture Notes in Computer Science*, pages 527–537. Springer-Verlag, 1997.

[FZ97] M. Feighlstein and S. Zaks. Duality in chain ATM virtual path layouts. In *4th International Colloquium on Structural Information and Communication Complexity (SIROCCO)*, Monte Verita, Ascona, Switzerland, July 1997.

[GCZ96] O. Gerstel, I. Cidon, and S. Zaks. The layout of virtual paths in ATM networks. *IEEE/ACM Transactions on Networking*, 4(6):873–884, 1996.

[GWZ95] O. Gerstel, A. Wool, and S. Zaks. Optimal layouts on a chain ATM network. In *3rd Annual European Symposium on Algorithms*, volume LNCS 979, pages 508–522. Springer Verlag, 1995.

[GZ94] Ornan Gerstel and Shmuel Zaks. The virtual path layout problem in fast networks. In *Symposium on Principles of Distributed Computing (PODC '94)*, pages 235–243, New York, USA, August 1994. ACM Press.

[KG98] Daniel Kofman and Maurice Gagnaire. *Réseaux Haut Débit, réseaux ATM, réseaux locaux et réseaux tout-optiques.* InterEditions-Masson, 1998. 2eme édition.

[KKP97] E. Kranakis, D. Krizanc, and A. Pelc. Hop-congestion trade-offs for high-speed networks. *International Journal of Foundations of Computer Science*, 8:117–126, 1997.

[SV96] L. Stacho and I. Vrt'o. Virtual path layouts for some bounded degree networks. In *Structure, Information and Communication Complexity, 3rd Colloquium, SIROCCO*, pages 269–278. Carleton University Press, 1996.

[Zak97] S. Zaks. Path layout in ATM networks - a survey. In *The DIMACS Workshop on Networks in Distributed Computing, DIMACS Center, Rutgers University*, Oct. 1997. manuscript.