Minimizing Congestion of Layouts for ATM Networks with Faulty Links

Leszek Gąsieniec * Evangelos Kranakis †§ Danny Krizanc †§
Andrzej Pelc ‡§

Abstract

We consider the problem of constructing virtual path layouts for an ATM network consisting of a complete network K_n of n processors in which a certain number of links may fail. Our main goal is to construct layouts which tolerate any configuration of up to f layouts and have the least possible congestion. First, we study the minimal congestion of 1-hop f-tolerant layouts in K_n. For any positive integer f we give upper and lower bounds on this minimal congestion and construct f-tolerant layouts with congestion corresponding to the upper bounds. Our results are based on a precise analysis of the diameter of the network $K_n[\mathcal{F}]$ which results from K_n by deleting links from a set $\mathcal{F}$ of bounded size. Next we study the minimal congestion of h-hop f-tolerant layouts in K_n, for larger values of the number h of hops. We give upper and lower bounds on the order of magnitude of this congestion, based on results for 1-hop layouts. Finally, we consider a random, rather than worst case, fault distribution. Links fail independently with constant probability $p < 1$. Our goal now is to construct layouts with low congestion that tolerate the existing faults with high probability. For any $p < 1$, we show such layouts in K_n, with congestion $O(\log n)$.

Key Words and Phrases: ATM networks, Complete network, Congestion, Faulty links, Hop, Virtual path.
1980 Mathematics Subject Classification: 68Q99
CR Categories: C.2.1
Technical Report RR 95/05-8, Université du Québec à Hull

*Instytut Informatyki, Uniwersytet Warszawski, Banacha 2, 02-097 Warszawa, Poland. E-mail: lechu@mimuw.edu.pl (This research was done during the author's stay at the Université du Québec à Hull as a postdoctoral fellow.)

†Carleton University, School of Computer Science, Ottawa, ON, K1A 5B6, Canada. E-mail: {kranakis,krizanc}@scs.carleton.ca

‡Département d'Informatique, Université du Québec à Hull, Hull, Québec J8X 3X7, Canada. E-mail: pelc@uqah.uquebec.ca

§Research supported in part by NSERC (Natural Sciences and Engineering Research Council of Canada) grants.

1 Introduction

Broadband Integrated Digital Services (ISDNs, for short) are meant to accomodate various kinds of data traffic, including voice, video, image, file transfer, as well as interactive. Such systems provide true high rate file transfer, video conferencing, on demand HDTV, interface with high-speed LANs, etc. Requirements in new and emerging information services require a new transfer mode for broadband ISDN. ATM (Asynchronous Transfer Mode) was developed [3] as a packet structure for broadband ISDNs. ATM is a new multiplexing and switching technology that results in more cost effective solutions of greater flexibility than several separate individually optimized technologies [3]. Because of this, it has significant commercial as well as public service applications. This technology is thoroughly described in the literature [8, 7].

For standard networks, routing is based either on variable or large size data units, and has been the topic of extensive studies in the literature, e.g., see [9, 1, 2]. However, this is not suitable for present day multimedia environments, which have large throughputs of different types of data services. By contrast, in ATM networks routing is based on relatively small fixed-sized packets. Such packets are routed through a layout of "virtual" paths, as well as sequences of such virtual paths, also called "virtual" channels. Although such layouts may be time-expensive to set-up from scratch by the network user, they remain fixed for relatively long time.

However, as the size of the networks grows ATM studies must take into account the possibility of component failures. In a network where links may fail it is important to establish a virtual path layout which guarantees fault-free transmission of the packets. The construction has to take into account the available capacity of the existing links, i.e., the congestion bounds of the links cannot be exceeded. Hence it is important to construct fault-tolerant virtual path layouts with the least possible congestion.

The model we use in this paper is based on the Virtual Path Layout model introduced by Gerstel and Zaks [5, 6]. Messages may be transmitted through arbitrarily long virtual paths. Packets are routed along those paths by maintaining a routing field whose subfields determine intermediary destinations of the packet, i.e., end-points of virtual paths on its way to the final destination. In such a network it is important to construct path layouts that take into account tradeoffs between the hop number (i.e. the number of virtual paths used to travel between any two nodes) and edge-congestion (i.e. the number of virtual paths passing through a link).

1.1 Notation and definitions

In this paper we consider the problem of constructing virtual path layouts for a complete network of n processors in which a certain number, say f, of links may fail. We also assume that the number of faults may be arbitrary but the network resulting after removing the faulty links remains connected. Before proceeding with an outline of the main results of the paper we give the following definitions.

- K_n is the complete network on a set X of n nodes.
- For any set $\mathcal{F}$ of links of K_n, $K_n[\mathcal{F}]$ is the network resulting from K_n by deleting all links from $\mathcal{F}$.
- A virtual path (VP) in a network is a simple chain in this network (i.e., a non-repetitive sequence $(v_1, ..., v_k)$ of nodes).
- A virtual channel (VC) of length k, joining nodes u and v, is a sequence $p_1, p_2, \ldots, p_k$ of VPs such that p_1 begins at node u, p_k ends at node v and the beginning of p_{i+1} coincides with the end of p_i, for $i < k$.
- An h-hop (virtual path) layout $\mathcal{P}$ in K_n is a collection of virtual paths, such that every pair of nodes u and v is joined by a VC of length at most h, composed of VPs from $\mathcal{P}$.
- The congestion of a layout $\mathcal{P}$ in K_n is the maximum number of VPs from $\mathcal{P}$ passing through any link of K_n.
- A layout $\mathcal{P}$ in K_n is h-hop f-tolerant, for positive integers h and f, if, for any set $\mathcal{F}$ of links such that $|\mathcal{F}| \leq f$ and $K_n[\mathcal{F}]$ is connected, any pair of nodes u and v is joined by a VC of length at most h, composed of VPs from $\mathcal{P}$ not containing links from $\mathcal{F}$.
- The length of a simple path in a network is the number of links in this path.
- The diameter of a connected network is the maximum over all pairs of distinct nodes of minimal lengths of paths joining these nodes.
- By $[n]_k$, for positive integers n and k, we denote the descending factorial $n(n-1)\cdots(n-k+1)$. For $k \leq 0$, we define $[n]_k = 1$.

1.2 Results of the paper

First, we study the minimal congestion of 1-hop f-tolerant layouts in K_n. For any positive integer f we give upper and lower bounds on this minimal congestion and construct f-tolerant layouts with congestion corresponding to the upper bounds. Our results are based on a precise analysis of the diameter of the network $K_n[\mathcal{F}]$ when the number $|\mathcal{F}|$ of faulty links is bounded by $F(n,k) = k(n - \frac{k+3}{2})$. The bounds $F(n,k)$ on the number of faults play an important role in our considerations, as they yield thresholds for the diameter of $K_n[\mathcal{F}]$. We show that

- if the number of faults is $f = F(n,k)$ then the minimal congestion of a 1-hop f-tolerant layout is $\Theta(k[n-2]_k)$.

Next we study the minimal congestion of h-hop f-tolerant layouts in K_n, for larger values of the number h of hops. We give upper and lower bounds on the order of magnitude of this congestion. More precisely, assuming that $F(n,k) < f \leq F(n,k+1)$ and $0 < h < n$, we show that the congestion depends on the ratio $\frac{k}{h}$, in the following way:

- there exists a h-hop f-tolerant layout in K_n with congestion

$$O\left(\frac{k}{h}[n-2]_{\lceil\frac{k+2}{h}\rceil-1}\right),$$

- every h-hop f-tolerant layout in K_n has congestion

$$\Omega\left(\frac{k}{h^2}[n-2]_{\lfloor\frac{k}{h}\rfloor-1}\right).$$

In the last section, we consider an alternative assumption on the fault distribution. Instead of imposing an upper bound f on the number of faulty links and assume their worst case location, we adopt a random approach. Assume that links fail independently with constant probability $p < 1$. Our goal is to construct layouts with low congestion that tolerate the existing faults with high probability. For any $p < 1$, we show such layouts in K_n, with congestion $O(\log n)$.

2 Single-hop Layouts

In this section we study the minimal congestion of 1-hop f-tolerant layouts in K_n. For any positive integer f we give upper and lower bounds on this minimal congestion and construct f-tolerant layouts with congestion corresponding to the upper bounds.

We first define a sequence of integers which will play an important role in our considerations. These integers are called *key-points*. For natural numbers n and $k \leq n-2$ define $F(n,k) = k(n-\frac{k+3}{2})$. Notice that $F(n,k+1) = F(n,k)+n-k-2$. Any natural number $f \leq \frac{(n-1)(n-2)}{2}$ can be uniquely represented as $F(n,k)+r$, where $k \leq n-2$ and $0 < r \leq n-k-2$.

The following result gives a lower bound on the congestion of any 1-hop f-tolerant layout in K_n.

Theorem 2.1 *Let $f = F(n,k)+r$, where $k \leq n-2$ and $0 < r \leq n-k-2$. Every 1-hop f-tolerant layout has congestion $\Omega(rk[n-2]_k)$.*

Proof: Let $f = F(n,k)+r$, where $k \leq n-2$ and $0 < r \leq n-k-2$. Let $(v_0, v_1, \ldots, v_k, v_{k+1})$ be any simple path of length $k+1$ in K_n. Consider the network G on the set X of all nodes, whose links are those from the above path and those of the complete graph on $X \setminus \{v_0, v_1, \ldots, v_{k-1}\}$ (see figure 1). Let $\mathcal{F}$ be the set of all links in K_n which are not links of G. Thus $G = K_n[\mathcal{F}]$ and

$$|\mathcal{F}| = \frac{n(n-1)}{2} - \frac{(n-k)(n-k-1)}{2} - k = F(n,k).$$

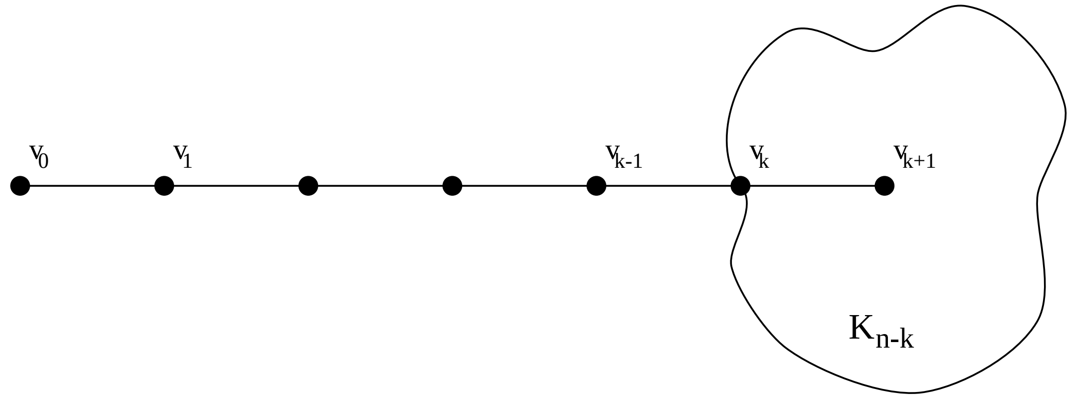

Figure 1: The network G

Consider any 1-hop f-tolerant layout $\mathcal{P}$ in K_n. We claim that there are at least $r+1$ VPs in $\mathcal{P}$ joining v_0 and v_{k+1}, which have the prefix $(v_0, v_1, ..., v_k)$. Suppose not and let $p_1, ..., p_t$, for $t \leq r$, be all such VPs in $\mathcal{P}$ with this prefix. Let Y be the set of links in all those VPs following link (v_{k-1}, v_k). Clearly $|Y| \leq r$ and hence $|\mathcal{F} \cup Y| \leq f$. However, all VPs in $\mathcal{P}$ joining v_0 and v_{k+1} must contain links from $\mathcal{F} \cup Y$, which contradicts 1-hop f-tolerance of $\mathcal{P}$.

It follows that for any pair u and v of distinct nodes and for any simple path $(u, v_1, ..., v_k, v)$ of length $k+1$ joining them, there are at least $r+1$ VPs in $\mathcal{P}$ joining u and v, which have the prefix $(u, v_1, ..., v_k)$. There are $\frac{n(n-1)}{2}$ pairs of nodes and $[n-2]_k$ simple paths of length $k+1$ joining each pair. For each such path, the $r+1$ VPs (of length at least $k+1$) contribute at least $(r+1)(k+1)$ links. Thus the sum of numbers of links in all VPs of the layout $\mathcal{P}$ is at least $(r+1)(k+1)\frac{n(n-1)}{2}[n-2]_k$. Consequently the congestion of $\mathcal{P}$ is at least the average of this total per link, i.e., $(r+1)(k+1)[n-2]_k = \Omega(rk[n-2]_k)$. □

Since $F(n, k+1) = F(n, k) + n - k - 2$, we obtain the following corollary concerning congestion of 1-hop f-tolerant layouts for f being a key-point.

Corollary 2.1 *If $f = F(n, k)$, where $k \leq n-2$ then every 1-hop f-tolerant layout has congestion $\Omega(k[n-2]_k)$.*

Next we focus attention on the upper bound on congestion of 1-hop f-tolerant layouts. We first prove the following lemma.

Lemma 2.1 *If $|\mathcal{F}| \leq F(n, k)$ then the diameter of $K_n[\mathcal{F}]$ is at most $k+1$.*

Proof: Let $\mathcal{F}$ be any set of links in K_n such that $K_n[\mathcal{F}]$ is connected. Suppose that the diameter of $K_n[\mathcal{F}]$ exceeds $k+1$. Consider any pair of nodes u and v and the shortest path P in $K_n[\mathcal{F}]$ joining u and v. First suppose that P has length $k+2$. Let Y be the set of all nodes not belonging to this path. Thus $|Y| = n-k-3$.

For any node w in Y, the distance between nodes v_1 and v_2 in the path, such that links (w, v_1) and (w, v_2) do not belong to $\mathcal{F}$, is at most 2: otherwise using the detour v_1, w, v_2 instead of the segment between v_1 and v_2 would create a shortcut, thus contradicting the minimality of P. It follows that for any node $w \in Y$ and all but three consecutive nodes v_i from P, the link (w, v_i) must belong to $\mathcal{F}$. On the other hand, all links joining nodes from the path P, except links in this path, must be in $\mathcal{F}$: otherwise a shortcut would again be possible. Thus the size of $\mathcal{F}$ is at least

$$(n-k-3)k + \frac{(k+3)(k+2)}{2} - (k+2) = F(n,k) + 1.$$

Consequently, if the shortest path in $K_n[\mathcal{F}]$ joining u and v has length $k+2$, then $|\mathcal{F}| \geq F(n,k)+1$. The same argument shows that if this shortest path has length $k+2+x$, for a positive integer x, then $|\mathcal{F}| \geq F(n, k+x)+1$. Since $F(n,k) < F(n,k+1)$, this implies that if the diameter of $K_n[\mathcal{F}]$ exceeds $k+1$ then $|\mathcal{F}| \geq F(n,k)+1$. This concludes the proof. □

The proofs of theorem 2.1 and lemma 2.1 provide the following characterization of the key-point $F(n,k)$: this is the maximum number f such that, whenever f links are deleted from the complete network K_n, the resulting network has diameter at most $k+1$.

Lemma 2.1 implies an upper bound on congestion of 1-hop f-tolerant layouts which matches the lower bound from theorem 2.1 when f is a key-point (cf. corollary 2.1).

Theorem 2.2 *Let $f \leq F(n,k)$, where $k \leq n-2$. Then the layout consisting of all VPs of length at most $k+1$ is 1-hop f-tolerant and has congestion $O(k[n-2]_k)$.*

Proof: Lemma 2.1 implies that the layout is 1-hop f-tolerant. In order to estimate its congestion, notice that, due to symmetry, the number of VPs containing link l is the same, for any link l. Each VP of length i contributes i to the congestion count. There are $[n]_{i+1}$ such VPs, hence the total contribution is $\sum_{i=1}^{k+1} i[n]_{i+1}$. Since all links are equally loaded, in order to obtain congestion, this sum should be divided by $\frac{n(n-1)}{2}$, thus giving congestion equal to $2\sum_{i=1}^{k+1} i[n-2]_{i-1}$. Let $a_i = i[n-2]_{i+1}$. We have $\frac{a_{i+1}}{a_i} \geq 2$ for $i \leq n-3$, hence the series in question grows faster than geometric and consequently the order of magnitude of the sum of its initial segment is the same as that of the last term. It follows that the congestion of our layout is $O(k[n-2]_k)$. □

The upper bound given in the above theorem is not tight for values of f which are not key-points. In fact we conjecture that it is the lower bound from theorem 2.1 which is tight for such values, up to a multiplicative constant.

Conjecture. Let $f = F(n, k) + r$, where $k \leq n-2$ and $0 < r \leq n-k-2$. There exists a 1-hop f-tolerant layout with congestion $O(r(k+1)[n-2]_k)$.

Notice that we formulated the conjecture putting $O(r(k+1)[n-2]_k)$ instead of $O(rk[n-2]_k)$ to include the case $k = 0$. Although we cannot prove this conjecture in general, we construct an appropriate layout for $k = 0$, i.e. when $f \leq n-2$. Let $v_0, ..., v_{n-1}$ be a labeling of all nodes of the complete network K_n. In the sequel all operations on node indices are performed modulo n. We say that two nodes v_i and v_j are k-neighbors if $k = \min(j-i, i-j)$. Note that $k \leq \lfloor \frac{n}{2} \rfloor$.

Consider the following 1-hop f-tolerant layout $\mathcal{P}_f$ in K_n, for $f \leq n-2$. If f is equal to $n-3$ or $n-2$ then $\mathcal{P}_f$ consists of all VPs of length 1 and 2. For $f \leq n-4$ the layout $\mathcal{P}_f$ is the union of $f+3$ disjoint layouts ${\mathcal{P}_f}^0, ..., {\mathcal{P}_f}^{f+2}$, defined as follows. The layout ${\mathcal{P}_f}^0$ consists of all VPs of length 1 in K_n. All other layouts ${\mathcal{P}_f}^i$, for $i = 1, ..., f+2$, consist of the following VPs of length 2. Any pair of k-neighbors in X, for $k = 1, ..., \lfloor \frac{n}{2} \rfloor$, except pairs (v_{i-1}, v_{k+i-1}) and (v_{k+i-1}, v_{2k+i-1}), is joined by a VP of length 2 with middle node v_{k+i-1}. This node is called *central* for fixed i and k. (see figure 2). Notice that, for fixed i and k, every link is in at most two VPs of ${\mathcal{P}_f}^i$ (cf. figure 2).

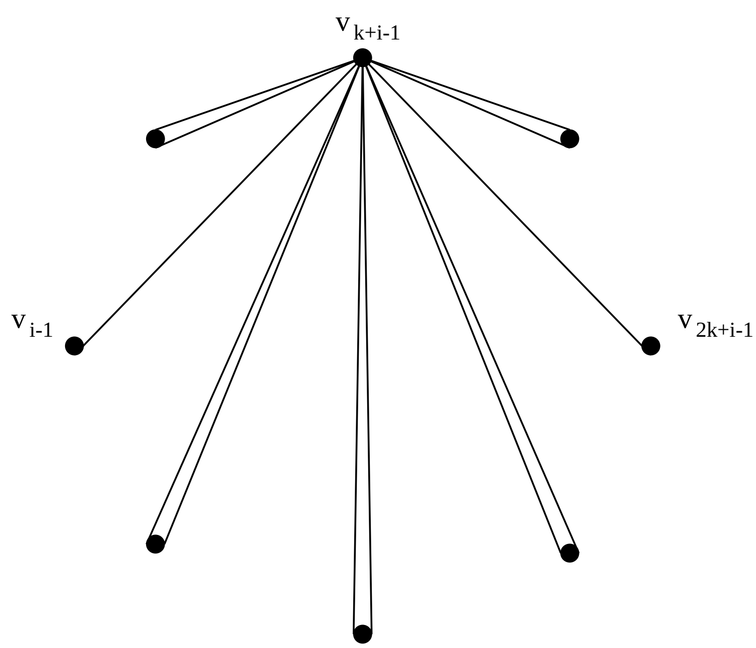

Figure 2: VPs joining k-neighbors in ${\mathcal{P}_f}^i$

Lemma 2.2 *For any $f \leq n-2$, the layout $\mathcal{P}_f$ is 1-hop f-tolerant and has congestion $O(f)$.*

Proof: For any pair of nodes u and v and all $i = 1, ..., f+2$, VPs in layouts $\mathcal{P}_f{}^i$ joining u and v have distinct middle nodes and hence are link disjoint. For at most two values of i the layout $\mathcal{P}_f{}^i$ does not contain a VP joining these nodes. Hence there are at least f link disjoint VPs of length 2 in $\mathcal{P}_f$ joining u and v. Together with the VP of length 1 this gives at least $f+1$ link disjoint VPs joining these nodes. This proves that the layout $\mathcal{P}_f$ is 1-hop f-tolerant.

It remains to estimate the congestion of layout $\mathcal{P}_f$. Fix a link $l = (u, v)$ and $i \in \{1, .., f+2\}$. Each of nodes u and v can be central for at most one value of k. For a fixed k, the link l is in at most two VPs of $\mathcal{P}_f{}^i$. Hence there are at most four VPs in $\mathcal{P}_f{}^i$ containing link l. This link is also in one VP from $\mathcal{P}_f{}^0$. This gives the upper bound $4(f+2)+1 = 4f+9$ on congestion of the layout $\mathcal{P}_f$.

□

Notice that the gap between the upper and the lower bounds obtained in theorems 2.1 and 2.2 for general values of $f = F(n, k) + r$, is a factor of $\frac{n-k}{r}$. For k close to n, as well as for f close to the nearest larger key-point, the orders of magnitude of the upper and lower bounds meet. In some cases, however, e.g. for $r = 1$ and $k = \frac{n}{2}$, this gap becomes a factor of $\Theta(n)$. The exact order of magnitude of the minimal congestion of 1-hop f-tolerant layouts for such values f remains open (cf. the conjecture).

3 Multi-hop Layouts

In this section we study the minimal congestion of h-hop f-tolerant layouts in K_n, for larger values of the number h of hops. We give upper and lower bounds on the order of magnitude of this congestion.

Theorem 3.1 *Assume that* $F(n, k) < f \leq F(n, k+1)$ *and* $0 < h < n$.
1. There exists a h-hop f-tolerant layout in K_n with congestion

$$O\left(\frac{k}{h}[n-2]_{\lceil \frac{k+2}{h} \rceil - 1}\right).$$

2. Every h-hop f-tolerant layout in K_n has congestion

$$\Omega\left(\frac{k}{h^2}[n-2]_{\lfloor \frac{k}{h} \rfloor - 1}\right).$$

Proof:

1. Let f be as assumed and let $|\mathcal{F}| \leq f$. According to lemma 2.1, there exists a path P of length $k+2$ in $K_n[\mathcal{F}]$, between any two nodes u and v. There exists a virtual channel

$C = (P_1, ..., P_h)$ of length h, joining u and v, such that all VPs P_i are segments of P and every path P_i has length at most $\lceil \frac{k+2}{h} \rceil$. It follows that the layout $\mathcal{Q}$ consisting of all virtual paths of length not greater than $\lceil \frac{k+2}{h} \rceil$ is h-hop f-tolerant. In view of theorem 2.2, the congestion of layout $\mathcal{Q}$ is $O(\frac{k}{h}[n-2]_{\lceil \frac{k+2}{h} \rceil - 1})$.

2. Fix a h-hop f-tolerant layout $\mathcal{P}$. Let $P = (v_0, v_1, ..., v_k)$ be any simple path of length k in K_n. Consider the unique virtual channel of length h joining v_0 and v_k whose VPs are consecutive segments $S_1, ..., S_h$ of P such that $S_1, ..., S_i$ have length $\lceil \frac{k}{h} \rceil$ and $S_{i+1}, ..., S_h$ have length $\lfloor \frac{k}{h} \rfloor$, for some $i \leq s$.

Consider the set $\mathcal{F}$ of links of size $F(n,k)$, defined in the proof of theorem 2.1. In the network $G = K_n[\mathcal{F}]$ the path P is the only simple path joining v_0 and v_k. Since $f \geq F(n,k)$, there must exist a VP in the layout $\mathcal{P}$ which contains one of the segments S_i. The value i must be the same, $i = i_0$, for at least $\frac{[n]_{k+1}}{h}$ paths P. For a given simple path S of length L, where $L = \lceil \frac{k}{h} \rceil$ or $L = \lfloor \frac{k}{h} \rfloor$, there exist $[n-L-1]_{k+1-L-1}$ simple paths of length k in which S is the i_0th segment. It follows that there are at least

$$N = \frac{[n]_{k+1}}{h[n - \lfloor \frac{k}{h} \rfloor - 1]_{k+1-\lfloor \frac{k}{h} \rfloor - 1}}$$

VPs of length $\Omega(\frac{k}{h})$ in the layout $\mathcal{P}$. Consequently, the congestion of $\mathcal{P}$ is at least

$$\Omega\left(\frac{k}{h} N \frac{2}{n(n-1)}\right) = \Omega\left(\frac{k}{h^2}[n-2]_{\lfloor \frac{k}{h} \rfloor - 1}\right).$$

□

If $h \geq k+2$, both the upper and the lower bound in the above theorem become constant. Indeed, in this case the layout consisting of VPs of length 1 (i.e., individual links only) is h-hop f-tolerant and has congestion 1.

4 Random faults

In this section we consider an alternative assumption on fault distribution. Instead of imposing an upper bound f on the number of faulty links and assume their worst case location, we adopt a random approach. Assume that links fail independently with constant probability $p < 1$. Our goal now is to construct layouts with low congestion that tolerate the existing faults with high probability. More precisely, a layout $\mathcal{P}$ is called h-hop *p-safe*, for a given positive real $p < 1$, if whenever links in the set $\mathcal{F}$ are chosen randomly with probability p, the probability that there exists a VC of length h in $K_n[\mathcal{F}]$, composed of VPs from $\mathcal{P}$, between any pair of nodes, is at least $1 - \frac{1}{n}$.

Theorem 4.1 *For any $p < 1$, there exists a 1-hop p-safe layout in K_n, with congestion $O(\log n)$.*

Proof: Let $q = 1 - (1-p)^2$, $c = \frac{-3}{\log q}$ and $f = \lceil c \log n \rceil$. In view of lemma 2.2, the layout $\mathcal{P}_f$ defined in section 2 has congestion $O(f) = O(\log n)$. It remains to show that this layout is p-safe.

Let $\mathcal{F}$ be a set of links chosen according to the definition of a p-safe layout. Consider any pair of distinct nodes u and v. There are f (link disjoint) VPs of length 2 joining u and v, in the layout $\mathcal{P}_f$. Consider any of those VPs, call it P. The probability that at least one of the two links of P is in $\mathcal{F}$ is equal to q. For distinct link disjoint paths P the above events are independent. Hence the probability that for every VP in $\mathcal{P}_f$, joining u and v, at least one of its links is in $\mathcal{F}$, is at most $q^f \leq q^{c \log n}$. Since there are less than n^2 pairs (u, v), the probability that this happens for at least one such pair is at most $n^2 q^{c \log n}$. We have

$$n^2 q^{c \log n} = n^2 n^{c \log q} = \frac{1}{n},$$

in view of $c = \frac{-3}{\log q}$. This proves that the layout $\mathcal{P}_f$ is p-safe. □

Since layout $\mathcal{P}_f$ considered in the above proof consists of VPs of length at most two, the following corollary holds.

Corollary 4.1 *For any $p < 1$, the layout consisting of all VPs of length one is 2-hop p-safe.*

In fact, in the 2-hop case we have $n - 2$ VC's for each pair of nodes. The probability that all of them are faulty for some pair is not only at most $\frac{1}{n}$ but indeed decreases exponentially in n.

The comparison of theorems 2.1 and 4.1 shows a dramatic difference between the worst case and random scenarios. Suppose that a fixed fraction, say $\frac{1}{4}$, of all links are faulty, and let $h = 1$. If we require worst case fault-tolerance from a layout, theorem 2.1, applied for a k linear in n, implies that the layout must have congestion exponential in n. However, if we assume that this fraction of faulty links are distributed randomly in the complete network K_n, theorem 4.1 shows that we can construct a layout fault-tolerant with high probability, which has only logarithmic congestion.

5 Conclusion

We presented upper and lower bounds on the minimal congestion of fault-tolerant layouts in complete networks, assuming the worst case fault distribution among links. We also showed that congestion can dramatically decrease if faults are distributed randomly and almost certain, rather than worst case, fault tolerance is required.

In the worst case scenario our bounds on congestion of 1-hop layouts are tight, up to a multiplicative constant, for some values of the number of faults, called key-points. In between these values there remain gaps between upper and lower bounds: we conjecture that the lower bounds are tight. Another interesting problem is to generalize the study of fault-tolerant layouts from the complete network to other graphs.

References

[1] B. Awerbuch, A. Bar-Noy, N. Linial and D. Peleg, "Improved Routing with Succinct Tables", Journal of Algorithms 11, 307-341 (1990).

[2] D. Bertsekas and R. Gallager, "Data Networks", Prentice-Hall Inc., 1992.

[3] J. Y. Le Boudec, "The Asynchronous Transfer Mode: A Tutorial", Computer Networks and ISDN Systems, 24:279-309, 1992.

[4] I. Cidon, O. Gerstel and S. Zaks, "A Scalable Approach to Routing in ATM Networks", WDAG 94, Springer Verlag LNCS.

[5] O. Gerstel and S. Zaks, "The Virtual Path Layout Problem in Fast Networks", ACM-PODC 94.

[6] O. Gerstel and S. Zaks, "The Virtual Path Layout Problem in ATM Ring and Mesh Networks", in Proceedings of SIROCCO 94, Carleton University Press, 1995, Ottawa.

[7] D. E. McDysan and D. L. Spohn, "ATM: Theory and Applications", McGraw-Hill Series on Computer Communication, 1995.

[8] M. de Prycker, "Asynchronous Transfer Mode: Solutions for Broadband ISDN", Ellis Horwood Limited, 1993 (2nd edition).

[9] N. Santoro and R. Khatib, "Labeling and Implicit Routing in Networks", Comput. J. 28(1985), 5-8.