# Adaptive Source Routing in High-Speed Networks

Alon Itai* and Hadas Shachnai†
Department of Computer Science
Technion, Haifa 32000, Israel

*e-mail: itai@cs.technion.ac.il
†e-mail: hadas@cs.technion.ac.il

### Proposed Running Head:

ITAI AND SHACHNAI, ADAPTIVE SOURCE ROUTING.

### Address for Proofs:

Prof. Alon Itai
Department of Computer Science
The Technion IIT, Haifa 32000
Israel.

**Abstract**

We study algorithms for adaptive source routing in high speed networks, where some of the links are unreliable. Thus, the delivery of a single message to its destination may require trying several paths. Assuming *a priori* knowledge of the failure probabilities of links, our objective is to devise routing strategies which minimize the expected delivery cost of a single message.

We describe optimal strategies for two cases: A tree-like network and a general serial/parallel topology. Whereas in the first case, the greedy algorithm is shown to be optimal (i.e., it is best to try the paths by decreasing order of their success probabilities), there is no simple decision rule for the second case. However, using some properties of serial/parallel graphs we show, that an optimal strategy can be easily derived from a fixed sequence of paths. We give an algorithm, polynomial in the number of links, for finding this sequence.

For a general network we show, that the problem of devising an optimal strategy can be solved in polynomial space and is #P-Hard, and that the minimal expected delivery cost in a given network is also hard to approximate. Finally, we show, that for scenarios of adaptive source routing, the common greedy strategy is not even a constant approximation to the optimal strategy.

# 1 Introduction

Routing protocols in networks of fixed topologies are typically evaluated by the time required to deliver a message, where dominant factors are the length of the path selected for delivery and delays caused by queuing. However, it is often the case, that even though no topology updates are initiated, at any time some links may malfunction. Thus, the efficiency with which a routing protocol adapts to the (unpredictable) changes in topology becomes a significant measure of performance.

In this paper we focus on the problem of adaptive source routing in high speed networks [5, 2, 6]. The nodes in these networks contain switching elements with limited computational resources. The usage of source routing allows a wide range of routes to be used in the network without burdening the intermediate nodes with agreeing on, or precomputing, all possible routing choices. Thus, the delivery of each message starts with a specification of the entire path from the source to the destination. If all the links on that path are nonfaulty then the message is transmitted; otherwise an indication is given to the first faulty link. The source, upon learning of a failure, chooses another path and retransmits the message. Since the transmission is very fast, the time it takes to deliver a message is proportional to the number of paths tried (i.e., it does not depend on the lengths of these paths).

Therefore, the problem of minimizing the expected cost of delivery reduces to devising a routing strategy which minimizes the expected number of paths tried for a single message.

We assume that recovery time is slow relative to sending a message. Thus, if while trying to route a single message a link has failed, it remains unoperational for the entire duration of the attempts to send this message. However, the failure patterns for the next message are independent of those of the current one. Hence, the trials for sending a single message are not statistically independent, whereas the trials for different messages are.

In the context of graph theory, we study the problem of searching an unknown graph: the above problem is a variation of the stochastic Canadian Traveler Problem ($CTP$) introduced in [12]. In the original problem a traveler is given a map, where each road has a probability of being untraversable. The goal is to minimize the expected length of the total distance traversed by the traveler towards destination. Papadimitriou and Yannakakis showed, that the problem of minimizing the expected ratio to the minimal path length is #P-Hard and solvable in polynomial space.

Other variations of this problem were studied in [4]: in the stochastic Recoverable CTP, each faulty link may recover. Bar-Noy and Schieber gave a polynomial-time strategy, which minimizes the expected delivery time of a message, if recovery times are not long relative to the transmission rates.

Topkis [15] considered the single source single destination routing problem, where paths are *link disjoint* and of equal lengths. The minimization criterion is the expected number of trials on paths for a single message. It was shown, that the optimal sequence of trials with respect to this measure, is by decreasing order of success probabilities on the paths, where the success probability of a path is the product of the success probabilities of its links. We show below that this strategy, which we call *greedy*, is optimal also for configurations with *dependent* paths, such as a path converging network (see definition in Section 3). Thus, Topkis' results fit a special case of our problem.

Bar-Noy et al. [3] considered the above model and showed that the problem of computing the expected cost of the optimal strategy is #P-Hard.

Some of the previous works consider a model in which failure probabilities are *not* known, and the adaptive routing process dynamically learns the desirable order of trials on the paths, thereby decreasing, in each transmission, the expected number of trials in subsequent transmissions [1, 10, 14]. In the present work we assume, that failure probabilities are known, and that links do not fail or recover within a sequence of trials for a *single message*.

Our study of unreliable high-speed networks suggests an alternative measure for evaluating the performance of routing strategies as well as search schemes in geometric contexts: Consider the transmission of objects towards a target, in a setting with obstacles, where motion is fast, and decisions are made at the launching point. Further discussion of this aspect of the problem is given in [12].

We first study the greedy algorithm, which is known to be efficient in other routing scenarios [9]. We characterize topologies for which the greedy algorithm is optimal, and demonstrate by a simple counterexample its inefficiency for a large set of topologies. In the sequel, we focus on the subclass of serial/parallel (*sp*) networks and present an algorithm which constructs an optimal routing scheme in time polynomial in the number of links, for any network of this type. We view the study of sp networks as a first step in the investigation of other common subclasses of topologies, such as layered graphs and planar networks. We conclude with some complexity results concerning the hardness and the approximability of the above optimization problem for general networks.

The rest of this paper is organized as follows:

Section 2 presents in detail the above network model and the routing approach adopted in the following sections. In Section 3 we show the optimality of the greedy strategy for *path converging* networks. For a general topology we show, that the greedy strategy does not provide a *constant* approximation to the optimal strategy.

In Section 4 we focus on the subclass of *serial/parallel* networks. Section 4.1 gives some properties of these networks, which enable us to reconstruct any optimal routing strategy from a sequence of at most $|E|$ $s$-$t$ paths, In Section 4.2 we present an $O(|E|^3)$ algorithm, $A_{sp}$, to find this sequence. We show that the sequence for the parallel composition of $G_1$ and $G_2$ (see definitions below) is a merge of the sequences of each of the $G_i$'s. The details concerning the implementation of $A_{sp}$ are given in [8].

In Section 5 we discuss general topologies, for which we show, that the problem of devising a strategy which minimizes the expected delivery time of a message is in PSPACE and #P-Hard. We also show, that there is no polynomial time approximation for the minimal expected cost of delivering a single message unless P=NP. Finally, in Section 6 we present some open problems.

# 2 Preliminaries

## 2.1 Model Formulation

In our model, the *Known Failures* model, a network is identified with a directed graph $G = (V, E)$, vertices $s, t \in V$ (the source and the destination), and success probabilities on the links – $p(e)$ is the probability that link $e$ has not failed. These probabilities are independent, and thus the success probability of a path is the product of the success probabilities of its links. Note that the success probabilities of paths are *not* independent. Without loss of generality, each link in $G$ belongs to some simple path from $s$ to $t$ (i.e. an $s$-$t$ path with no cycles), and there are no nodes of indegree 1 and outdegree 1.

Given a graph $G$, with source $s$, a destination $t$, and a vector $\bar{p}$ of a priori success probabilities of the links, a *routing strategy* (or strategy, for short), $S(G, s, t, \bar{p}) = m$, is a function, that returns the next path to be tried.

### 2.2 Modifying the Network Topology

We assume the following mode of operation: A vector $\bar{p}$ gives the a priori success probabilities of the links. To simplify the presentation, we assume, that $0 < p(e_i) < 1$. After each trial, if $m$ is an $s$-$t$ path, then we are done; otherwise, let $m = (e_1, \ldots, e_r)$ and $e_i = (v_{i-1}, v_i)$ be the first faulty link (i.e., the entire subpath $(s, v_1, ..., v_{i-1})$ is nonfaulty). We update the network topology as follows:

(i) Contract the subpath $(s, v_1, ..., v_{i-1})$, i.e., remove all links entering $v_1, ..., v_{i-1}$, and then replace any remaining link of the form $(v_j, u)$ (for $j < i$) by a link $(s, u)$ with the same success probability.

(ii) Remove the link $e_i$ from the network.

(iii) Remove all links that do not belong to any simple $s$-$t$ path.

(iv) If $v$ becomes a vertex with indegree and outdegree equal to 1, the pair of links $(u, v)$, $(v, w)$ is replaced by the link $(u, w)$ whose success probability is $p(u, w) = p(u, v) \cdot p(v, w)$.

Thus, given the a priori probabilities, the routing strategy $S$ finds a first path $m_1$. If $m_1$ fails, the process is continued on the new $(G, \bar{p})$ until an $s$-$t$ path is found or all paths have failed, in which case $S$ fails, and the message is *blocked*.

Let $E_S(G, \bar{p})$ denote the expected number of trials on paths under $S$ for a graph $G$ and success probabilities $\bar{p}$. Our objective is to find an *optimal routing strategy* – a strategy with the minimal expected number of trials, $E_{OPT}(G, \bar{p})$.

### 2.3 Trial Trees

Let $(G, \bar{p})$ be a Known Failures network with $n$ paths from $s$ to $t$. Any routing strategy $S$ on $G$ may be identified with a *trial tree* $\mathcal{T}$ constructed as follows:
The vertices of the tree correspond to the states of $S$, where each state signifies a series of trials on the network links with the recorded outcome (success/failure) for every trial. At each state, the strategy $S$ chooses the next path to be tested, with which we label the vertex. Each tree-edge corresponds to an outcome of the trial. If the path consists of $k$ links, then the corresponding vertex has $k + 1$ children, where the $i$-th child indicates that link $i$ was the first link to fail, and the $k + 1^{\text{st}}$ child indicates success. Hence, the *trial tree* prescribes a fixed series of trials on paths for every possible configuration of failures on the links.

Figure 1 gives an example of a network topology and a possible trial tree. A trial tree is *optimal* if it corresponds to an optimal strategy.

## 3 Optimality of the Greedy Algorithm for Path Converging Topologies

Our first result refers to a subclass of topologies, in which dependencies of paths do not decrease the total number of trials per message. Thus no strategy could do any better than the *greedy* algorithm,

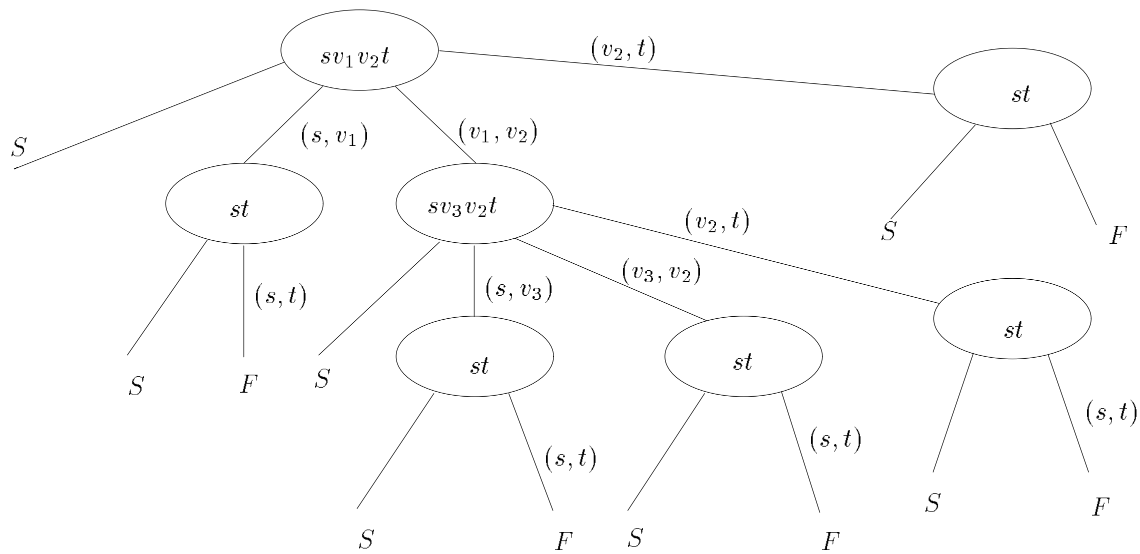


(a)

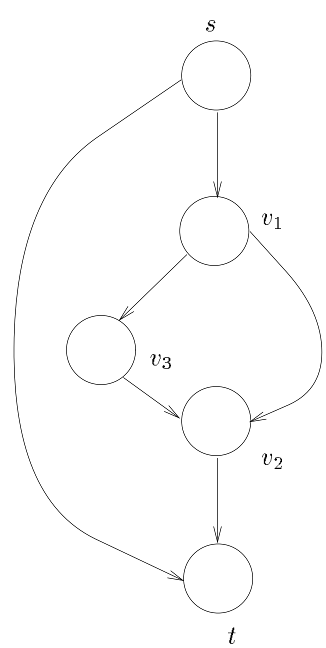


(b)

Figure 1: A Trial Tree (a) for *serial/parallel* Network (b)

which tries in each stage the path with maximal success probability (where $\bar{p}$ is updated as described in Section 2.2).

**Definition 1:** A network $G = (V, E)$ with the source and destination $s, t$ respectively is *path converging* from $s$ to $t$, if $G \setminus \{s\}$ consists of a tree rooted at $t$, and $s$ is connected to all the leaves.

**Lemma 1:** If the topology of $G$ is path converging from $s$ to $t$, then the greedy algorithm is optimal.

The proof uses Theorem 1 and a simple interchange argument – we give it in the Appendix.

Though the greedy approach was shown to be expedient in other routing scenarios (see ex. in [9]), in solving our optimization problem, the greedy algorithm may be far from optimal. We demonstrate it first by the following counterexample:

**Example 1:** Consider the network of Figure 2. The set of paths is
$\{(s, v_1, v_2, t), (s, v_1, v_3, t), (s, t)\}$,
and the success probabilities are

$$p_{(s,v_1)} = \frac{1}{2}\ ,\ \ p_{(s,t)} = \frac{1}{3}\ ,\ \ p_{(v_1,v_2)} = p_{(v_1,v_3)} = \frac{2}{3} - \varepsilon\ ,\ \ p_{(v_2,t)} = p_{(v_3,t)} = 1\ \ .$$

The path $(s, t)$ has the highest success probability, but under any optimal routing scheme this path is tried *last*, regardless of the states of the other paths in the network.

In fact, for a general topology, the greedy algorithm does not even provide a constant approximation to the optimal strategy, as suggested in the following

**Lemma 2:** For any $c > 1$ there exists a network $(G, \bar{p})$, such that

$$\frac{E_G(G, \bar{p})}{E_{OPT}(G, \bar{p})} > c\ \ , \tag{1}$$

where $E_G(G, \bar{p})$, $E_{OPT}(G, \bar{p})$ denote the expected number of trials on paths under the greedy algorithm and an optimal strategy respectively, for a network $G$ and the success probabilities given by the vector $\bar{p}$.

**Proof:** Consider the network $(G, \bar{p})$ in Figure 3: For simplicity, we present the set of $n_L$ disjoint paths from $s$ to $t$ by $n_L$ parallel links. Similarly, for the $n_R$ disjoint $s$-$u$ paths and for the $n_R$ disjoint paths from $u$ to $t$.
We assume the following success probabilities:

$$\begin{array}{lll} p(e_i) & = p & 1 \le i \le n_L \\ p(f_i) & = q > p & 1 \le i \le n_R \\ p(g_i) & = \frac{p}{q} - \varepsilon & 1 \le i \le n_R \end{array}$$

Denote by $E_S(G_L, \bar{p_L}) \equiv E_L$, $E_S(G_R, \bar{p_R}) \equiv E_R$, the expected number of trials in the left component (the $n_L$ parallel $s$-$t$ paths) and in the right component (the $n_R^2$ $s$-$t$ paths through the node $u$) respectively, under a strategy $S$. Observe, that the expected number of trials on each of these components

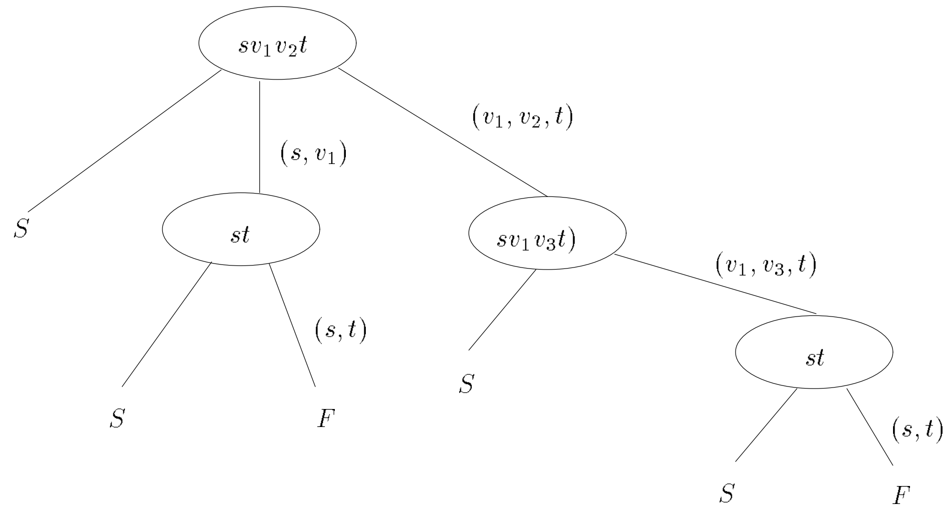


(a)

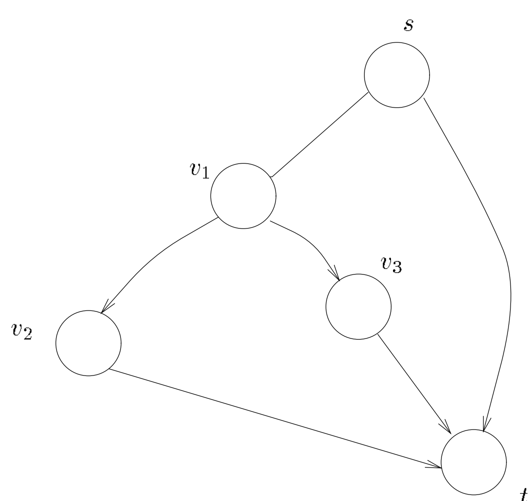


(b)

Figure 2: The Optimal Trial Tree (a) for Example 1 (b)

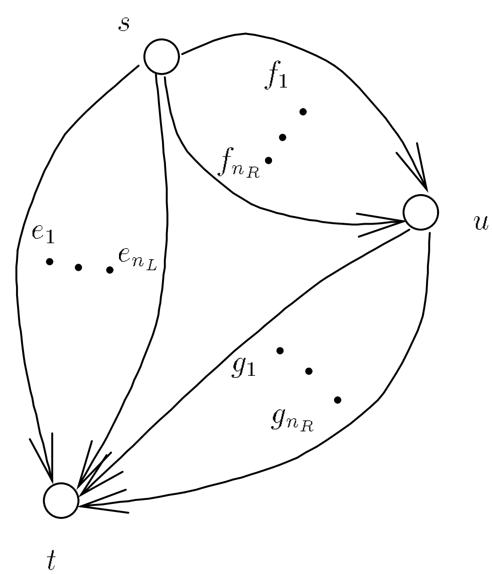


Figure 3: The Suboptimality of the Greedy Algorithm

does not depend on the strategy $S$.
Let $p_L$, $p_R$ be the probabilities that there is no $s$-$t$ path in the left and right components of $G$ respectively.
Obviously, the greedy algorithm tries first all the paths in the left component before passing to the right component, hence

$$E_G(G, \bar{p}) = E_L + p_L \cdot E_R \quad .$$

We denote by $H$ the strategy which tries first all the paths in the right component, before passing to the left component.
Then,

$$E_H(G, \bar{p}) = E_R + p_R \cdot E_L \quad ,$$

and

$$\frac{E_G(G, \bar{p})}{E_{OPT}(G, \bar{p})} \geq \frac{E_G(G, \bar{p})}{E_H(G, \bar{p})} = \frac{E_L + p_L \cdot E_R}{E_R + p_R \cdot E_L} \quad .$$

It is sufficient then to require, that

$$\frac{\frac{E_L}{E_R}}{1 + p_R \frac{E_L}{E_R}} > c \quad .$$

We now compute $E_L, E_R$ and $p_R$. Let the random variable $X_{s,t}$ denote the number of trials of $s$-$t$ links until finding the first nonfaulty link. We use the notation $\mathcal{G}(p, n)$ for the *truncated* Geometric distribution with a success probability $p$ and where the maximal number of trials is $n$. Then $X_{s,t} \sim$

$\mathcal{G}(p, n_L)$,
and

$$E_L = \quad E(X_{s,t}) = \quad \tfrac{1-(1-p)^{n_L}}{p}$$

We observe, that a transmission of a message in the right component involves finding a nonfaulty $s$-$u$ link $l$, and then testing $s$-$t$ paths on which $l$ is the first link, until finding a nonfaulty $u$-$t$ link. The probability that there is a nonfaulty link between $s$ and $u$ is $(1-(1-q)^{n_R})$. Let $X_{s,u}$ ($X_{u,t}$) be the number of trials of $s$-$u$ ($u$-$t$) links until the first nonfaulty link is found, then $X_{s,u} \sim \mathcal{G}(q, n_R)$, $X_{u,t} \sim \mathcal{G}(p/q - \varepsilon, n_R)$, and

$$\begin{aligned} E_R &= E(X_{s,u}) + (1-(1-q)^{n_R}) \cdot E(X_{u,t}) \\ &= \sum_{j=0}^{n_R-1} (1-q)^j + (1-(1-q)^{n_R}) \sum_{j=0}^{n_R-1} (1-\frac{p}{q}+\varepsilon)^j \\ &= \frac{1-(1-q)^{n_R}}{q} + (1-(1-q)^{n_R}) \cdot \frac{1-(1-\frac{p}{q}+\varepsilon)^{n_R}}{\frac{p}{q}-\varepsilon} \end{aligned}$$

We note, that $p_R$ is the probability that all $s$-$u$ links are faulty or there exists a nonfaulty $s$-$u$ link, with all the $u$-$t$ links faulty.
Thus,

$$p_R = \quad (1-q)^{n_R} + (1-(1-q)^{n_R}) \cdot (1-\tfrac{p}{q}+\varepsilon)^{n_R} \quad .$$

Let $x = E_L/E_R$, then

$$\frac{1}{x} < \left(\frac{1}{q} + \frac{1}{\frac{p}{q}-\varepsilon}\right) \cdot p \ .$$

For

$$q = \frac{3(1-4c\varepsilon)}{4c}, \qquad p = \frac{3(1-4c\varepsilon)}{(4c)^2}, \qquad \varepsilon < \frac{1}{8c}$$

and sufficiently large $n_R$, we have

$$\frac{1}{x} + p_R < \frac{1}{c} \ ,$$

thus

$$\frac{E_G(G,\bar{p})}{E_H(G,\bar{p})} > c \quad .$$

□

# 4 SP - Networks

In this section we discuss a subclass of topologies for which there exists an optimal trial tree, in which all the orders of trials of paths are derived from a unique permutation $\pi$ of $E$, the set of network links. Hence, the problem of constructing a *trial tree* reduces to finding $\pi$.

**Definition 2:** Let $(G_1, s_1, t_1)$, $(G_2, s_2, t_2)$ be node disjoint networks, with $s_i, t_i \in V(G_i)$:
$G_1 \Longrightarrow G_2$, the *serial* composition of $G_1$ and $G_2$, is the union of $G_1$ and $G_2$ where $t_1$ is identified with $s_2$. $G_1||G_2$, the *parallel* composition of $G_1$ and $G_2$, is the union of $G_1$ and $G_2$, where $s_1$ is identified with $s_2$ and $t_1$ is identified with $t_2$.
$G$ is a *serial/parallel (sp)* network, if it can be constructed using only the operations $\Longrightarrow$ and $||$ from networks containing a single link. (See also [7, 16]).

**Lemma 3:** The update procedure in Section 2.2, when applied to an sp network $G$ of $|E|$ links, produces an sp network of at most $|E| - 1$ links.

For any $e \in E$, we denote by $G \cdot e$ and $G - e$ the networks resulting from the contraction and deletion of $e$ respectively. The proof of the lemma follows from the fact that if $G$ is an sp network, then so are $G \cdot e$ and $G - e$. This can be shown by induction on $|E|$.
We introduce below some properties of optimal trial trees for sp networks.

### 4.1 Identifying Strategies with Complete Orders on Paths

A *fault configuration* of a network is a partition of its links into faulty and nonfaulty links. Along a series of trials, links are *tested* and their status (faulty or nonfaulty) becomes known. A fault configuration is *complete* for strategy $S$, if during the application of $S$ with this configuration every link of $G$ is tested.
In the following discussion, we assume, that the success probabilities on the links satisfy $0 < p(e) < 1$, $\forall\, e \in E$. (We assumed this holds initially, and obviously, the update procedure of Section 2.2 maintains this property.)

**Lemma 4:** Let $(G, \bar{p})$ be a network, and $S$ a strategy. Then there exists a complete fault configuration on $(G, \bar{p})$ for $S$. Moreover,there exist exactly two such fault configurations, which differ only in the status of the last link tested by $S$.

**Proof:** Existence: By induction on the number of links in $G$.
*Basis*: The result is obvious for a network consisting of a single link.
*Induction step*: Let $l = (s, v)$ be the first link tried by $S$. If $G - l$ contains an $s$-$v$ path then we make $l$ fail. Otherwise, we make $l$ nonfaulty. In either case, we obtain a graph $G'$ as described in Section 2.2. By construction, $G'$ has one link less than $G$, and thus we can use the induction hypothesis.
W.l.o.g., we may assume, that the last link tested by $S$ is faulty, then it remains to show that there is a unique complete fault configuration for $S$.
Uniqueness: Let $\mathcal{C}_1$, $\mathcal{C}_2$ be complete fault configurations and $l = (u, v)$ the first link tested by $S$ for which $\mathcal{C}_1$, $\mathcal{C}_2$ differ. Observe, that since the strategy $S$ is deterministic, the order of trials on the network links is identical under $\mathcal{C}_1$, $\mathcal{C}_2$ until the link $l$ is tested. Without loss of generality, we assume, that $l$ is nonfaulty under $\mathcal{C}_1$ and faulty under $\mathcal{C}_2$. By the above assumption, $l$ is not the last link tested by $\mathcal{C}_1$, $\mathcal{C}_2$, therefore $v \neq t$.
We handle two cases:

(i) All $s$-$v$ paths have already been tested (and were found faulty) by $\mathcal{C}_1$, $\mathcal{C}_2$: Hence, under the configuration $\mathcal{C}_2$ the strategy $S$ will never reach $v$ again, and all links emanating from $v$ will not be tested by $S$ under $\mathcal{C}_2$, contradicting the completeness of $\mathcal{C}_2$.

(ii) At least one $s$-$v$ path was not tested: Then, since under $\mathcal{C}_1$ no other $s$-$v$ path will be tested, $\mathcal{C}_1$ is not complete – a contradiction.

□

Let $\mathcal{C}$ be a complete fault configuration on $(G, \bar{p})$ for the strategy $S$, denoted by $cfc_S(G, \bar{p})$. Then we order the links of $G$ in the order by which they were tested under this configuration. By Lemma 4 the order is unique. We denote it by $\pi_L(G, S) = (l_1, \ldots, l_{|E|})$, and call it the *complete order of trials on links* for strategy $S$. An order $(l_1, \ldots, l_{|E|})$ is an *optimal complete order*, if it is the complete order on the links for some optimal strategy $S^*$.

**Definition 3:** A sequence of paths $(m_1, \ldots, m_k)$ is a *complete order on paths for strategy S*, if it is the sequence chosen by $S$ for the complete fault configuration on links, $cfc_S(G, \bar{p})$. We denote this sequence of paths by $\pi_M(G, S)$.

**Definition 4:** A link $l_1$ *precedes* a link $l_2$ in the trial tree $\mathcal{T}$, if there exists a sequence of trials on paths given by $\mathcal{T}$, in which $l_1$ is tested before $l_2$. (Note that it is possible that $l_1$ precedes $l_2$ because of one sequence of trials and on another sequence $l_2$ precedes $l_1$). A trial tree $\mathcal{T}$ is *compatible* with the order $\prec$, if $l_i \prec l_j$ implies, that in every sequence of trials given by $\mathcal{T}$, $l_i$ is tested before $l_j$.

**Definition 5:** $\widehat{G}$ is a *minimal serial component*, if in any presentation of $\widehat{G}$ of the form $\widehat{G} = (G_1 \Longrightarrow G_2)$ , either $G_1 = \emptyset$ or $G_2 = \emptyset$. A *minimal parallel component* is defined similarly.

In the remainder of the paper, we assume an optimal strategy $S^*$ and denote $\pi_M(G, S^*)$ and $\pi_L(G, S^*)$ by $\pi_M^*(G)$ and $\pi_L^*(G)$.

**Definition 6:** The subgraph $\tilde{G}$ of $G$ is a *prefix of G with respect to an optimal strategy* $S^*$, if there is a prefix of $\pi_M^*(G)$ in which every path consists solely of links in $\tilde{G}$.

Given the two permutations $\pi_1 = (i_{11}, \ldots, i_{1r})$, $\pi_2 = (i_{21}, \ldots, i_{2s})$ on disjoint sets $(\{i_{11}, ..., i_{1r}\} \cap \{i_{21}, ..., i_{2s}\} = \emptyset)$, we denote by $\pi_1 \circ \pi_2$ the concatenation of the permutations, i.e., the sequence $(i_{11}, \ldots, i_{1r}, i_{21}, \ldots i_{2s})$.

**Lemma 5:** Let $G_1, G_2$ be sp networks, and $S_1^*, S_2^*$ be optimal strategies for $G_1, G_2$, with the corresponding complete orders on the links $\pi_L^*(G_1)$ and $\pi_L^*(G_2)$. Then $\pi_L^*(G_1) \circ \pi_L^*(G_2)$ is an optimal complete order for $G_1 \Longrightarrow G_2$.

**Proof:** Observe, that in any complete fault configuration corresponding to an optimal strategy $S^*$ for $G_1 \Longrightarrow G_2$, all the links in $G_1$ are tested before $S^*$ tests any of the links in $G_2$. In any other case, it can be easily shown that $S^*$ is not optimal. The lemma follows from the optimality of $S_1^*$ and $S_2^*$. □

Thus it remains to see how to deal with parallel compositions. In the following lemma we give three properties of *sp* networks, that will enable us to find an algorithm for recursive construction of an optimal strategy.

**Lemma 6:** For any *sp* network $(G, \bar{p})$:

(i) There exists an optimal trial tree $\mathcal{T}^0$, in which for any two links $l_1$ and $l_2$ either $l_1$ always precedes $l_2$ or vice versa.

Let $S^*$ denote the optimal strategy corresponding to $\mathcal{T}^0$, with the complete order on links $\pi_L^*(G)$. For any subgraph $G'$ of $G$, we denote by $\pi_L^*(G)_{|G'}$ the thinning of $\pi_L^*(G)$ that includes only the links of $G'$.

(ii) If $G = (G_1 \parallel G_2)$, and

$$\pi_L^*(G) = \pi_L^*(G_1) \circ \pi_L^*(G_2) \ , \tag{2}$$

then for every prefix $\tilde{G}_2$ of $G_2$, $\pi_L^*(G_1) \circ \pi_L^*(G_2)_{|\tilde{G}_2}$ is an optimal complete order for the graph $G_1 \parallel \tilde{G}_2$.

(iii) If $G = (G_1 \parallel G_2)$, and there exists a prefix $\tilde{G}_1$ of $G_1$, such that

(a)

$$\pi_L^*(G) = \pi_L^*(\tilde{G}_1) \circ \pi_L^*(G - \tilde{G}_1) \ , \tag{3}$$

(where $G - \tilde{G}_1$ is the subgraph of $G$ resulting from the deletion of the links in $\tilde{G}_1$);

(b) $\tilde{G}_1$ is the *maximal* prefix of $G_1$ satisfying (3),

then $\pi_L^*(\tilde{G}_1) \circ \pi_L^*(G - \tilde{G}_1)_{|G_2}$ is an optimal complete order for $\tilde{G}_1 \parallel G_2$.

**Proof:** By induction on the number of links in $G$.
The base case is $|E| = 0$.
For the induction step, we assume that properties $(i)$-$(iii)$ hold in a network with smaller number of links and show that $G$ satisfies property $(i)$. Then we use property $(i)$ for $G$ and the induction hypothesis to show, that $G$ satisfies properties $(ii)$ and $(iii)$.
To simplify the proof we use an equivalent representation of a trial tree: Each node in the original tree (which corresponds to a trial of a given path) is replaced by a sequence of trials of the links on that path. Thus, the trial tree becomes a *decision tree* (see Figure 4).

(i) We construct $\mathcal{T}^0$ from two subtrees as follows: Let $OUT(s)$ denote the set of links emanating from $s$, i.e.,

$$OUT(s) = \{(s, v_1) \mid v_1 \in V \setminus \{s\} \ , \ (s, v_1) \in E\} \ .$$

For any link $l = (s, v_1) \in OUT(s)$ let $S_l$ be a strategy of minimal cost, among all the strategies that test $l$ first. Note that the optimal strategy $S^*$ satisfies

$$E_{S^*}(G, \bar{p}) = \min\{E_{S_l}(G, \bar{p}) : l \in OUT(s)\} \ . \tag{4}$$

We denote by $G_l, G_{\neg l}$ the graphs resulting from that trial, where $l$ is nonfaulty and faulty respectively. By Lemma 3, $G_l, G_{\neg l}$ are sp networks. Using the induction hypothesis, there exist two optimal trial trees $\mathcal{T}_l^0, \mathcal{T}_{\neg l}^0$ for $G_l$ and $G_{\neg l}$. We show the existence of $\mathcal{T}_l^0, \mathcal{T}_{\neg l}^0$ that are mutually compatible, with respect to the relative order by which the links are tested. Let $\mathcal{T}^0(l)$ be the

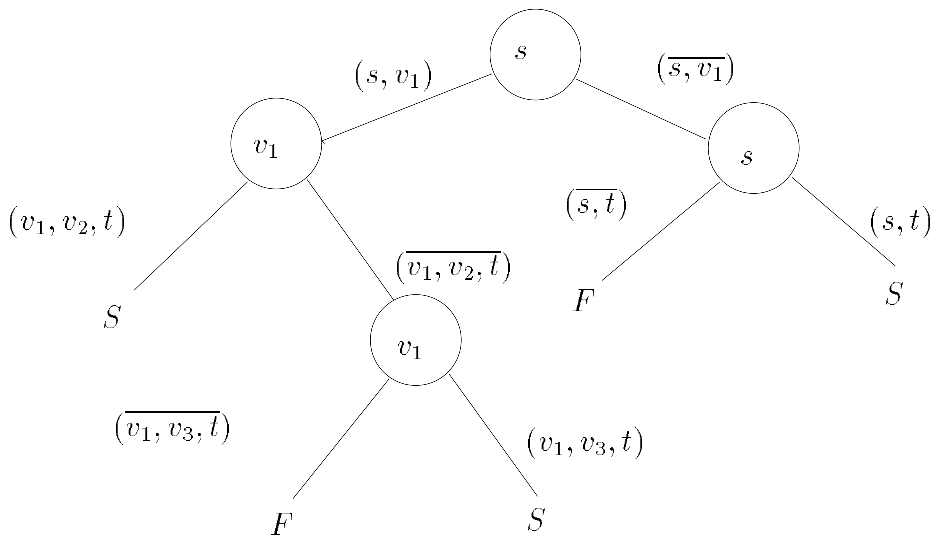


Figure 4: The Optimal Decision Tree for Example 1

decision tree consisting of a root connected to the roots of $\mathcal{T}_l^0$ and $\mathcal{T}_{\neg l}^0$, and the edges emanating from the root of $\mathcal{T}^0(l)$ are labeled $l$ and $\bar{l}$. Using the fact that the trees $\mathcal{T}_l^0$ and $\mathcal{T}_{\neg l}^0$ are optimal, we identify the tree $\mathcal{T}^0(l)$ with $S_l$, for every $l \in OUT(s)$, and choose $\mathcal{T}^0$ to be the tree $\mathcal{T}^0(l)$ for which the expected number of trials is minimal, i.e., the corresponding strategy $S^*$ satisfies (4). We start by showing

**Claim 1:** There exist optimal trial trees $\mathcal{T}_l^0, \mathcal{T}_{\neg l}^0$ for $G_l, G_{\neg l}$, such that for any pair of links $l_i, l_j$ in $G_l \cap G_{\neg l}$, if $l_i \prec l_j$ in $\mathcal{T}_{\neg l}^0$ then $l_i \prec l_j$ in $\mathcal{T}_l^0$.

**Proof:** We handle the following two cases separately:
Case 1. There is no alternative path from $s$ to $v_1$. We further distinguish between the following sub-cases:

(a) $l$ is the only link emanating from $s$, then $G_{\neg l} = \emptyset$ and the claim holds.

(b) There exist $\mathcal{T}_l^0, \mathcal{T}_{\neg l}^0$, such that the corresponding $\pi_L^*(G_{\neg l})$ is a suffix of $\pi_L^*(G_l)$. I.e., if $\pi_L^*(G_l) = (i_1, \ldots, i_r)$ (where $r$ is the number of links in $G_l$), then $\pi_L^*(G_{\neg l}) = (i_j, \ldots, i_r)$, for some $1 < j \le r$. Thus the trees $\mathcal{T}_l^0$ and $\mathcal{T}_{\neg l}^0$ satisfy the statement in the claim.

(c) Neither (a) nor (b). In this case $v_1 \neq t$.
Observe, that the links emanating from $s$ in $G_l$ either belong to $G$, (and hence belong also to $G_{\neg l}$) or are derived from links that emanated from $v_1$. Also, $v_1 \notin G_{\neg l}$.
Let $\widehat{G}_l$ be the minimal serial component of $G_l$ containing all the links emanating from $s$ in $G_l$. $\widehat{G}_l$ can be represented as $\widehat{G}_l = G_{l,1} \parallel G_{l,2}$, and $G_{l,1}$ contains all the links that emanated in $G$ from $v_1$ (see Figure 5). Since $v_1 \neq t$, $G_{l,1} \neq \emptyset$ and $G$ contains $s$-$t$ paths that avoid $v_1$. Hence $G_{l,2} \neq \emptyset$.

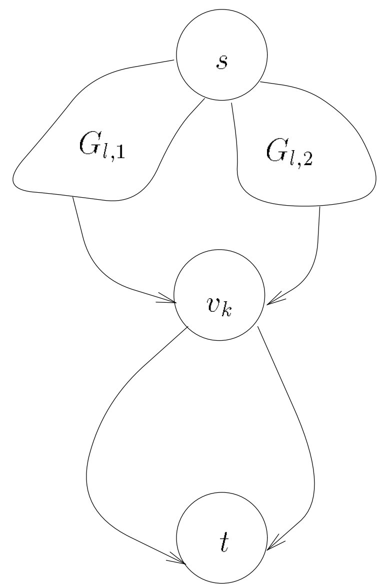


Figure 5: The Graph $G_l$ obtained after the trial of $(s, v_1)$ in Case 1 $(c)$

By applying property $(iii)$ of the induction hypothesis to $\widehat{G}_l$, if $\widetilde{G}_{l,1}$ is the prefix of $G_{l,1}$ which is to be tried before passing to $G_{l,2}$ under some optimal complete order $\pi^*_L(\widehat{G}_l)$ for $\widehat{G}_l$, then

$$\pi^*_L(\widetilde{G}_{l,1}) \circ \pi^*_L(\widehat{G}_l)_{|G_{l,2}}$$

is an optimal complete order for the subgraph $\widetilde{G}_{l,1} \parallel G_{l,2}$. This implies that $\pi^*_L(\widehat{G}_l)_{|G_{l,2}}$ is an optimal complete order for $G_{l,2}$, since $G_{l,2}$ is obtained after trying the prefix $\widetilde{G}_{l,1}$. In addition, $G_{\neg l}$ is a serial composition $G_{l,2} \Longrightarrow G_3$ for some subgraph $G_3$ of $G$. Assume that $\pi^*_L(G_3)$ is an optimal complete order for $G_3$ , then by Lemma 5 $\pi^*_L(\widehat{G}_l)_{|G_{l,2}} \circ \pi^*_L(G_3)$ is an optimal complete order for $G_{\neg l}$. (In the special case where $G_3 = \emptyset$ we have $\pi^*_L(G_3) = \emptyset$). This order is consistent with $\pi^*_L(\widehat{G}_l) \circ \pi^*_L(G_3)$, that is optimal for $G_l$. Hence there exist $\mathcal{T}^0_l, \mathcal{T}^0_{\neg l}$, that are mutually compatible.

<u>Case 2</u>. There exists an alternative path from $s$ to $v_1$.
Define $G_l, G_{\neg l}$ as in Case 1, then $G_l$ is obtained by updating the graph when $l$ is nonfaulty, and $G_{\neg l}$ is obtained by omitting only $l$ from $G$. Hence, identifying the links in $G_l$, that emanate from $s$ due to the contraction of $l$, with the corresponding links (that emanate from $v_1$) in $G_{\neg l}$, we note that every link of $G_l$ also belongs to $G_{\neg l}$.
Let $\widehat{G}_{\neg l}$ be the minimal serial component of $G_{\neg l}$, such that $\widehat{G}_{\neg l} = G_{\neg l,1} \parallel G_{\neg l,2}$ and $G_{\neg l,1}$ consists of all the paths in $\widehat{G}_{\neg l}$ which include $v_1$, and let $\widehat{G}_l = G_{l,1} \parallel G_{l,2}$ and the subgraph $G_3$ be defined as in Case 1 $(c)$.
We denote by $\pi^*_L(G_{\neg l,1})$,$\pi^*_L(G_{\neg l,2})$ some optimal complete orders for $G_{\neg l,1}, G_{\neg l,2}$ respectively, then $\pi^*_L(G_{\neg l,1})_{|G_{l,1}}$ is an optimal complete order for $G_{l,1}$, and $\pi^*_L(G_{\neg l,2})$ is an optimal complete order

for $G_{l,2}$. Using properties $(i)-(iii)$ for $G_l, G_{\neg l}$, there exists an optimal complete order $\pi_L^*(\widehat{G}_{\neg l})$ for $\widehat{G}_{\neg l}$, such that $\pi_L^*(\widehat{G}_{\neg l})_{|\widehat{G}_l}$ is an optimal complete order for $\widehat{G}_l$. Then the optimal complete orders for $G_l, G_{\neg l}$, given by $\pi_L^*(\widehat{G}_{\neg l})_{|\widehat{G}_l} \circ \pi_L^*(G_3)$ and $\pi_L^*(\widehat{G}_{\neg l}) \circ \pi_L^*(G_3)$ respectively, are mutually compatible. $\square$

Using Claim 1, the tree $\mathcal{T}^0(l)$ can be constructed from the compatible subtrees $\mathcal{T}_l^0, \mathcal{T}_{\neg l}^0$ by attaching their roots to the endpoints of the edges $l, \neg l$ that emanate from $s$. Let $S_l$ be the strategy corresponding to $\mathcal{T}^0(l)$, and let $l^*$ satisfy

$$E_{S_{l^*}}(G,\bar{p}) = \min\{E_{S_l}(G,\bar{p}) : l \in OUT(s)\} \ .$$

Then $\mathcal{T}^0 = \mathcal{T}^0(l^*)$.

(ii) Let $G = G_1 \parallel G_2$, and $\pi_L^*(G) = \pi_L^*(G_1) \circ \pi_L^*(G_2)$ be a complete order compatible with some optimal tree $\mathcal{T}^0$ for $G$. Let $\tilde{G}_2$ be any prefix of $G_2$, and $perm(\pi_L^*(G_1) \circ \pi_L^*(G_2)_{|\tilde{G}_2})$ be any permutation of $(\pi_L^*(G_1) \circ \pi_L^*(G_2)_{|\tilde{G}_2})$.
Let the order

$$\mathrm{perm}(\pi_L^*(G_1) \circ \pi_L^*(G_2)_{|\tilde{G}_2}) \circ \pi_L^*(G_2)_{|G_2-\tilde{G}_2}$$

be a complete order on the links of $G$ for some strategy $S$, then from the optimality of $S^*$,

$$E_{S^*}(G,\bar{p}) \le E_S(G,\bar{p}) \ . \tag{5}$$

Denote by $\mathcal{C}_{G_1\parallel\tilde{G}_2}$ the set of all configurations of failures on the links in $G_1 \parallel \tilde{G}_2$, and let $\mathcal{C} \in \mathcal{C}_{G_1\parallel\tilde{G}_2}$. We define the indicator function $f$ as follows: $f(\mathcal{C}, G) = 1$ if all paths in $G$ are faulty under $\mathcal{C}$, and 0 otherwise. We use the notation $E_S(G,\bar{p}\,|\,\mathcal{C})$ for the expected number of trials under the strategy $S$ for $G$, with the fault configuration $\mathcal{C}$.
Using the subtree of $\mathcal{T}^0$ that includes only links in $G_1 \parallel \tilde{G}_2$, the expected number of trials for $S^*$ is given by

$$E_{S^*}(G,\bar{p}) = E_{S^*}(G_1 \parallel \tilde{G}_2, \bar{p}) + \sum_{\mathcal{C}_i \in \mathcal{C}_{G_1\parallel\tilde{G}_2}} f(\mathcal{C}_i, G_1 \parallel \tilde{G}_2) \cdot E_{S^*}(G_2 - \tilde{G}_2, \bar{p}\,|\,\mathcal{C}_i) \cdot \Pr(\mathcal{C}_i) \ . \tag{6}$$

We proceed by evaluating the right hand side of Inequality (5). Observe, that the inequality holds for any trial tree that corresponds to the strategy $S$, therefore we may assume a trial tree $\mathcal{T}_S$, that is compatible with the order $\mathrm{perm}(\pi_L^*(G_1) \circ \pi_L^*(G_2)_{|\tilde{G}_2}) \circ \pi_L^*(G_2)_{|G_2-\tilde{G}_2}$. Using the subtree of $\mathcal{T}_S$ that includes only links in $G_1 \parallel \tilde{G}_2$, and the fact that $S^*$ and $S$ use the same complete orders on the links for the subgraph $G_2 - \tilde{G}_2$, we have

$$E_S(G,\bar{p}) = E_S(G_1 \parallel \tilde{G}_2, \bar{p}) + \sum_{\mathcal{C}_i \in \mathcal{C}_{G_1\parallel\tilde{G}_2}} f(\mathcal{C}_i, G_1 \parallel \tilde{G}_2) \cdot E_{S^*}(G_2 - \tilde{G}_2, \bar{p}\,|\,\mathcal{C}_i) \cdot \Pr(\mathcal{C}_i) \ . \tag{7}$$

Substituting (6) and (7) into Inequality (5) we have the desired property for $G_1 \parallel \tilde{G}_2$.

(iii) Using the induction hypothesis $(ii)$ for $\tilde{G}_1 \parallel G_2$, if $\pi_L^*(\tilde{G}_1) \circ \pi_L^*(G - \tilde{G}_1)_{|G_2}$ is optimal for $\tilde{G}_1 \parallel G_2$, then for any prefix $\tilde{G}_2$ of $G_2$, the order $\pi_L^*(\tilde{G}_1) \circ \pi_L^*(G - \tilde{G}_1)_{|\tilde{G}_2}$ is optimal for $\tilde{G}_1 \parallel \tilde{G}_2$.

Assume by way of contradiction, that there is a prefix $\tilde{G}_2$ of $G_2$, such that the complete order $\pi_L^*(\tilde{G}_1)\circ\pi_L^*(G-\tilde{G}_1)_{|\tilde{G}_2}$ is not optimal for the subgraph $\tilde{G}_1 \parallel \tilde{G}_2$.
This implies that there exists a strategy $S'$ for $\tilde{G}_1 \parallel \tilde{G}_2$, whose complete order on the links is

$$\pi_L(\tilde{G}_1 \parallel \tilde{G}_2, S') = \text{perm}(\pi_L^*(\tilde{G}_1) \circ \pi_L^*(G-\tilde{G}_1)_{|\tilde{G}_2}) \ ,$$

such that

$$E_{S'}(\tilde{G}_1 \parallel \tilde{G}_2, \bar{p}) < E_{S^*}(\tilde{G}_1 \parallel \tilde{G}_2, \bar{p}) \ . \tag{8}$$

Let

$$\pi_L(G, S') = \text{perm}(\pi_L^*(\tilde{G}_1) \circ \pi_L^*(G-\tilde{G}_1)_{|\tilde{G}_2})\circ\pi_L^*(G)_{|G-\tilde{G}_1-\tilde{G}_2}$$

be the complete order on the links for $G$ under $S'$, and let $\mathcal{T}_{S'}$ be a trial tree that is compatible with this order.
We note, that $\mathcal{T}_{S'}$ can be obtained from $\mathcal{T}^0$ by replacing the subtree in $\mathcal{T}^0$ that includes the links in $\tilde{G}_1 \parallel \tilde{G}_2$ with the corresponding subtree in $\mathcal{T}_{S'}$. Using Inequality (8) and a computation similar to the computation given in (6) and (7), we have

$$E_{S'}(G, \bar{p}) < E_{S^*}(G, \bar{p}) \ .$$

Contradiction to the optimality of $S^*$.

$\square$

Combining the three properties in Lemma 6 we have the following *merge* property:

**Corollary 1:** For any parallel composition of $G_1, G_2$, with the optimal orders on paths $\pi_M^*(G_1)$,$\pi_M^*(G_2)$ respectively, there exists an optimal order $\pi_M^*(G)$ that is consistent with both $\pi_M^*(G_1)$ and $\pi_M^*(G_2)$.

The following theorem shows how to reduce the search space of the optimal trial trees.

**Theorem 1:** For any sp graph there exists an optimal strategy $S^*$, such that the corresponding trial tree is compatible with $\pi_L(G, S^*)$.

The theorem follows from property $(i)$ in Lemma 6.

Observe, that by Definition 3, if $\pi_L(G, S^*)$ is the complete order on links under an optimal strategy $S^*$ for $G$, then $\pi_M^*(G) = \pi_M(G, S^*)$ is the set of $s$-$t$ paths tried within the sequence $\pi_L(G, S^*)$. Therefore, Theorem 1 implies, that if we know the complete order on paths, $\pi_M^*(G)$, for some unknown optimal strategy $S^*$, then we can reconstruct $S^*$ (i.e., determine the trials on paths made by $S^*$ under any possible scenario of failures) as follows:

Reconstruct a Strategy $S^*$ from $\pi_M^*(G)$

While $\pi_M^*(G) \neq \emptyset$ do

1. Try the first path in $\pi_M^*(G)$.
2. Upon failure, $\bar{p}$ and $G$ are updated as described in Section 2, and $\pi_M^*(G)$ changes in the following manner:

(a) If the link $e_i$ was found faulty in the last trial, then all the paths in $\pi^*_M(G)$ containing it are deleted from $\pi^*_M(G)$,

(b) If $e_j \notin E$ (i.e., $e_j$ was contracted after the last trial), then it is deleted from all the paths in which it appeared in $\pi^*_M(G)$.

(c) For any two links $(u, v), (v, w)$ in $G$, that were concatenated into the single link $(u, w)$ (see $(iv)$ in Section 2.2), we update the paths including $(u, v)$ and $(v, w)$ to include the link $(u, w)$, that is, on those paths we replace the subpath $(u, v, w)$ by the single link $(u, w)$. (Observe, that our update procedure guarantees, that if $(u, v), (v, w)$ are concatenated into a single link then there is no path in $\pi_M(G)$ that contains only one of these links, or in which the two links are not consecutive).

Hence, for any sp graph $(G, \bar{p})$, it is sufficient to find a complete order on paths, which corresponds to an optimal strategy.
Given $\pi_M(G, S)$, let $\tau(\pi_M(G, S))$ be the strategy resulting from the above construction. We proceed with a proof of correctness.

**Theorem 2:** For any strategy $S$,

$$S = \tau(\pi_M(G, S)) \quad . \tag{9}$$

**Proof:** Let $S' = \tau(\pi_M(G, S))$, then both $S$ and $S'$ try first the same path, $m_1 = (e_1, \ldots, e_l)$, in $\pi_M(G, S)$ ($e_i$ denotes the $i$th link on the path). Denote by $e_i$ the first faulty link in $m_1$. Let $m_j$ be the first path in $\pi_M(G, S)$ not containing $e_i$. By the above construction, the next path tried by $S'$ is $m'_j$ resulting from $m_j$ by contraction of the nonfaulty links $e_1, \ldots, e_{i-1}$.
Assume that the next path tried by $S$ is $m \neq m'_j$. Denote by $v_j$ the last vertex in $G$ common to both $m$ and $m'_j$, then $e = (v_j, u)$ is the first link in $m - m'_j$ and $e' = (v_j, u')$ is the first link in $m'_j - m$. Since both $S$ and $S'$ are compatible with $\pi_L(G, S)$ for $cfc_S(G, \bar{p})$, it follows that in $\pi_L(G, S) \quad e' \prec e$, therefore after the trial of $e$ in the above fault configuration ($e_1, \ldots, e_{i-1}$ are nonfaulty and $e_i$ is faulty), $S$ never tries $e'$.
Now, assume a configuration of failures in which $e_1, ..., e_{i-1}$ are nonfaulty, $e_i$ is faulty and the only nonfaulty path includes $e$, then $S$ does not find it, in contradiction to its validity.
By an inductive argument it follows, that for any configuration of failures on the links, $S$ and $S'$ choose the same sequence of trials on paths. □

## 4.2 The Construction of $\pi^*_M(G)$

The length of an optimal sequence of paths is at most $O(|E|)$, since each trial of a path adds at least one link to the set of tested links. We show below how to use Corollary 1 to find this sequence in polynomial time.

**Definition 7:** Given a parallel composition $G = (G_1 \parallel G_2)$ and any strategy $S$, we say that $G_1 \longrightarrow G_2$ under $S$, if for every configuration of failures on the links, $S$ tries a path in $G_2$ only after finding that every path in $G_1$ includes a faulty link.

**Definition 8:** Let $\pi_M^*(G)$ be an optimal order on paths for $G$ with the corresponding optimal complete order on links $\pi_L^*(G)$. We call the graph $SG$ a *suffix* of $G$, if it can be obtained from $G$ as follows:

1. For some prefix $\tilde{\pi}_M^*(G)$ of $\pi_M^*(G)$, represented by a prefix $\tilde{\pi}_L^*$ of $\pi_L^*(G)$, assign to each link its status in $cfc_{S^*}(G,\bar{p})$.
2. Apply the update procedure (as given in Section 2.2) to $G$ using the status (faulty/nonfaulty) assigned to the links in $\tilde{\pi}_L^*$.

**Example 2:** For the network of Example 1, an optimal complete order on links is

$$\pi_L^*(G) = ((s, v_1), e_1, e_2, (s, t)),$$

where $e_1 = (v_1 v_2, t)$ and $e_2 = (v_1 v_3, t)$. (Recall that we replace the links attached to vertices of degree 2 by a *single* link.)
A complete fault configuration for the network under the corresponding optimal strategy is $(1, 0, 0, 1)$. The subgraph $\tilde{G} = (V, \tilde{E})$ with $\tilde{E} = \{(s, v_1), e_1\}$ is a prefix of $G$. Taking $\tilde{\pi}_L^* = ((s, v_1), e_1)$ with the link $(s, v_1)$ nonfaulty, and the link $e_1$ faulty we obtain the graph $SG = (V', E')$, with $V' = \{s, t\}$ and $E' = \{(s, t), (s, t)\}$, which is a suffix of $G$.

The construction of an optimal trial tree for an *sp* graph involves handling *serial* and *parallel* compositions of subgraphs. For these two cases we give the following algorithm:
Algorithm $A_{sp}$:
Input: The optimal orders on paths $\pi_M^*(G_1), \pi_M^*(G_2)$ for $(G_1, \bar{p}_1), (G_2, \bar{p}_2)$ respectively.
Output: An optimal order on paths $\pi_M^*(G)$ for $(G, (\bar{p}_1, \bar{p}_2))$, where $G$ is a serial or parallel composition of $G_1$ and $G_2$.

Case 1: $G = (G_1 \Longrightarrow G_2)$
Let $\pi_M^*(G_1) = (m_1^1, \ldots, m_k^1), \quad \pi_M^*(G_2) = (m_1^2, \ldots, m_l^2)$ then

$$\pi_M^*(G) = (m_1^1 m_1^2, m_2^1 m_1^2, \ldots, m_k^1 m_1^2, m_k^1 m_2^2, \ldots, m_k^1 m_l^2)\ .$$

Case 2: $G = (G_1 \parallel G_2)$

1. Let $\pi_M^*(G) =$ the empty sequence; $i = 1, j = 2$.
2. Let $SG_i = G_i \backslash \pi_M^*(G)$ (i.e., contract and delete from $G_i$ the links in $\pi_M^*(G)$ according to their state in $cfc_S(G_i, \bar{p}_i)$); $SG_j = G_j \backslash \pi_M^*(G)$.
3. Let $\widetilde{SG}_i^*$ be the maximal prefix $\widetilde{SG}_i$ of $SG_i$ satisfying
$$E(T|\widetilde{SG}_i \longrightarrow \widetilde{SG}_j) \leq E(T|\widetilde{SG}_j \longrightarrow \widetilde{SG}_i) \tag{10}$$
for every prefix $\widetilde{SG}_j$ of $SG_j$.
4. Append $\widetilde{SG_i}^*$ to $\pi_M^*(G)$; Exchange $i$ and $j$.

5. Repeat steps 2-4 until both suffixes are empty graphs.

**Theorem 3:** $A_{sp}$ produces an optimal trial tree for any sp graph $G$.

**Proof:** For $G = (G_1 \Longrightarrow G_2)$, where $G_1$, $G_2$ are *sp*, the proof is immediate. We therefore assume that $G = (G_1 \parallel G_2)$. From Corollary 1, it is sufficient to consider all the prefixes of $\pi^*_{M1}$ and $\pi^*_{M2}$. By Definition 6, a prefix $\tilde{G}$ of a graph $G$ represents a sequence of paths in $G$. Let $\pi^*_M(G) = (\widetilde{SG}_{11}, \widetilde{SG}_{21}, \widetilde{SG}_{12}, \widetilde{SG}_{22}, \ldots)$ be the sequence of paths produced by the algorithm, where $\widetilde{SG}_{1k}, \widetilde{SG}_{2k}$ are prefixes of the suffixes of $G_1, G_2$ left from iteration $k-1$ of the algorithm, and we pass from $\widetilde{SG}_{1k}$ to $\widetilde{SG}_{2k}$ when all paths in $\widetilde{SG}_{1k}$ were found faulty, and from $\widetilde{SG}_{2k}$ to $\widetilde{SG}_{1(k+1)}$, when all paths of $\widetilde{SG}_{2k}$ were found faulty.

It remains to show, that if $\widetilde{SG}^*_{1k}$ is the maximal prefix satisfying (10) at iteration $k$, then it is the next term in $\pi^*_M(G)$ (and the same for $\widetilde{SG}^*_{2k}$ after the concatenation of $\widetilde{SG}^*_{1k}$ to $\pi^*_M(G)$).

By property $(iii)$ in Lemma 6, we look for a prefix $\widetilde{SG}_{1k}$ such that the order of trials $\widetilde{SG}_{1k} \longrightarrow SG_{2k}$ is optimal for $\widetilde{SG}_{1k} \parallel SG_{2k}$. Since the next term in $\pi^*_M(G)$ belongs to $G_1$, there exists such a prefix in $SG_{1k}$. In order to find the desired prefix, we check for every prefix $\widetilde{SG}_{1k}$ of $SG_{1k}$, if

$$E(T|\widetilde{SG}_{1k} \longrightarrow SG_{2k}) \leq E(T|\widetilde{SG}_{1k} \not\longrightarrow SG_{2k}). \tag{11}$$

The right hand side of inequality (11) is computed by using properties $(ii)$ and $(iii)$ of Lemma 6:

If inequality (11) is satisfied for a prefix $\widetilde{SG}_{1k}$ of $SG_{1k}$, then by property $(ii)$ of the lemma, for *every* prefix $\widetilde{SG}_{2k}$ of $SG_{2k}$, the order $\widetilde{SG}_{1k} \longrightarrow \widetilde{SG}_{2k}$ is optimal for $\widetilde{SG}_{1k} \parallel \widetilde{SG}_{2k}$, therefore, it is sufficient to check inequality (10). If there exists a prefix $\widetilde{SG}_{2k}$ for which inequality (10) is not satisfied with $\widetilde{SG}_{1k}$, then obviously inequality (11) is not satisfied for $\widetilde{SG}_{1k}$.
There may be several prefixes of $SG_{1k}$ satisfying (10), but since the next term in $\pi^*_M(G)$ is a prefix of $SG_{2k}$, we take the maximal among the $\widetilde{SG}_{1k}$'s, denoted above by $\widetilde{SG}^*_{1k}$. □

**Theorem 4:** The construction of an optimal trial tree for an *sp*-graph $G = (V, E)$, using the algorithm $A_{sp}$ requires $O(|E|^3)$ steps.

**Proof:** For a given sequence of serial/parallel compositions, we consider first the number of steps required in constructing an optimal trial tree for a single composition.

1. $G' = (G_1 \Longrightarrow G_2)$: $O(|E|)$ steps are required in the implementation of the algorithm, for the case $G$ is a serial composition.

2. $G' = (G_1 \parallel G_2)$: We have to consider the number of operations required in Step 3 of $A_{sp}$. Writing $\pi^*_M(G_1)$ and $\pi^*_M(G_2)$ as two sequences of paths, $(i_1, \ldots, i_k)$ and $(i_1, \ldots, i_l)$ respectively, we start with the complete sequences. Each time a prefix of a sequence is appended to $\pi^*_M(G)$, it is deleted from the appropriate sequence. Therefore, any prefix of $\pi^*_M(G_1)$ is compared at most twice with every prefix of $\pi^*_M(G_2)$, and the sequences are merged in $O(|E|^2)$ steps. Observe, that the evaluation of the two sides of inequality (10), in each comparison of two prefixes, does not affect the overall number of operations, since the desired expectations may be obtained iteratively in constant time, using the expectations calculated in previous stages.

As the number of compositions producing an *sp* graph $G = (V, E)$ is $O(|E|)$, the overall time complexity of the algorithm $A_{sp}$ is $O(|E|^3)$. □

# 5 The Complexity of Adaptive Source Routing in a General Network

In this section we state some complexity results concerning the minimization of the cost of delivering a single message in a general high-speed network. These results account for our difficulty in devising an optimal strategy for a wider class of topologies.

**Lemma 7:** ([3]) For a given network $(G, \bar{p})$, computing $E_{OPT}(G, \bar{p})$ is #P-Hard.

**Theorem 5:** The problem of devising a strategy, which minimizes the cost of delivery for a single message can be solved in polynomial space and is #P-Hard.

**Proof:** The above problem is in PSPACE, since it is a problem of decision-making in conditions of uncertainty (follows from Proposition 1 and Theorem 1 in [11]).
From Lemma 7 it follows, that devising an optimal strategy is #P-Hard (since the computational problem of finding the first path to be tried under an optimal strategy is #P-Hard). $\square$

We now claim that the minimal expected cost is even hard to approximate.
Our optimization problem is closely related to the $s$-$t$ reliability problem defined as follows:
Input: A network $(G, \bar{p})$, and $s, t \in V$.
Output: The probability that $G$ has a nonfaulty $s$-$t$ path.

Denote this probability by $p$, then given $\epsilon > 0$, $\tilde{p}$ is an *$\epsilon$-approximation to $p$*, if

$$(1-\epsilon)p \leq \tilde{p} \leq (1+\epsilon)p$$

**Theorem 6:** For any polynomial time algorithm A, and $\varepsilon > 0$, there exists a network $(G, \bar{p})$, such that

$$\frac{E_A(G, \bar{p})}{E_{OPT}(G, \bar{p})} > 1 + \varepsilon \ .$$

**Proof:** We use a reduction similar to that given in the proof of Lemma 7 ([3]): For a given network $(G, \bar{p})$, where $G = (V, E)$, $e = |E|$, such that the success probability of every link is $\frac{1}{2}$, let $s, t \in V$, and $p$ be the probability that there is a nonfaulty $s$-$t$ path in $G$. Obviously, $p = \frac{k}{2^e}$ for some integer $k \geq 1$. The reduction is to the network $(G', \bar{p}')$, with $V' = V$, $E' = E \cup \{(s,t)\}$. The success probability on $(s,t)$ is $\frac{1}{2^e}$. Then an optimal strategy tries the link $(s,t)$ last in any series of trials on paths.
Thus,

$$E_{OPT}(G', \bar{p}') = E_{OPT}(G, \bar{p}) \cdot p + (E_{OPT}(G, \bar{p}) + 1) \cdot (1-p),$$

or

$$p = E_{OPT}(G, \bar{p}) + 1 - E_{OPT}(G', \bar{p}').$$

Let $\alpha = \frac{\epsilon}{(2e+1)2^e}$, and $B_\alpha$ be a polynomial time algorithm, that provides $\alpha$-approximations $B_\alpha(G)$ and $B_\alpha(G')$ for $E_{OPT}(G, \bar{p})$, $E_{OPT}(G', \bar{p}')$ respectively. Let $\tilde{p}$ denote an approximation to $p$.
Then,

$$\begin{aligned} \tilde{p} &= B_\alpha(G) + 1 - B_\alpha(G') \\ &\geq (1-\alpha)E_{OPT}(G, \bar{p}) + 1 - (1+\alpha)E_{OPT}(G', \bar{p}') \\ &= p - \alpha(E_{OPT}(G', \bar{p}') + E_{OPT}(G', \bar{p}')) \geq p - \alpha(2e+1) \ . \end{aligned}$$

Hence, the existence of a polynomial time $\alpha$-approximations to $E_{OPT}(G,\bar{p})$,$E_{OPT}(G',\bar{p}')$ implies the existence of an $\epsilon$-approximation for $p$. The statement of the theorem follows from the fact that it is hard to approximate $p$ [13]. $\square$

# 6 Conclusions and Open Problems

We examined a variant of the CTP [12], which applies to high speed networks, where failures on links may occur between transmissions of distinct messages, and the expected transmission time of a single message depends on the number of paths tried until its successful delivery.

We have shown, that the problem of devising an optimal routing strategy, based on the failure probabilities of the links and the accumulated knowledge on faulty links is #P-Hard.
A polynomial time algorithm for constructing an optimal routing strategy was presented for the subclass of serial-parallel networks.

**Open Problems:**

- Is there a subclass of topologies for which the greedy algorithm is a constant approximation to an optimal strategy?
- Are there polynomial time algorithms for devising efficient routing schemes for other subclasses of topologies (such as layered graphs and planar networks)?
- Can we achieve a bounded ratio to the minimal expected cost for a general network using efficient decomposition to serial/parallel components?
- How do the above results extend to resolving the problem of adaptive routing in an arbitrary (unreliable) network, where the network properties call for *source routing* , and distances are incorporated into the cost function?

# Acknowledgments

We would like to thank Shai Ben-David for the contribution of Example 1 and for many illuminating comments throughout the early stages of this work. We also thank Amotz Bar-Noy, Amir Herzberg, Ilan Kessler and Shay Kutten, for some helpful discussions, and for communicating to us Lemma 7, that motivated the discussion of general network topologies. The presentation of the paper benefited much from the careful reading and detailed comments of two anonymous referees.

# References

[1] V. Ahuja, "Design and Analysis of Computer Networks", *McGraw-Hill Book Company, New York*, 1982.

[2] B. Awerbuch, I. Cidon, I. Gopal, M. Kaplan and S. Kutten "Distributed Control for PARIS", proceedings of In *Proc. of the 9th ACM Symposium on Principles of Distributed Computing*, pp. 145-160, Quebec City, Canada, August 22-24, 1990.

[3] A. Bar-Noy, A. Herzberg, I. Kessler and S. Kutten, Private Communication.

[4] A. Bar-Noy, B. Schieber, "The Canadian Traveler Problem", In *Proc. of the 2nd ACM-SIAM Symposium on Discrete Algorithms*, pp. 261-270, San Francisco, CA, January 28-30, 1991.

[5] I. Cidon, and I. Gopal, "Paris – an Approach to Integrated Private Networks", *International J. of Digital and Analogue Cabled Systems*, 1984.

[6] D. Clark et al., "Aurora Giga-bit Test Bed" Comp. Networks in ISDN, 1991.

[7] R.J. Duffin, "Topology of Series Parallel Network", *J. Math. Anal. Appl.* **10**, pp. 303-318, 1965.

[8] A. Itai, H. Shachnai, "Optimal Source Routing in a Network with Dependent Paths", TR #690, Technion IIT, CS Dept., October 1991.

[9] Y. Mansour, B. Patt-Shamir, "Greedy Packet Scheduling on Shortest Paths", In *Proc. of the 10th ACM Symposium on Principles of Distributed Computing*, pp. 165-175, Montreal, Canada, August 19-21, 1991.

[10] M. S. Meketon, D. M. Topkis, "Adaptive Routing for Damaged Networks", *Proc. IEEE Mil. Comm. Conf.*, pp. 282-288, 1983.

[11] C. H. Papadimitriou, "Games Against Nature". *J. of Computer and System Sci.* **31**, pp. 288-301, 1985.

[12] C. H. Papadimitriou, M. Yannakakis, "Shortest Paths without a Map". *Theoretical Comp. Sci.* **84**, pp. 127-150, 1991.

[13] A. Rosenthal, "Computing the Reliability of Complex Networks". *SIAM J. Appl. Math.* **32**, #2, pp. 384-393, 1977.

[14] M. Schwartz, T. E. Stern, "Routing Techniques Used in Computer Communication Networks", *IEEE Transactions on Communication*, COM-28, pp. 539-552, 1980.

[15] D. M. Topkis, "Reordering Heuristics for Routing in Communication Network", *J. of Applied Prob.*, pp. 130-143, 1986.

[16] J. Valdes, R.E. Tarjan and E.L. Lawler, "The Recognition of Series Parallel Digraphs", In *Proc. of the 11th Annual ACM Symposium on Theory of Computing*, 1979.

## Appendix

**Proof of Lemma 1:** The network topology implies, that any link emanating from $s$ belongs to a *single* $s$-$t$ path. The proof is by interchange argument:
Since any path converging network is an sp graph, we use Theorem 1 to identify an optimal strategy with a unique order of trials on the network links. Assume that $\pi^*_M(G) = (m_1, \ldots, m_n)$ is a complete order on paths for an optimal strategy $S^*$ on $G$, then $\pi^*_M(G)$ contains every $s$-$t$ path in $G$. Let $m_i, m_{i+1}$ be two consecutive elements in $\pi^*_M(G)$, with the success probabilities $p_{m_i}, p_{m_{i+1}}$ respectively, then it is sufficient to show, that

$$p_{m_i} < p_{m_{i+1}} \Longrightarrow \quad S^* \text{ is not optimal.}$$

Denote by $\text{Prob}_{S^*}(T > r)$ the probability, that for a given network, the number of trials of paths under the strategy $S^*$ is higher than $r$, then

$$E_{S^*}(G, \bar{p}) = \sum_{r=0}^{n-1} \text{Prob}_{S^*}(T > r)\,. \tag{12}$$

We denote by $S'$ the strategy, which switches $m_i$ and $m_{i+1}$, i.e.

$$\pi_M = (m_1, \ldots, m_{i+1}, m_i, \ldots, m_n)\,. \tag{13}$$

We handle the following two cases separately:

1. $m_i$ and $m_{i+1}$ are disjoint, then it follows from [15], that

$$E_{S^*}(G, \bar{p}) - E_{S'}(G, \bar{p}) > 0 \;\;. \tag{14}$$

2. If $m_i$ and $m_{i+1}$ are dependent, then let

$$m_i = (e_{i1}, \ldots, e_{ik}) \qquad m_{i+1} = (e_{(i+1)1}, \ldots, e_{(i+1)l}) \quad .$$

   For some $1 \leq j \leq min(l, j)$, the last $j$ links in $m_i, m_{i+1}$ are identical.
   We use the notation $p_{ir}, p_{(i+1)r}$ for the success probabilities on $e_{ir} \;\; 1 \leq r \leq k \;\; , e_{(i+1)r} \;\; 1 \leq r \leq j$ respectively, then

$$p_{i(k-j)} = p_{(i+1)(l-j)}, \;\; \ldots, \;\; p_{ik} = p_{(i+1)l},$$

   and since failure probabilities on the links are independent,

$$p_{m_{i+1}} > p_{m_i} \Longrightarrow \prod_{r=1}^{l-j-1} p_{(i+1)r} > \prod_{r=1}^{k-j-1} p_{ir} \;\;,$$

   thus, conditioning on the location of the first fault on each of the paths under the strategies $S^*$ and $S'$, we find

$$E_{S^*}(G, \bar{p}) - E_{S'}(G, \bar{p}) > 0 \;\;.$$

$\square$