Local Restoration for Trees and Arborescences under the Spanning Tree Protocol

Gaia Nicosia* Gianpaolo Oriolo⋆ Laura Sanità° Ezio Sperduto*

*Dipartimento di Informatica e Automazione, Università degli studi Roma Tre
Via della Vasca Navale 79, I 00146 Roma, Italia.

⋆Dipartimento di Ingegneria dell'Impresa, Università degli studi di Roma "Tor Vergata"
Via del Politecnico 1, I 00133 Roma, Italia.

°Institute of Mathematics, Ecole Polytechnique Federale de Lausanne
1015 Lausanne, Switzerland.

Abstract

Protocols belonging to the Spanning Tree Protocol (STP) family, used for instance in Ethernet networks, route traffic demands on tree topologies that are evaluated through shortest path procedures. In this paper we deal with the problem of assigning costs to the arcs of a network in order to guarantee that the tree that is defined by the SPT protocol allows for a restoration of traffic demands in failure situations such that the following property holds: traffic demands that are not affected by the failure are not redirected. The latter property, called minimum service disruption property, is a goal pursued by network operators in order to minimize transmission delays due to the restoring process, and therefore guarantee adequate levels of quality of service (QoS).

We prove that for each spanning tree T of an undirected connected graph, there is a suitable cost vector such that the Spanning Tree Protocol routes the traffic on T and, when there are links and/or nodes failures, does not re-route the traffic demands that are not affected by the failure. We then generalize our results to directed graphs.

Keywords: *Spanning Tree Protocol, Ethernet networks, Restoration.*

Introduction

A crucial issue for network design is the survivability, that is the capability of a network to remain operational even when some components, like links, fail. Survivable communication networks guarantee therefore that traffic demands can be routed even in failure situations, by *restoring* the traffic that is affected by the failure. Since restoring is often expensive and may cause long delays in transmissions [2, 3], in order to meet QoS constraints a key property of the restoring strategy is the so-called *minimum service disruption property* [4] or *strong resilience* [1], requiring that traffic which is not affected by the failure is *not* redirected in the process of restoring the affected traffic.

In this paper we investigate how to implement the minimum service disruption property in networks where the traffic is routed by some protocol on trees. Namely, we consider the following setting, typical for networks, such as Ethernet based networks, where the traffic is routed according to the Spanning Tree Protocol [5].

We are given an undirected connected graph $G(V, E)$ and we assume that there are traffic demands between each pair of vertices. The Spanning Tree Protocol routes the traffic between each pair of vertices

u, v on the unique $u - v$ path of a spanning tree: the spanning tree is a shortest path tree rooted at some suitable node $r \in V$ with respect to a cost vector $w : E \mapsto \mathcal{Z}_+$. The protocol has some simple rules that break ties, so we may assume that this tree is unique and we denote it by $T_G(w, r)$. Note that the cost vector w is usually defined by traffic engineers and, for our purposes, once it is defined, it does not change, even when some failures occur.

Now suppose that some set of edges $F \subset E$ fail. We deal with edge failures, since the failure of a vertex can be reduced to the failure of all the edges incident to it. For sake of simplicity, we first assume that $G \setminus F$ is still connected.

The STP protocol will route the traffic on $T_{G\setminus F}(w, r)$, i.e. a shortest path tree of $G\setminus F$, with root r and with respect to w (we may assume that $w(e)$ switched to ∞, if $e \in F$). Trivially, $T_{G\setminus F}(w, r) = T_G(w, r)$ if and only if no edge of $T_G(w, r)$ is failed. Suppose therefore that some edges of $T_G(w, r)$ belong to F and so fail. Note that, unless *all* the edges of $T_G(w, r)$ fail, for some pair of vertices $\{u, v\}$, the $u - v$ path on $T_G(w, r)$ will not be affected by the failure: the minimum service disruption property amounts to requiring that, for any such pair, the $u - v$ paths on $T_G(w, r)$ and $T_{G\setminus F}(w, r)$ are indeed the same. Note that this holds if and only if each edge $\{x, y\} \in E(T) \setminus F$ belongs to $T_{G\setminus F}(w, r)$. Our discussion leads to the following definition:

Definition 1 *Let $G(V, E)$ be an undirected connected graph and $q \leq\mid E \mid$ an integer. A cost vector $w : E \mapsto \mathcal{Z}_+$ is q-consistent if there exists a spanning tree $T(w)$ such that:*

(i) For each $r \in V$, $T(w)$ is the unique shortest path tree of G rooted at r (i.e. $T(w) = T_G(w, r)$).

(ii) For each $F \subseteq E$, with $\mid F \mid\leq q$, and each edge $\{x, y\} \in E(T(w)) \setminus F$, $\{x, y\}$ belongs to every shortest path tree of $G \setminus F$ that is rooted at a vertex in the same component of u and v.

In this case, we say that w defines $T(w)$.

Note that this definition allows for properties that are slightly stronger than those discussed so far. In particular, $T(w)$ is the unique shortest path tree of G, with respect to *any* root $r \in V$ (and therefore we do not need tie-breaking rules for defining the root). Moreover, $G \setminus F$ needs *not* to be connected: in this case, we require that, for each pair of vertices $\{u, v\}$ whose path P on $T(w)$ is not affected by the failure of F, every shortest path of $G \setminus F$ that is rooted at a vertex in the same component of u and v, includes P. Note in particular that the failure of a node, i.e. the failures of all the edges incident to it, will certainly cause the lost of the connectivity.

Our main result shows that, for each spanning tree T of an undirected graph $G(V, E)$ and an integer $q \leq\mid E \mid$, one may build in linear time a q-consistent cost vector $w : E \mapsto \mathcal{Z}_+$ such that $T = T(w)$. In other words, for each spanning tree of an undirected graph, there is a suitable cost vector w such that the Spanning Tree Protocol routes the traffic on T and, when there are failures, routes the traffic demands according to the minimum service disruption property.

Our results can be generalized to directed graphs. In this case we assume that there are traffic demands only from some common source node r to every other node.

References

[1] G. Brightwell, G. Oriolo and F,B. Sheperd, *Reserving Resilient Capacity in a Network*, Siam Journal on Discrete Mathematics, **14** (2001) pp. 524–539.

[2] P. Flocchini, L. Pagli, G. Prencipe, N. Santoro, P. Widmayer, *Computing all the best swap edges distributively*, Journal of Parallel and Distributed Computing, **68** (2008) pp. 976–983.

[3] E. Nardelli, G. Proietti, P. Widmayer, *Swapping a Failing Edge of a Single Source Shortest Paths Tree Is Good and Fast*, Algorithmica, **35** (2003) pp. 56–74.

[4] A. F. de Sousa, G. Soares, *Improving Load Balance and Minimizing Service Disruption on Ethernet Networks using IEEE 802.1S MSTP*, Proc EuroFGI Workshop on IP QoS and Traffic Control, Lisbon, Portugal, **1** (2007) pp. 25–35.

[5] Standard IEEE 802.1D - Spanning Tree protocol.