Scheduled Hot-Potato Routing

Joseph (Seffi) Naor* Ariel Orda† Raphael Rom‡

Abstract

This paper is concerned with fast, hot-potato routing, performed according to a predetermined schedule. At each time period each node selects an outgoing link, through which an incoming packet is sent. No buffers are used. We investigate first the problem of how to route a network-wide demand of packets, given the predetermined schedule. We show that certain versions of the problem have efficient solutions, while other versions are intractable. We then consider the problem of finding an optimal schedule given a network-wide demand of packets. We indicate that the problem is tractable for either a single source or single destination. However, for the multi-source multi-destination case we show that it is an NP-complete problem. The problem remains intractable even for a simple topology of nodes arranged on a bi-directional line. We present an efficient heuristic for directed tree-networks, and adapt it to general topologies through a recursive scheme, for which an efficient performance bound is shown.

1 Introduction

Due to the development of high-speed networks, there has been recently growing interest in efficient routing techniques that would allow a fast decision at the node regarding the edge through which an incoming packet should be routed. An efficient routing technique for a high-speed network is one that demands little "thinking" from the node. Ideally, a node would take an immediate routing decision (e.g., [1]). In addition, an efficient scheme for high-speed routing would avoid the use of buffering at intermediate nodes: ideally, a packet would flow throughout the network without queuing at nodes (e.g., [2]). Indeed, queuing means more delay and also a more complicated strategy from the node's standpoint.

A possible solution that addresses the above issues is that of scheduled hot-potato routing. "Scheduled" means that each node routes incoming packets to outgoing links according to a predetermined schedule, i.e.: the identity of the outgoing link is a function solely of time. This means that the action taken by the node is immediate, and in particular it does not need even to look at the packet's header in order to perform the routing. Moreover, the routing decision becomes ever more simple if we let just one outgoing link to be "active" at every time instant. "Hot-Potato" routing means that packets are not queued enroute. Thus, with scheduled hot-potato routing we let packets move smoothly and fastly among nodes, through high-speed links, without encountering intra-nodal queuing, nor switching delays. Nonetheless, such a strategy demands careful planning, both of the predetermined schedule and of its actual use once it has been set.

Hot-potato (also called Deflection) routing has been given attention in recent years, e.g., [3, 4, 5, 6, 7, 8, 9, 10], mainly in the area of multiprocessing. Typically, however, in the above papers, hot-potato routing does consider the identity (label) of the packet for making its routing decision. Scheduled routing also started to gain attention, e.g. [11, 12, 13]. However, these works considered the use of store-and-forward buffering.

In this work we consider the combination of both principles, namely scheduling and hot-potato, in order to achieve an efficient method for routing in high-speed networks. Modern high-speed networks are oriented towards the establishment of long duration sessions, and are not datagram based. In such networks, a session is set up in advance, determining both the routes and the rate in which data will subsequently flow. In particular, ATM networks [14] are based on the establishment of virtual paths and virtual circuits whose lifetime is relatively long. Because of the timing balance involved, there is typically enough time to gather the data that is necessary to make the routing decision. Similar ideas to those presented in this paper have been applied to practical pilot networks, e.g., Isochronets [2], in which network bandwidth is

*Computer Science Department, Technion, Haifa 32000, Israel.

†Department of Electrical Engineering, Technion, Haifa 32000, Israel.

‡Department of Electrical Engineering, Technion , Haifa 32000, Israel, and Sun Microsystems, Mountain View, CA.

time-divided among routing trees in order to allocate "green-band" paths for traffic destined to a common root. However, until now, these ideas were tested only at the experimental level, while the present work is the first to provide formal analysis.

Two typical problems related to routing through pre-established connections are the *design* problem and the *usage* problem. The first is concerned with the efficient design of the long-term routing plan, once some data on the source-destination demands is collected. The second is concerned with the best use of an existing routing plan, for traffic patterns that were not necessarily accounted for at the time that the routing plan was devised, e.g., low-priority traffic classes. We note that the usage problem follows the philosophy of the best-effort traffic class in ATM networks.

The present work is an attempt at providing an analytical framework for this novel method of routing. We dwell on the problems that should be addressed, provide optimal solutions to some while proving the intractability of others, for which efficient heuristics are described and analyzed. It is interesting to note that our framework is somewhat related to satellite-switched TDMA [15]. (Note, though, that we deal with a single "transponder" and a multi-hop topology, whereas the classical satellite switching problem deals with multiple transponders and a single-hop topology.)

Specifically, we address the following problems. First, we consider a network for which a schedule has been predetermined, and investigate the usage problem, i.e., how to route a given demand of packets (for various source-destination pairs), through a pre-established schedule, in the most efficient way. We then consider the design problem, i.e., planning the best schedule for an expected demand of packets. We give an optimal solution for the first problem. We show that the second problem is intractable (even for very simple topologies). We indicate its tractability in some special cases and present efficient heuristics for the general case.

2 Model

The network is a directed graph $G(V, E)$, where V is the set of nodes and $E \subseteq V \times V$ is the set of edges, or links. For a node $i \in V$, let IN_i be the set of neighbors k such that $(k, i) \in E$, and OUT_i be the set of neighbors k such that $(i, k) \in E$.

We consider a discrete domain of time and a synchronous mode of operation. It is assumed that transmission delay on all links is of unit time, and refer to a *packet* as the amount of data that can be transmitted in a unit of time. Thus, at most one packet can be carried on a link at any time. Furthermore, it is assumed that nodes do not have buffers, and thus, an incoming packet has to be switched immediately to an outgoing link.

Considering a time period of P, we define the *nodal schedule* of a node i, S^i, to be a sequence $\{k^i_0, k^i_1, \cdots k^i_{P-1}\}$ such that, for $0 \leq \tau \leq P-1$, $k^i_\tau \in OUT_i$, i.e., the nodal schedule of i assigns to each time-slot within $[0, P-1]$ a single neighbor of node i (routing is discussed next). Note that P is a given quantity independent of any other network parameter (notably of $|V|$). Typically, one can assume that P is small compared to $|V|$. Given a nodal schedule for each node, we define the *network schedule* S to be the set comprising all nodal schedules, i.e., $S = \{S^i | i \in V\}$.

We now discuss the *scheduled hot-potato routing* for a given network schedule S. Consider a node $i \in V$ and a time-slot t[1]. If exactly one packet arrives at i at the beginning of that time-slot, then it is sent (immediately) through link (i, k), where $k = k^i_\tau \in S^i$, and $\tau = t \bmod P$. If more than one packet arrives at the same time, then an arbitrarily chosen (single) packet is routed according to the schedule, while the rest are *eliminated*.

In communication networks it is common to specify the traffic requirements as the average amount of flow (bits/sec) between source-destination pairs. In particular, modern networks use the notion of connections (or sessions) and associate flow requirements with these connections. Given a flow requirement between a certain pair of nodes, a period P and a slot size, this requirement can be expressed in units of slots/period. Thus, a flow requirement for a connection can be expressed as a number of packets per period P (note that a connection may have more than a single slot assigned to it in a period). With this in mind, the traffic requirements are formally defined as follows.

A *demand* D of size K on the network is a sequence of (not necessarily different) K source-destination pairs $D = \{(s_1, d_1), (s_2, d_2), \cdots (s_K, d_K)\}$; each pair belongs to a certain connection and designates one packet transmission of that connection. At times, we denote by p_j the packet associated with the j-th pair of D. In other words, $D = \{p_1, p_2, \cdots p_K\}$, where p_j is a packet to be routed from node s_j to node d_j. If, for all j, $s_j \equiv s$ then we say that D is a *single source*

[1] All times are global relative to some starting time denoted by $t = 0$.

demand; similarly, if for all j $d_j \equiv d$ then we say that D is a *single destination* demand; if both hold, we say that D is a *single-source/single-destination* demand. For ease of presentation we shall assume that each node may be a source of at most one packet, i.e., $s_i \neq s_j$ for $i \neq j$ (this implies $K \leq |V|$). Our results are easily adapted to handle the more general case, by splitting a source s with K_s packets into a sequence of K_s independent nodes each connected to s. Denote by $M_{G,D}$ the diameter of graph G with respect to D, i.e., the maximal minimum-hop distance among all source-destination pairs of D.

Given a network demand D of size K, a *departure plan* Δ is a sequence of (integer) times $\Delta = \{\delta_1, \delta_2, \cdots \delta_K\}$ such that δ_j is the departure time of packet p_j i.e., p_j leaves the source node s_j (for the first time) at the beginning of time-slot δ_j.

A packet is routed between nodes until it arrives at its destination or it is eliminated. For given S and D, we say that a departure plan Δ is *proper* if no elimination occurs. For S, D and a proper Δ, the *arrival time* α_j of a packet p_j is the time at which it reaches the destination; if the packet does not reach its destination (i.e., it moves endlessly among nodes), then we set $\alpha_j = \infty$. For given S and D, the *termination time* of a proper departure plan Δ is the highest value of arrival times (possibly infinity).

We note that, even if p_j is the only packet in the network, for a given network schedule S, there may still be no departure time for p_j that would bring it to its destination. Namely, for any departure time the packet will loop in the network forever. We say that S is *feasible* for packet p_j if there is some finite time t, such that if p_j departs at t, then it arrives at its destination in finite time; t is said to be a feasible departure time for p_j, given S. We say that a schedule S is feasible for demand D if it is feasible for each packet of D.

From the previous description it is not clear how packets are ever removed from the network. A useful way to view this is to assume that each node has a link emanating from it that terminates with the end users that are attached to that node. In other words, for a given node i, the set OUT_i contains a link to the attached end users. Thus, if a packet arrives at node i at time τ, and if schedule S^i indicates that k^i_τ is the edge that points to the end user, then we say that the packet is delivered to its destination. A variation on this approach is the one described in [1], where, by default, each packet, while being switched according to the nodal schedule, is also delivered to the end user attached to every node along the path, and is discarded by those nodes for which it is not intended (as is done in most LANs)

3 Use of a Given Schedule

In this section we consider the situation in which the network and its schedule are given and we look for ways to route packets through the network under various timing circumstances.

The first question is whether a proper departure plan of finite termination time exists at all. The following lemma gives a necessary and sufficient condition.[2]

Lemma 3.1 *Given a network G, a schedule S and a demand D, there is a proper departure plan Δ with finite termination time if and only if S is feasible for D.*

Lemma 3.2 *Given a network G, a schedule S with period P and a demand D, S is feasible for D iff there is a proper departure plan Δ whose termination time is at most $T_{\max} = K \cdot (|V| - 1) \cdot P$.*

The above corollary enables us to consider a finite domain of time, whose size $T_{\max}$ is polynomial in $|V|$, K, and P.

It is clear that if packets are allowed to enter the network freely upon arrival, some of them may be eliminated. In an attempt to maximize throughput, it would be useful to exert some admission control, i.e., forbid the transmission of some of the packets so that throughput is enhanced. More formally, consider the following problem:

Throughput Maximization (TM) Problem: Given a network G, a schedule S, a demand D, and a departure plan Δ, find a maximal set of packets that can be delivered to their respective destinations.

To attack this problem, we observe that the dynamic behavior of the network G with schedule S over a finite domain of time T can be represented by an equivalent static network, called the *Timed-Exploded (TE)* version of G. Note that T is implicitly determined by the other parameters of the TM problem and is bounded by the value $T_{\max}$ as computed in Lemma 3.2. Specifically, a time-exploded version $TE_{G,S,T}$ of G and S, and of size $T \leq T_{\max}$, is obtained in the following way.

The graph $TE_{G,S,T}$ consists of $T + 1$ "layers" $0, \cdots, T$, where each contains $|V|$ nodes, such that

[2] Proofs omitted from the conference version appear in the full version, which can be obtained from the authors.

there is a $1-1$ correspondence between the nodes of each layer and nodes in V; for each $v \in V$, and $0 \leq t \leq T$, we denote by v^t the node at level t that corresponds to node v. v^t is said to be a *descendant* of v. The set of edges of $TE_{G,S,T}$ is defined as follows. First, for each $v \in V$, and $1 \leq t < T$, we insert an edge from v^t to w^{t+1}, where w is the next node from v at time t according to schedule S, i.e.: $w = k_t^v$. Then, for each packet $p_j = (s_j, d_j) \in D$ we insert an edge from s_j^0 to $s_j^{t_j}$, where t_j is the departure time of p_j according to the departure plan Δ. If $t_j = 0$, then we insert an edge from s_j^0 to w^1, where w^1 is the next node from s_j at time 0 according to schedule S.

Consider a path in $TE_{G,S,T}$ that starts at the 0th-level descendant of the source node in G of some packet $p_j \in D$, and ends at a descendant of the destination node of p_j. Such a path in $TE_{G,S,T}$ corresponds to the unique route in G that conforms to the schedule S, departure plan Δ, and upper bound T on the arrival time, and is referred to as a *relevant path*. The graph $TE_{G,S,T}$ has the additional following properties, that can easily be verified:

- $TE_{G,S,T}$ is the union of at most $|V|$ directed relevant paths.
- Two relevant paths in $TE_{G,S,T}$ that meet at some node v^t, continue from v^t together until one of them (or both) reaches its destination. Notice that two such paths correspond to the routes of two packets that collide at node v at time t (causing one of them to be eliminated).

Given these properties, solving problem TM amounts to answering the following question: what is the maximum number of disjoint relevant paths in $TE_{G,S,T}$?

In order to answer this question, we construct the corresponding (undirected) *Path Graph* of $TE_{G,S,T}$ ($PG(TE_{G,S,T})$) as follows. In $PG(TE_{G,S,T})$, there corresponds a node to every relevant path in $TE_{G,S,T}$; there is an edge between every two nodes that correspond to two (relevant) paths in $TE_{G,S,T}$ that are not edge-disjoint.

Clearly, finding the maximum number of relevant edge-disjoint paths in $TE_{G,S,T}$ is equivalent to finding a maximum independent set in $PG(TE_{G,S,T})$, i.e., a maximum set of nodes such that any two nodes in the set are not adjacent. In general, both the problem of finding edge disjoint paths (between various source-destination pairs), and that of finding a maximal independent set, are known to be NP-Complete [16]. Nonetheless, we now show that in our case the problems are tractable, since the graph $PG(TE_{G,S,T})$ is chordal. A graph is called *chordal* if every cycle of length 4 or more contains a chord. Chordal graphs constitute a graph family for which it is well known [17] that many NP-complete problems become tractable when restricted to it. In particular, a maximum independent set can be computed in linear time in a chordal graph [17].

Lemma 3.3 *$PG(TE_{G,S,T})$ is a chordal graph.*

Theorem 3.1 *Problem TM can be solved by an $O(|V| \cdot T)$ algorithm.*

Proof: Since $TE_{G,S,T}$ is the union of at most $|V|$ directed relevant paths, the graph $PG(TE_{G,S,T})$ contains at most $|V|$ vertices. Every relevant path can intersect at most T other relevant paths (based on time of departure), or at most $|V|$ relevant paths (based on the number of vertices). In other words, every relevant path can intersect at most $\min\{|V|, T\}$ other relevant paths, yielding that the number of edges in $PG(TE_{G,S,T})$ is at most $(\min\{|V|^2, |V|T\}$. Therefore, $PG(TE_{G,S,T})$ can be constructed in $O(\min\{|V|^2, |V|T\})$ time. A maximum independent set in $PG(TE_{G,S,T})$ can be computed in linear time, i.e., in $O(\min\{|V|^2, |V|T\})$ time, yielding the theorem. □

Since $T \leq T_{\max} = O(K \cdot |V| \cdot P)$ and $K = O(|V|)$, we conclude that the above solution is polynomial in $|V|$ and P.

A natural question that arises in this context is that of finding a better use of a given schedule, i.e.: assuming all packets in D are available at time $t = 0$, when should each depart. More formally:

Schedule Usage (SU) Problem: Given a network G, a network schedule S, and a demand D, find a proper departure plan Δ of minimum termination time.

Unfortunately, this problem is intractable. To demonstrate this we construct the TE graph a little differently. In addition to the edges already defined for $TE_{G,S,T}$, for each $v \in V$ we insert an edge between v^0 and each v^t, $0 < t \leq T$. From this graph TE we construct a path graph as before. We note that there are at most $K \cdot T = O(|V| \cdot T)$ relevant paths in $TE_{G,S,T}$, and thus the construction and size of $PG(TE_{G,S,T})$ remain tractable. However, it can easily be verified that this path graph is not chordal anymore.

The intractability of Problem SU can be proved by a reduction from 3-SAT. We omit the details, and outline the proof very briefly, since the reduction needed is similar to the one used in Theorem 4.1. We have previously shown that a solution to the SU problem corresponds to a maximum independent set in the graph

$PG(TE_{G,S,T})$. Let H denote the graph generated in the reduction of 3-SAT to the Vertex Cover Problem in [16, pp. 54-55]. The idea of the NP-completeness proof of Problem SU is to generate a schedule for the vertices in H and a demand function such that a minimum departure plan corresponds to a minimum vertex cover in H, which in turn corresponds to a satisfying assignment to a 3-SAT formula.

Given that Problem SU is intractable we offer the following heuristic algorithm (Algorithm HSU) to find a good departure plan, which is based on a greedy approach. The main idea is to start at time $t = 0$, and schedule the maximal number of departures for that time. We then attempt to schedule the maximal number of departures for each successive time slot, using a residual time exploded graph, to avoid elimination by previously scheduled packets.

Algorithm HSU:

1. Initialization: $D' = D$, $TE' = TE_{G,S,T_{\max}}$, $t = 0$.
2. Solve problem TM for D', TE' and a departure plan Δ' such that $t_j = t$ for all packets $p_j \in D'$. Let $\widehat{D}$ be the resulting maximal set.
3. Set $\delta_j = t$ for all packets $p_j \in \widehat{D}$.
4. Remove from TE' all the edges of paths that correspond to the packets of $\widehat{D}$.
5. Set $D' = D' \backslash \widehat{D}$; $t = t + 1$.
6. If $D' \neq \emptyset$ then repeat Step 2, otherwise Stop: the departure time is $\Delta = \{\delta_1, \ldots, \delta_K\}$.

Algorithm HSU terminates in at most $T_{\max}$ iterations, since at each iteration, at least one path is selected and scheduled individually (See Lemma 3.2). Note also that algorithm HSU is not restricted to a single packet per source; it will also work if D includes a separate element for each (duplicate) packet.

4 Designing an Optimal Schedule

In this section we consider the complementary problem to that considered in the previous section, namely, how to design a schedule rather than use a given one.

Given demand D and schedule S, call the termination time, corresponding to the solution of Problem SU, the *optimal termination time* of S with respect to D. We refer to the *optimal termination time* for D as the minimum, among all optimal termination times, taken over all possible schedules. A schedule S that induces an optimal termination time (with respect to D) is called an *optimal schedule*. Note that the optimal termination time is bounded by Lemma 3.2.
We are now ready to present our problem formally:

Schedule Design (SD) Problem: Given a network G and a demand D, find an optimal schedule with respect to D.

In other words, our aim is to find a schedule, with respect to demand D, that achieves the best possible termination time (through the solution of problem SU). Due to the finiteness of the optimal termination time, (Lemma 3.2), problem SD can be transformed into the following decision problem:

T-constrained Schedule Design (TSD) Problem: Given a network G and a demand D, does there exist a schedule S for which there is a departure time Δ with termination time at most T?
We note that Problem SD can be optimally solved via Problem TSD, by performing a binary search over the feasible range of T.

4.1 Intractability

Unfortunately, problem TSD (and thus also problem SD) is intractable, as shown by the following theorem.

Theorem 4.1 *Problem TSD is NP-Complete.*

It turns out that problem SD is intractable even for some very simple special cases. Consider a network topology in which nodes are arranged on a line, such that each two adjacent nodes are connected by a bidirectional edge. Let D be such that traffic flows in both ways, i.e., some destinations are located to the right of their sources, while others are located to the left. It can be shown [18] that even in this setting problem SD is intractable.

We note, though, that a simple but efficient heuristic exists for the linear topology: first, all sources that have destinations to their right send (at time 0) their packets. Once these packets arrive, the other sources send (all at once) their packets. It is clear that by using this heuristic packets never collide and the termination time is at most twice that of the optimal termination time.

4.2 Tractable Special Cases

Consider the single destination case, i.e., where all packets have the same destination. Then, problem TSD can be solved in polynomial time in the following way. We construct $TE_{G,T}$, a time-exploded graph

of G and time T (that does not correspond to any schedule S), similar to the construction of $TE_{G,S,T}$ as follows. As in graph $TE_{G,S,T}$, $TE_{G,T}$ consists of $T+1$ "layers" $0, \cdots, T$, where each contains $|V|$ nodes, such that there is a $1:1$ correspondence between the nodes of each layer and nodes in V. The edge set of $TE_{G,T}$ is defined differently. Let v, w be vertices in G; there is an edge between every two nodes $v_t, w_{t+1} \in TE_{G,T}$, for $0 \leq t < T-1$, iff $(v,w) \in E$.

Thus, for every path in $TE_{G,T}$, a schedule S can be chosen, such that this path can be realized as a relevant path. Notice, however, that there exists a schedule that realizes two given paths in $TE_{G,T}$ (as relevant paths), if and only if the two paths are *disjoint*. This implies that the problem of satisfying a given demand D is equivalent to that of finding a set of disjoint paths, from the sources to the destinations, in the graph $TE_{G,S,T}$. This problem can be cast as a $\{0,1\}$ multicommodity flow problem in $TE_{G,T}$, where edges and vertices of $TE_{G,T}$ are assumed to have unit capacity.

The general discrete multicommodity flow problem is intractable, even in its $\{0,1\}$ version. However, in the common destination case, the multicommodity flow problem can be transformed into a single source/single destination max-flow problem, by connecting all sources to a (fictitious) common source (usually referred to as super-source). Hence, for the case of a single destination, TSD can be solved in polynomial time. A similar argument can be made for the single source case.

The above implies that SD can also be solved in polynomial time for the case of a single destination or source. This is achieved by performing a binary search on the value of T. Since the arrival time of each packet, when transmitted alone, is bounded by $M_{G,D}$ (the diameter of the graph with respect to D, see Section 2). It follows that $K \cdot M_{G,D}$ is an upper bound on the value of T. This increases the complexity of the solution by a multiplicative factor of $O(\log(K \cdot M_{G,D})) = O(\log(|V|))$.

4.3 Heuristics

Consider first the special case of directed tree networks where edges are directed from children to parents. The only schedule that can be considered is the trivial one where each node always chooses its unique outgoing edge. We now devise a departure plan, which is not necessarily optimal, but for which we can bound the termination time. This plan will serve us in building a heuristic for general graph topologies.

If the problem has at all a feasible solution, then the destination of each source-destination pair must be on the path between the source and the (common) root of the tree. Consider first a modified problem in which the root of the tree is the destination for all the packets in D. This is a common destination problem for which an optimal departure plan can be found. We use this departure plan for our (unmodified) problem, but route the packets only to their destination (instead of the root). It is not hard to see that this is a feasible departure plan which we now analyze.

Lemma 4.1 *For common destination directed tree networks (where the destination is the root of the tree), there is a departure plan whose termination time is bounded by* $K + M_{G,D}$.

Proof: We claim that departure times can be scheduled such that packets arrive at the root, one packet per time-slot, after some initialization time which is not greater than $M_{G,D}$. We prove this by induction on the number of vertices in the tree. It clearly holds for a tree containing two vertices. Let r be the root of the tree, let $r_1, r_2, \ldots, r_\ell$ be its neighbors in the tree, and let $T_1, \ldots, T_\ell$ be subtrees such that r_i is the root of T_i. For each subtree T_i, we denote by D_i the restricted demand D of the packets originating within T_i. If the destination of a packet is the root r, then its destination in D_i will be r_i. By the inductive hypothesis, for each subtree T_i, there exists a departure plan with the above property. Since the subtrees interact only in the root r, it is clear that by performing "sequentially" the departure plans computed for each of the subtrees, it is possible to obtain a departure plan in which packets arrive at the root r, one packet per time-slot. Clearly, in the worst case, the first packet will arrive at r after $M_{G,D}$ time slots. □

We refer to the algorithm implied by this proof as the *Train Algorithm*.

We note that the above upper-bound also applies to the inverse case, where edges are directed from the root to the leaves, and each destination is on a directed path from the corresponding source. In order to see that, reverse the direction of edges (so that now the root is a sink on the tree), and reverse the roles of sources and destinations. We thus have:

Corollary 4.1 *If G is a directed tree, in which the orientation is either towards the root, or, else, outwards from the root, and if there is a directed path from each source to each destination, then the Train Algorithm provides a departure plan with termination time of at most* $K + M_{G,D}$.

Consider now the case of an undirected tree. The problem is that since there is no single orientation, source-destination pairs may need conflicting orientations. However, we show how the Train Algorithm can be used in this case, such that termination times are still within a reasonable bound.

We begin by choosing a node in the above tree, which we refer to as the *separator*. A good way to choose a separator is discussed below. When the separator is removed, the tree is separated into components where each is a tree; denote by $C_1, C_2, \cdots, C_m$ these components with the separator attached to each. Denote by k_{ij} $(i \neq j)$ the number of packets whose source belongs to C_i while its destination belongs to C_j. The total number of intercomponent packets is therefore $f = \sum_{ij} k_{ij}$. We are interested in bounding the time required to deliver this intercomponent traffic.

Lemma 4.2 *The entire intercomponent traffic f for a given set of components $C_1, \ldots, C_m$ can be delivered within $f + 2M_{G,D} \log f$ time steps.*

We now proceed to provide a good solution for the case of an undirected tree by a recursive application of the Train Algorithm. It is known that every tree has a separator, that if removed, separates the tree into components, where each component is a tree containing at most half of the nodes of the original tree [19]. In a similar way it can be shown that every tree, some of whose nodes are sources, has a separator such that if removed, separates the tree into components, such that each component is a tree containing at most half of the sources of the original tree. We choose such a node as our separator.

We first send the intercomponent packets in the manner described in Lemma 4.2 and then apply the whole procedure, recursively, to the (undirected) tree defined by each component. Note that, in this case, different components do not share nodes, so the recursion can be applied concurrently to all components, i.e., packets will be routed concurrently as part of the recursive solution. Call this scheme the *Undirected Train Algorithm.*

Lemma 4.3 *The Undirected Train Algorithm guarantees a termination time of*

$$K + 2M_{G,D} \log^2 K$$

Proof: Let $\Gamma(k)$ be the execution time of the Undirected Train Algorithm for the given tree topology and for k packets. By Lemma 4.2 the intercomponent packets can be delivered within

$$\sum_{ij} k_{ij} + 2M_{G,D} \log(\sum_{ij} k_{ij}) \leq \sum_{ij} k_{ij} + 2M_{G,D} \log k$$

Due to the manner in which the separator is chosen, each component has at most $k/2$ packets. Hence, the recursive structure of the algorithm gives

$$\Gamma(k) \leq \Gamma\left(\frac{k}{2}\right) + \sum_{i,j} k_{ij} + 2M_{G,D} \log k$$

We note that the summation of the term k_{ij} in the recursive expression cannot exceed k. Also, Since there are at most $\log k$ stages to the recursion, the third term appears at most $\log k$ times. Thus $\Gamma(k)$ is bounded by $k + 2\log^2(k) M_{G,D}$. For the entire graph, $k = K$, which yields the required result. □

The Undirected Train Algorithm suggests a heuristic algorithm (HSD) for any general topology: compute a spanning tree Θ for G, and then run the Undirected Train Algorithm on Θ. This gives us:

Theorem 4.2 *For any graph G and demand D, algorithm HSD terminates in $O(|V|^2 \log^2 K)$ steps and provides a departure plan whose termination time is bounded by $K + 2|V| \log^2 K$.*

We note that, on a general graph, algorithm HSD makes use of just part of the edges. While the above analysis provides a guaranteed performance on general graphs, further improvements can be achieved by perturbating with the output of HSD, in an attempt to make use of other, non-tree edges.

5 Conclusion

In this paper we investigated efficient schemes for high-speed routing, that rely on techniques of both hot-potato and scheduled routing. Such schemes are advantageous because they require neither packet buffering nor routing decisions based on packet contents. Our analysis was carried in two directions. First, we indicated how routing should be planned after a schedule has been fixed and then analyzed the problem of designing an efficient schedule. For some of the problems tractable solutions that are optimal have been found. Other problems were shown to be intractable. For these problems we presented heuristic solutions. In particular, for the schedule design problem we obtained a recursive algorithm for which an efficient performance bound has been proved.

Several problems are still left open for further research. Additional versions of the general problem is the most obvious. Obtaining better heuristics is another one. For example, the heuristic algorithm presented for the design problem on a general topology currently achieves a bound of $K + 2|V| \log^2 K$. We would like to improve this bound to that of $O(K + M_{G,D} \log^2 K)$ which requires that $M_{\Theta,D} = O(M_{G,D})$. Whether such a spanning tree Θ exists and whether it can be constructed seem to be nontrivial problems. Spanner techniques [20] might prove to be effective in such a research effort.

Acknowledgement

We would like to thank Efim Dinitz for several useful discussions.

References

[1] I. Cidon, I. Gopal, and S. Kutten, "New models and algorithms for future networks," in *Proceedings of the 7th ACM Symposium on Principles of Distributed Computing*, pp. 74–89, 1988.

[2] Y. Yemini and D. Florissi, "Isochronets: A high-speed network switching architecture," in *Proceedings of IEEE INFOCOM*, pp. 740–747, 1993.

[3] U. Feige and P. Raghavan, "Exact analysis of hot-potato routing," in *Proceedings of IEEE Symposium on Foundations of Computer Science*, pp. 553–562, November 1992.

[4] A. Acampora and S. Shah, "Multihop lightwave networks: a comparison of store-and-forward and hot-potato routing," in *Proceedings of IEEE INFOCOM*, pp. 10–19, 1991.

[5] A. Greenberg and J. Goodman, "Sharp approximate models of deflection routing in mesh networks," in *Proceedings of IEEE INFOCOM*, pp. 307–318, 1983.

[6] A. Greenberg and B. Hajek, "Deflection routing in hypercube networks," *IEEE Transactions on Communications*, vol. 40, pp. 1070–1081, 1992.

[7] B. Hajek, "Bounds on evacuation time for deflection routing," *Distributed Computing*, vol. 5, pp. 1–6, 1991.

[8] N. Maxemchuck, "Comparison of deflection and store and forward techniques in the manhattan street and shuffle exchange networks," in *Proceedings of IEEE INFOCOM*, pp. 800–809, 1989.

[9] T. Szymanski, "An analysis of hot potato routing in a fiber optic packet switched hypercube," in *Proceedings of IEEE INFOCOM*, pp. 918–925, 1990.

[10] Z. Zhang and A. Acampora, "Performance analysis of multihop lightwave networks with hot potato routing and distance age priorities," in *Proceedings of IEEE INFOCOM*, pp. 1012–1021, 1991.

[11] Y. Mansour and B. Patt-Shamir, "Greedy packet scheduling on shortest paths," *Journal of Algorithms*, vol. 14, pp. 449–465, 1993.

[12] I. Cidon, S. Kutten, Y. Mansour, and D. Peleg, "Greedy packet scheduling," in *Proceedings of the 4th International Workshop on Distributed Algorithms*, (Bari), September 1990.

[13] P. Rivera-Vega, R. Varadarajan, and S. Navathe, "Scheduling data redistribution in distributed databases," in *Proceedings of the 6th International Conference on Data Engineering*, (Los Angeles), February 1990.

[14] J. J. Bae and T. Suda, "Survey of Traffic Control Schemes and Protocols in ATM Networks," *Proceedings of the IEEE*, vol. 79, pp. 170–189, February 1991.

[15] M. Bonuccelli, I. Gopal, and C. Wong, "Incremental time-slot assignment in SS/TDMA satellite systems," *IEEE Transactions on Communications*, vol. 39, pp. 1147–1156, 1991.

[16] M. Garey and D. Johnson, *Computers and Intractability.* San Francisco: Freeman, 1979.

[17] M. Golumbic, *Algorithmic Graph Theory and Perfect Graphs.* New York: Academic Press, 1980.

[18] E. Dinitz. Private communication, May 1993.

[19] B. Korte, L. Lovasz, and H. Promel, *Paths, Flow, and VLSI-Layout.* Berlin New-York: Springer-Verlag, 1990.

[20] D. Peleg and A. Schaffer, "Graph spanners," *Journal of Graph Theory*, vol. 13, pp. 99–116, 1989.