

Minimax Open Shortest Path First Routing Algorithms in Networks Supporting the SMDS Service

Frank Y.S. Lin

Jonathan L. Wang

Bell Communications Research

3 Corporate Place

Piscataway, New Jersey 08854

(908)699-4220 (908)758-5187

lin1@cc.bellcore.com jwang@cc.bellcore.com

Abstract

In this paper, we present and compare two quasi-static minimax Open Shortest Path First (OSPF) routing algorithms in networks supporting the SMDS service (which we refer to as the SMDS networks). In OSPF routing, the network is modeled as a graph and each link is associated with a nonnegative arc weight (referred to as the link set metric). A shortest path spanning tree is calculated for each origin to carry both the individually addressed and the group addressed (multicast) traffic. The OSPF routing protocol is adopted as a major part of the default Inter-Switching System Interface (ISSI) routing algorithm for SMDS networks where arc weights are inversely proportional to the aggregate link set capacities. We consider the problem of choosing a set of link set metrics so that the maximum link utilization factor is minimized in an SMDS network. The problem is formulated as a nonlinear mixed integer programming problem.

The emphasis of this paper is (i) to consider a quasi-static ISSI routing scheme and present a formal problem formulation and (ii) to develop an efficient and effective algorithm to solve the problem. A dual (based upon Lagrangean relaxation) and a primal approach are taken to solve the minimax OSPF routing problem. In computational experiments, the primal algorithm performs better for most of the test cases. However, it is suggested that the primal and the dual approach be applied in a joint fashion to achieve better performance. The joint algorithm determines good primal solutions for networks with up to 26 nodes in minutes of CPU time of a RISC-based file server. Compared with the default ISSI routing scheme, the minimax routing algorithm results in a 13.67% to 133.33% improvement in the maximum link utilization factor for 8 test networks.

1. Introduction

The SMDS (Switched Multi-megabit Data Service)^{[1] [2] [3]} service is a public, high speed, connectionless (datagram), packet switched data service that the Regional Bell Operating Companies (RBOCs) have offered^[4]. SMDS provides performance and features over a wide area similar to those found in local area networks, and it is regarded as the first phase of the Broadband Integrated Services Digital Networks (B-ISDNs). It will furnish subscribers with high speed access (1.5 Mbps and 45 Mbps in current specifications and SONET rates in the future) as well as multipoint connection (instead of point-to-point connection by today's private line facilities). Subscribers would access the service via a dedicated interface (Subscriber Network Interface (SNI)). The SMDS Interface Protocol (SIP)^[2], which is based on the DQDB (Distributed Queue Dual Bus) protocol^[5], running over SNIs is used to transport variable length user packets between the network (a set of Switching Systems (SSs) with SMDS interface capability) and the Customer Premises Equipment (CPE).

SIP only defines the protocol between CPE and the network (switching) equipment. SSs in the network are connected across the Inter-Switching System Interface (ISSI) using the ISSI protocol (ISSIP)^[6]. The ISSIP encapsulates SIP level 3 packets into ISSIP Level 3 Data Transport Protocol Data Units (L3 DTPDUs), and segments them into 53-byte cells at level 2. However, level 2 of the ISSIP can also be based on the B-ISDN Network-Node Interface. Level 1 of the ISSIP currently supports both DS3 and SONET rates for an ISSI link set. An ISSI link set consists of all ISSI links directly connecting two SSs.

The ISSIP provides, among other features, the necessary features to route a packet from its source SS to its destination SS. This routing procedure is defined in the Routing Management Protocol (RMP) of the ISSIP. The RMP is derived from Version 2 of the Open Shortest Path First (OSPF) specification^[7]. The main features of the RMP are as follows:

- The Routing Management Entity (RME) in every SS have identical routing databases;
- Each RME's database describes the complete topology of the RME's domain;

- Each RME uses its database and the Shortest Path First (SPF) algorithm to derive the set of shortest paths to all destinations from which it builds its routing table.

Each link set is assigned a nonnegative number in the RMP called the link set metric. The default link set metric of each link set is inversely proportional to the aggregate link set capacity. One can then apply standard shortest path algorithms, e.g. Dijkstra's algorithm^[8] to calculate a shortest path spanning tree for every origin. Ties are broken by choosing the switch with the lowest router ID number. The RMP is used to support two types of traffic - individually addressed messages and multicast messages. The individually addressed message is transmitted from the origin to the destination over the unique path in the shortest path spanning tree. The multicast message is destined for more than one destination (may not be for all destinations, which is referred to as broadcast). However, exactly one copy of the multicast message will be transmitted over every link in the shortest path spanning tree according to the ISSIP. A multicast message will be discarded by a leaf (termination) switch in the shortest path spanning tree if the message is not for any user connected to the switch.

It can easily be verified that by using the default link set metrics as the arc weights in the shortest path algorithm, the total link set utilization factors of the network is minimized. However, one potential drawback of using the default link set metrics is that there is no embedded mechanism to prevent a link set from being overloaded. To react to network loads, a minimax utilization routing algorithm is developed in this paper. The proposed algorithm can simply be regarded as a way of updating the arc weights used in the RMP so that the network load can be better balanced. The RMP is still the underlying routing mechanism. The objective of the minimax utilization routing algorithm is to find a route between every origin and its multicast destination(s) in a network so that the maximum link utilization of the network is minimized.

The major advantages of using the minimum of the maximum link utilization as the performance objective include:

- The performance measure (utilization) is a linear function of the routing decision variables, as opposed to a nonlinear function when other performance measures, e.g. packet delay or blocking probability, are used.
- The routing decisions made by the minimax routing algorithm usually do very well with respect to other major performance criteria in various networks such as the call blocking probability in circuit-switched networks, the packet delay in virtual-circuit based packet networks and the cell loss probability in B-ISDNs.
- It is clear that an optimal routing assignment (with respect to the minimax criterion) remains optimal if the traffic requirements grow uniformly.
- The minimax criterion can be treated as a goal from the viewpoint of the system to provide a balanced and robust operating point.
- A single performance indicator (the maximum link utilization factor) is provided. This single value can be used to derive upper bounds on other performance measures, e.g. end-to-end delay, call blocking rate and cell loss probability.
- For engineering tractability, end-to-end performance objectives are usually converted into link utilization constraints. The minimax routing then provides the most efficient utilization of the network capacity and precludes unnecessary capacity expansion.

Due to the discrete nature of the minimax routing problem where exactly one spanning tree is used to transmit packets for each root, there is no published research, to our knowledge, that has been attempted to solve this type of problems optimally. Tcha and Maruyama^[9] considered the problem of determining a path for each origin-destination (O-D) pair (the union of the selected path from a common origin is not necessarily a spanning tree) so as to minimize the maximum link utilization. In their formulation, no multicast traffic was considered. They developed a straightforward heuristic scheme, which is conceptually similar to the simplex method, to solve the minimax routing problem for connection-oriented networks. They tested the heuristic

on relatively small test problems with 15 to 95 O-D pairs (the test problems we considered are with 90 to 650 O-D pairs).

In this paper, the minimax OSPF routing problem is solved by two approaches. The first approach is based upon the Lagrangean relaxation technique which has been applied to solve the routing problem for virtual circuit networks^{[10] [11] [12] [13]}. The Lagrangean relaxation problem can be decomposed into three independent and easily solvable subproblems. The second approach is a primal heuristic which is simple but effective. In computational experiments, the primal approach is shown to perform better for most of the test cases. However, it is suggested that the primal and the dual approach be applied in a joint fashion to achieve better performance. The joint algorithm determines good primal solutions for networks with up to 26 nodes in minutes of CPU time on a file server.

In the next section, we give the notation used in the paper. Section 3 presents a formal definition of the routing problem addressed in this paper. We then propose two solution approaches in Section 4. Some computational results are reported in Section 5.

2. Notation

We model a network supporting the SMDS service as a graph $G(V, L)$ where the switches are represented by nodes and the communication channels are represented by links. We have the following notation:

- $V = \{1, 2, \dots, N\}$: the set of nodes in the graph (network).
- L : the set of links in the graph (network).
- W : the set of O-D pairs (with individually addressed traffic demand) in the network.
- γ_w (packets/sec): the mean arrival rate of new traffic for each O-D pair $w \in W$.
- α_r (packets/sec): the mean arrival rate of multicast traffic for each multicast root $r \in V$.
- P_w : the set of all possible simple directed paths from the origin to the destination for O-D pair w . (All the traffic for O-D pair w is transmitted over exactly one path in the set P_w .)
- P : the set of all simple directed paths in the network, that is, $P = \cup_{w \in W} P_w$.
- O_w : the origin of O-D pair w .
- T_r : the set of all possible spanning trees rooted at r for multicast root r . (The multicast traffic from r is transmitted over exactly one spanning tree in the set T_r .)
- T : the set of all spanning trees in the network, that is, $T = \cup_{r \in V} T_r$.
- C_l (packets per second): the capacity of link $l \in L$.
- a_l : the link set metric for link $l \in L$ (a decision variable).
- x_p : the routing decision variable which is 1 if path p is used to transmit the packets for O-D pair w and 0 otherwise.
- δ_{pl} : the indicator function which is 1 if link l is on path p and 0 otherwise.
- y_t : the routing decision variable which is 1 if tree $t \in T_r$ is used to transmit the multicast traffic originated at root r and 0 otherwise.
- σ_{tl} : the indicator function which is 1 if link l is on tree t and 0 otherwise.

Since all of the packets for an O-D pair are transmitted over exactly one path from the origin to the destination in OSPF routing, we have $\sum_{p \in P_w} x_p = 1$ and $x_p = 0$ or 1. The SMDS SSs are assumed to have the capability of duplicating multicast traffic for multiple downstream branches in a spanning tree. When a packet is multicasted from the root to the destinations using spanning tree t , it is assumed that exactly one copy of the packet is transmitted over each link in spanning tree t . Therefore, similar to the single-destination case, we have $\sum_{t \in T_r} y_t = 1$ and $y_t = 0$ or 1.

3. Problem Formulation

Based on the notation given in the previous section, the problem of determining a set of nonnegative link set metrics (and thus a spanning tree for each origin to transmit individually addressed and group addressed traffic) so as to minimize the maximum link utilization in the network can be formulated as the following nonlinear mixed integer programming problem:

$$Z_{IP'} = \min \max_{w \in W, p \in P_w} \frac{\sum_{r \in V} \sum_{t \in T_r} x_p \gamma_w \delta_{pl} + \sum_{r \in V} \sum_{t \in T_r} y_t \alpha_r \sigma_{tl}}{C_l} \quad (IP')$$

subject to:

$$\sum_{w \in W} \sum_{p \in P_w} x_p \gamma_w \delta_{pl} + \sum_{r \in V} \sum_{t \in T_r} y_t \alpha_r \sigma_{tl} \leq C_l \quad \forall l \in L \quad (1)$$

$$\sum_{p \in P_w} x_p = 1 \quad \forall w \in W \quad (2)$$

$$\sum_{t \in T_r} y_t = 1 \quad \forall r \in V \quad (3)$$

$$x_p = 0 \text{ or } 1 \quad \forall p \in P_w, w \in W \quad (4)$$

$$y_t = 0 \text{ or } 1 \quad \forall t \in T_r, r \in V \quad (5)$$

$$\sum_{w \in W} \sum_{p \in P_w} x_p \delta_{pl} \leq (N-1) \sum_{t \in T_r} y_t \sigma_{tl} \quad \forall l \in L, r \in V \quad (6)$$

$$\sum_{q \in P_w} \sum_{l \in L} a_l x_q \delta_{ql} \leq \sum_{l \in L} a_l \delta_{pl} \quad \forall p \in P_w, w \in W \quad (7)$$

$$a_l \geq 0 \quad \forall l \in L. \quad (8)$$

The objective function represents the minimax link utilization in the network. Terms in the left hand side of Constraint (1) denote the aggregate flow of packets over link l . Constraint (1) requires that the aggregate flow not exceed the capacity of each link. Constraints (2) and (4) require that all of the individually addressed traffic for an O-D pair be transmitted over exactly one path. Similarly, Constraints (3) and (5) require that all of the multicast traffic from one multicast root be transmitted over exactly one spanning tree. The left hand side of Constraint (6) (together with (2) and (4)) is the number of selected paths (for individually addressed traffic) that are rooted at origin r and pass through link l , while the right hand side of Constraint (6) (together with (3) and (5)) is equal to $N-1$ if link l is used in the spanning tree for root r to multicast messages and 0 otherwise. Recall that $N-1$ is the maximum number of selected paths originated at node r and passing through link l . Therefore, Constraint (6) requires that the union of selected paths from one origin to all the destinations for individually addressed traffic be the same spanning tree rooted at the origin to carry multicast traffic. (Note that this constraint implies that the selected paths from one origin to carry individually addressed traffic form a spanning tree.) The left hand side of (7) (together with (2) and (4)) is the routing cost for O-D pair w (for one unit of flow on the selected path). The right hand side of (7) is the cost of path $p \in P_w$. Constraint (7) requires that for each O-D pair a shortest path be used to carry the individually addressed traffic where the link set metric of link l is a_l . Constraint (8) requires that the link set metrics be nonnegative.

Let

$$s = \max_{l \in L} \frac{\sum_{w \in W} \sum_{p \in P_w} x_p \gamma_w \delta_{pl} + \sum_{r \in V} \sum_{t \in T_r} y_t \alpha_r \sigma_{tl}}{C_l}.$$

An equivalent formulation of IP' is

$$Z_{IP} = \min s \quad (IP)$$

subject to:

$$\sum_{w \in W} \sum_{p \in P_w} x_p \gamma_w \delta_{pl} + \sum_{r \in V} \sum_{t \in T_r} y_t \alpha_r \sigma_{tl} \leq C_l s \quad \forall l \in L \quad (9)$$

$$0 \leq s \leq 1 \quad (10)$$

$$\sum_{p \in P_w} x_p = 1 \quad \forall w \in W \quad (11)$$

$$\sum_{t \in T_r} y_t = 1 \quad \forall r \in V \quad (12)$$

$$x_p = 0 \text{ or } 1 \quad \forall p \in P_w, w \in W \quad (13)$$

$$y_t = 0 \text{ or } 1 \quad \forall t \in T_r, r \in V \quad (14)$$

$$\sum_{w \in W} \sum_{p \in P_w} x_p \delta_{pl} \leq (N-1) \sum_{t \in T_r} y_t \sigma_{tl} \quad \forall l \in L, r \in V \quad (15)$$

$$\sum_{q \in P_w} \sum_{l \in L} a_l x_q \delta_{ql} \leq \sum_{l \in L} a_l \delta_{pl} \quad \forall p \in P_w, w \in W \quad (16)$$

$$a_l \geq 0 \quad \forall l \in L. \quad (17)$$

Constraints (11)-(17) are the same as Constraints (2)-(8). Constraints (9) and (10) require that the utilization of each link not exceed s (and unity). Constraint (1) is therefore redundant and eliminated.

4. Solution Procedures

(IP), as formulated in the previous section, is a nonlinear and nonconvex mixed integer programming problem. Thus it is difficult to solve the problem optimally. In this section, two quasi-static algorithms to solve (IP) are proposed. The first approach is to solve (IP) using Lagrangean relaxation. The second approach is a primal heuristic which is simple but very effective as will be shown in the computational experiments.

4.1 Approach 1

The first approach to solving the mixed integer programming problem formulated in the previous section is Lagrangean relaxation. We dualize Constraints (9) and (15) of (IP) to obtain the following relaxation:

$$Z_D(u, b) = \min \left\{ s + \sum_{l \in L} u_l \left(\sum_{w \in W} \sum_{p \in P_w} x_p \gamma_w \delta_{pl} + \sum_{r \in V} \sum_{t \in T_r} y_t \alpha_r \sigma_{tl} - C_l s \right) + \sum_{r \in V} \sum_{l \in L} b_{rl} \left(\sum_{w \in W} \sum_{p \in P_w} x_p \delta_{pl} - (N-1) \sum_{t \in T_r} y_t \sigma_{tl} \right) \right\} \quad (\text{LR})$$

subject to:

$$0 \leq s \leq 1 \quad (18)$$

$$\sum_{p \in P_w} x_p = 1 \quad \forall w \in W \quad (19)$$

$$\sum_{t \in T_r} y_t = 1 \quad \forall r \in V \quad (20)$$

$$x_p = 0 \text{ or } 1 \quad \forall p \in P_w, w \in W \quad (21)$$

$$y_t = 0 \text{ or } 1 \quad \forall t \in T_r, r \in V \quad (22)$$

$$\sum_{q \in P_w} \sum_{l \in L} a_l x_q \delta_{ql} \leq \sum_{l \in L} a_l \delta_{pl} \quad \forall p \in P_w, w \in W \quad (23)$$

$$a_l \geq 0 \quad \forall l \in L. \quad (24)$$

(LR) can be decomposed into three independent subproblems. The first subproblem is to determine s . There are two cases to consider:

1. If $\sum_{l \in L} u_l C_l \geq 1$, then $s = 1$.
2. If $\sum_{l \in L} u_l C_l < 1$, then $s = 0$.

The second subproblem is to determine $\{x_p\}$. A solution to the second subproblem is for every O-D pair w (where $O_w = r$) to route all of the required traffic over a shortest path where the arc weight of link l is $\gamma_w u_l + b_{rl}$. All $a_l, l \in L$ are set to be 0 so that Constraint (23) is satisfied and no constraint is imposed on $\{x_p\}$ by (23). The third subproblem is to determine $\{y_t\}$. A solution to the third subproblem is for every root r to route all of the required multicast traffic over a minimum cost spanning tree where the arc weight of link l is $\alpha_r u_l - (N-1) b_{rl}$.

Dual (D):

For any $(u, b) \geq 0$, by the weak Lagrangean duality theorem^[14], the optimal objective function value of (LR), $Z_D(u, b)$, is a lower bound on Z_{IP} . The dual problem (D) is

$$Z_D = \max_{u, b \geq 0} Z_D(u, b). \quad (\text{D})$$

Therefore, in order to obtain the greatest lower bound, we solve the dual problem (D). There are several methods for solving the dual problem (D), of which the subgradient method^[15] is the most popular and is employed here. Let an $(|L| + |V| + |L|)$ -tuple vector g be a subgradient of $Z_D(u, b)$. Then, in iteration k of the subgradient optimization procedure, the multiplier vector $\lambda = (u, b)$ is updated by

$$\lambda^{k+1} = \lambda^k + t^k g^k. \quad (25)$$

The step size t^k is determined by

$$t^k = \delta \frac{Z_{IP}^k - Z_D(\lambda^k)}{\|g^k\|^2} \quad (26)$$

where Z_{IP}^k is the objective function value for a heuristic solution (upper bound on Z_{IP}) and δ is a constant, $0 < \delta \leq 2$.

The following property is useful for choosing an initial value for u and may improve the lower bounding procedure.

Property 1: If (IP) is feasible, then the following region contains an optimal solution $\{u_i^*\}$ to (D).

$$\sum_{l \in L} u_l^* C_l = 1. \quad (27)$$

Property 1 basically reduces the search region for an optimal solution to (D). One application of Property 1 is to choose the initial value of u_l to be $1/(|L| C_l)$. Another possible application of Property 1 is that an upper bound $1/C_l$ can be imposed on u_l when the subgradient method is applied. Moreover, since Property 1 provides a region that guarantees an optimal solution to (D) and $\{(u, b) | \sum_{l \in L} u_l C_l = 1\}$ is a break point of $Z_D(u, b)$ (s is discontinuous at this point), we employ the following mechanism to find better lower bounds. First, we temporarily choose the multiplier $u_l' = u_l (\sum_{l \in L} u_l C_l)^{-1}$ for link l so that $\sum_{l \in L} u_l' C_l = 1$. In addition, temporarily choose the multiplier $b_{rl}' = b_{rl} (\sum_{l \in L} u_l C_l)^{-1}$ for link l and origin r . Note that using the temporary set of multipliers (u', b') results in the same choice of paths as using the

original (u, b) . We also note that when all of the arc weights (multipliers) are multiplied by a positive scalar, the shortest paths found in solving (LR) do not change. It is only for the purpose of calculating $Z_D(u, b)$, we temporarily use the multiplier in order to find a higher dual objective function value through a simple calculation. The original value of the multipliers u_l and b_{rl} will still be used in the subgradient method to update the multiplier for the next iteration. The above procedure is for solving the dual problem and obtaining good lower bounds on the optimal objective function value. We next describe a procedure for finding good primal solutions.

Primal Solutions:

In each iteration of solving (D), a set of multipliers $\{u, b\}$ is used. A heuristic for determining the link set metrics is to let a_l be u_l . Dijkstra's shortest path algorithm is then applied to determine a shortest path spanning tree for each origin to deliver both the individually addressed and multicast traffic. The heuristic solution that results in the lowest maximum link utilization in the network is reported.

In Approach 1, Constraint (15) is dualized, which introduces $|L| + |V|$ multipliers (this dominates the complexity of (LR)) and $|V|$ minimum cost spanning tree problems in (LR). Since the amount of multicast traffic is expected to be small compared with the amount of individually addressed traffic and thus the significance of Constraint (15) is small, a more efficient solution procedure which is a variation of Approach 1 is introduced below to reduce the computational complexity.

4.2 A Variation of Approach 1

In this approach, only Constraint (9) is dualized and Constraint (15) is replaced by a "tree constraint" which states that the selected paths from one root for individually addressed traffic should form a spanning tree rooted at the root. (This "tree constraint" can be satisfied in the Lagrangean relaxation if Dijkstra's shortest path algorithm is applied.) With this change, the constraint that requires the spanning trees used to carry both individually addressed and group addressed traffic from the same root be the same is not considered in the modified (LR). We denote this modified (LR) problem by (LR'). Note that, despite the change, the optimal objective function value of (LR') is still a legitimate lower bound on Z_{IP} . In addition, it can be easily verified that Property 1 is still valid for the new dual problem.

Similar to Approach 1, A heuristic for determining the link set metrics is to let a_l be u_l . Dijkstra's shortest path algorithm is then applied to determine a shortest path spanning tree for each origin to deliver both the individually addressed and multicast traffic.

As mentioned earlier, this modified solution procedure has significantly lower complexity than Approach 1 -- $|L|$ multipliers versus $|L| + |L| + |V|$ multipliers in the respective dual problems and 2 versus 3 subproblems in the respective Lagrangean relaxation problems. Due to the expected small amount of multicast traffic, the compromise of primal solution quality with this modification is minimal. Consequently, this modified solution procedure is adopted in our computational experiments.

4.3 Approach 2

We next propose a primal approach to solving the minimax routing problem. The basic idea is to adjust a_l according to the current link flow. More precisely, if the utilization of link l is the maximum in the network, then we artificially increase a_l in an attempt to reduce the traffic flow of link l . It is clear that the utilization of link l does not increase when a_l is increased (if all the other link set metrics remain unchanged). However, in the event that three or more links share the same maximum link utilization, it is possible that the utilization of one of these most congested links increases after their corresponding link set metrics are increased simultaneously by the same value.

The overall algorithm is given below.

1. Assign an initial value to each a_l . Set the iteration counter k to be 1.
2. If k is greater than a prespecified counter limit, stop.
3. Apply Dijkstra's shortest path algorithm to calculate a shortest path spanning tree for each origin.
4. Calculate the aggregate flow for each link.
5. Identify the set of link(s) with the highest utilization, denoted by S .
6. For each $l \in S$, increase a_l by a positive value t^k .
7. Increase k by 1 and go to Step 2.

t^k can be chosen by different ways. However, the following two properties of $\{t^k\}$ are suggested: (i) $\sum_{k=1}^{\infty} t^k$ approaches infinity and (ii) t^k approaches 0 as k approaches infinity. The first property is meant to prevent the algorithm from being stalled. Whereas, the second property decreases the possibility of oscillation. If a sequence of t^k satisfies the first property, then every a_l will be unbounded when k approaches infinity. To avoid this difficulty, one may periodically normalize $\{a_l\}$, e.g. to make the sum of all link set metrics to be a given constant. However, in practice, only a finite number of iterations are allowed. If the initial value of each a_l is properly chosen, overflow will not occur.

There are a couple of advantages of the proposed primal heuristic:

1. The algorithm is simple and therefore may allow a larger number of iterations or shorter computational requirements.
2. Individually addressed traffic and multicast traffic are considered in a uniform way. More precisely, when one link is identified to be among the most congested, one simply increase the corresponding link set metric trying to reduce the amount of link traffic, which may be due to either traffic type.

5. Computational Results

In the computational experiments, we test the primal and the dual minimax routing algorithm with respect to their (i) computational efficiency, (ii) effectiveness in determining good solutions and (iii) capabilities of reducing the maximum link utilization (compared with the OSPF routing in conjunction with the default link set metrics).

The variation of Approach 1 discussed in Section 4.2 and Approach 2 discussed in Section 4.3 were coded in C¹ and run on a SUN SPARC file server². For the first set of experiments, the algorithms were tested on 8 networks. For each of the 8 networks, it was assumed that for each O-D pair the total individually addressed traffic rate at which packets are generated is 1 packet per second. It was also assumed that for each root the total group addressed traffic rate at which packets are generated is 0 (single destination traffic only) and 0.05 packets per second. The link capacities for each test network are assumed to be the same (homogeneous networks) with 100 packets per second for each link. Each experiment is run for 6000 iterations.

5.1 Performance of the Variation of Approach 1

For the variation of Approach 1, the subgradient method described in Section 4.1 was applied to solve (D). In our implementation, Z_p^* was initially chosen as 1 and updated to the best possible upper bound found so far in each iteration. In Equation (26), δ was initially set to 2 and halved whenever the objective function value did not improve in 80 iterations. The initial value of u_i was chosen to be $1/(|L| C_i)$. The mechanism (i.e., using temporary multipliers) described in Section 4.1 to improve the lower bounds were implemented. The routing assignment (set of paths) associated with the best heuristic to (IP) was used.

Table 1 summarizes the results of our computational experiments with the variation of Approach 1. The second column gives the total group addressed traffic rate at which each root generates. The third column is the largest lower bound on the optimal objective function value found in 6000 iterations. Note that this is the best objective function value of the dual problem. In addition to the mechanism described in Section 4.1 to improve the lower bound, we have attempted another mechanism based upon the fact that the set of possible objective function values for (IP) is discrete (equal to the number of O-D pairs using the most congested link divided by the link capacity when there is no group addressed traffic) in the experiments. The fourth column gives the best objective function value for (IP) in 6000 iterations. The percentage difference $(\text{upper-bound} - \text{lower-bound}) \times 100 / \text{lower-bound}$ is an upper bound on how far the best feasible solution found is from an optimal solution. The sixth column provides the CPU times which include the time to input the problem parameters.

Table 1 shows that the variation of Approach 1 is a sensible approach. However, the gaps between the lower and the upper bounds in some of the test cases are significant, which indicates that there may be still room to improve the solutions. Another observation from Table 1 is that the performance of the algorithm is comparable for the two demands of group addressed traffic. Moreover, the algorithm seems to perform better for larger networks, or more precisely, for networks with a larger number of active paths (carrying flows) passing through the most congested link in an optimal solution. For example, OCT, ARP2 and SWIFT have the highest three lower bounds and the lowest three error bounds. This observation can be explained from a linear programming point of view. Consider the linear programming relaxation of the integer programming problem (IP) with no group addressed traffic and without Constraints (12) and (14)-(17). Note that the constraint $x_p \leq 1$ is implied by Constraint (11) in the linear programming relaxation and can thus be ignored. It is clear that at most $|L| + |W|$ x_p 's can be nonzero in an optimal basic feasible solution to this problem. Next, we consider a Lagrangean relaxation of this linear programming relaxation problem where Constraint (9) is

dualized. It can be shown that in this case there exist an optimal solution to the primal problem (the linear programming relaxation problem) where routing assignments are determined by finding a shortest path for each O-D pair where the arc weight is the optimal Lagrangean multiplier. Our heuristic is to use the multipliers to determine a set of routing assignments and thus exactly $|W|$ x_p 's are 1 (and all the others are 0). Therefore, at most $|L|$ O-D pairs out of the $|W|$ O-D pairs can have multiple shortest paths. We thus can use $|L|/|W|$ as a measure to roughly predict the performance of the dual approach. The smaller the ratio is (and usually the larger the network is), the algorithm is expected to perform better. In our experiments, OCT, ARP2 and SWIFT, on which the best results are obtained, are indeed the networks with the three lowest $|L|/|W|$ ratios. Consequently, the dual approach is recommended for large networks (with low $|L|$ to $|W|$ ratios).

5.2 Comparison of the Variation of Approach 1 and the Default ISSI Routing

Next, we compare the variation of Approach 1 and the OSPF routing with the default link set metrics. For the purpose of illustration, we use the maximal aggregate link flow instead of the maximal link utilization factor as the performance measure (these two measures are equivalent for homogeneous networks). The traffic demands are assume to be the same as those considered in Table 1. The results are summarized in Table 2. The second column gives the total group addressed traffic rate at which each root generates. The third column reports Z_p^* , the best primal objective function value found by the proposed algorithm. The fourth column reports Z^{ISSI} which is the objective function value found by applying the default ISSI routing algorithm. The fifth column gives the percentage improvement of the variation of Approach 1 over the default ISSI routing algorithm. Table 2 shows that using the variation of Approach 1 results in an improvement in the maximum link utilization. The range of improvements is from 7% to 53.42%.

5.3 Performance of Approach 2

We next repeat the computational experiments performed in Tables 1 and 2 using Approach 2. The initial values of the link set metrics are chosen to be 1. The increase of link set metric for the link set(s) with the highest utilization in iteration k is $1/k$. This choice satisfies the two properties discussed in Section 4.3. Table 3 summarizes the results of our computational experiments with Approach 2. The columns are organized the same way as in Table 1.

Table 3 shows that Approach 2 is, in general, efficient and effective in finding a good set of link set metrics. From a comparison of Tables 1 and 3, Approach 2 is in general superior in terms of computation time and quality of solutions. However, for the cases corresponding to the SWIFT network where Approach 2 does not provide satisfactory solutions, the variation of Approach 1 solves the problem optimally (when the group addressed traffic demand is 0). This observation suggests that the dual and the primal approaches be applied in a joint fashion to achieve better performance. Note that, a side product of the dual approach is to provide lower bounds to evaluate the quality of the heuristic solutions.

Another observation from Table 3 is that the performance of Approach 2 is, in general, also comparable for the two demands of group addressed traffic. An interesting finding is that for the OCT network, a better set of link set metrics is found when the group addressed traffic demand for each root is 0.05 packet/sec (higher total network load but lower link utilization). On the other hand, for the SITA network, a better set of link set metrics is found when only single-destination traffic exists (in fact, an optimal solution is found in this case). If this set of link set metrics is used in the case where the multicast traffic demand is 0.05, the upper bound would become at most 0.0315.

5.4 Comparison of Approach 2 and the Default ISSI Routing

We also investigate the relative performance of Approach 2 and the OSPF routing with the default link set metrics. The performance measure is the maximum link utilization in the network. The results are summarized in Table 4. The traffic demands are assume to be the same as those considered in Tables 1 to 3. The columns are organized the same way as in Table 2. The results in Table 4 show that using Approach 2 results in an improvement in the maximum link utilization for up to 133%.

REFERENCES

1. J. Shantz, "SMDS: An Alternative to Private Network?" *Telecommunications Magazine*, pp. 52-54, May 1989.
2. "Generic System Requirements in Support of Switched Multi-Megabit Data Service," Bellcore Technical Reference, TR-TSV-000772, Issue 1, May 1991.
3. "Switched Multi-Megabit Data Service (SMDS) Operations Technology Network Element Generic Requirements," Bellcore Technical Reference, TR-TSV-000774, Issue 2, April 1990.
4. A. Ramirez, "Big Local Phone Companies to Offer Fast-Data Service," *The New York Times*, May 21, 1992.
5. R. M. Newman and J. L. Hullett, "Distributed Queuing: A Fast and Efficient Packet Access Protocol for QPSX," *Proc. 8th Int. Conf. Comp. Comm.*, pp. 294-199, Munich, 1986.

1. Approach 2 was originally coded in FORTRAN 77 and converted into a C program using the "f2c" utility for the purpose of comparing computational efficiency with Approach 1. It is expected that the computation time reported for Approach 2 should be lower if Approach 2 is coded in C directly.

2. Bellcore does not recommend or endorse products or vendors. Any mention of a product or vendor in this paper is to indicate the computing environment for the computational experiments discussed or to provide an example of a technology for illustrative purposes; it is not intended to be a recommendation or endorsement of any product or vendor. Neither the inclusion of a product or a vendor in a computing environment or in this paper, nor the omission of a product or vendor, should be interpreted as indicating a position or opinion of that product or vendor on the part of the authors or of Bellcore.

6. "Generic Requirements for SMDS Networking," Bellcore Technical Advisory, TA-TSV-001059, Issue 2, August 1992.
7. "The OSPF Specification, Version 2," Network Working Group Request for Comments (RFC 1131), OSPFIGP Working Group of the Internet Engineering Task Force, June 1990.
8. E.W. Dijkstra, *A Note on Two Problems in Connection with Graphs*, Numerical Mathematics, Vol. 1, pp. 269-271, October 1959.
9. D. Tcha and K. Maruyama, "On the Selection of Primary Paths for a Communication Network", *Computer Networks and ISDN Systems*, Vol. 9, pp. 257-265, 1985.
10. B. Gavish and S. L. Handler, "An Algorithm for Optimal Route Selection in SNA Networks", *IEEE Transactions on Communications*, COM-31, vol. 10, pp. 1154-1160, October 1983.
11. F. Y.S. Lin and J. R. Yee, "Three Algorithms for Routing and Flow Control in Virtual Circuit Networks", in *Proc. IEEE Globecom*, pp. 339-343, December 1990.
12. J. R. Yee and F. Y.S. Lin, "A Routing Algorithm for Virtual Circuit Data Networks with Multiple Sessions Per O-D Pair," *Networks*, May 1992.
13. F. Y.S. Lin and J. R. Yee, "A New Multiplier Adjustment Procedure for the Distributed Computation of Routing Assignments in Virtual Circuit Networks," to appear in *ORSA Journal on Computing*.
14. A. M. Geoffrion, "Lagrangean Relaxation and Its Use in Integer Programming", *Math. Programming Study*, Vol. 2, pp. 82-114, 1974.
15. M. Held, P. Wolfe and H. D. Crowder, "Validation of Subgradient Optimization", *Math. Programming*, Vol. 6, pp. 62-88, 1974.

Network ID	Group Addressed Traffic (%)	Lower Bounds	Upper Bounds	Percentage Difference (%)	CPU Time (sec)
OCT	0	0.570	0.590	3.5	357.55
	5	0.572	0.597	4.4	358.97
ARPA2	0	0.400	0.420	5.0	197.63
	5	0.403	0.443	9.9	195.87
NORDIC	0	0.100	0.130	30.0	56.76
	5	0.101	0.143	41.6	56.82
SWIFT	0	0.220	0.220	0.0	77.43
	5	0.222	0.226	1.8	77.87
SITA	0	0.030	0.050	66.7	33.78
	5	0.031	0.051	64.5	34.01
PSS	0	0.130	0.160	23.1	66.09
	5	0.135	0.153	13.3	66.48
GTE	0	0.060	0.080	33.3	47.60
	5	0.059	0.081	37.3	47.87
TRANSPAC	0	0.120	0.140	16.7	47.07
	5	0.118	0.141	19.5	46.53

Table 1. Summary of computational results of the variation of Approach 1

Network ID	Group Addressed Traffic (%)	Z_p^h (packets/sec)	Z^{ISSI} (packets/sec)	$\frac{Z^{ISSI} - Z_p^h}{Z_p^h}$ (%)
OCT	0	59.00	80.00	35.59
	5	59.65	80.65	35.21
ARPA2	0	42.00	47.00	11.90
	5	44.30	47.40	7.00
NORDIC	0	13.00	17.00	30.77
	5	14.30	17.25	20.63
SWIFT	0	22.00	26.00	18.18
	5	22.55	26.55	17.74
SITA	0	5.00	7.00	40.00
	5	5.10	7.35	44.12
PSS	0	16.00	18.00	12.50
	5	15.25	18.30	20.00
GTE	0	8.00	12.00	50.00
	5	8.05	12.35	53.42
TRANSPAC	0	14.00	17.00	21.43
	5	14.10	17.40	23.40

Table 2. Comparison of the maximal link utilization factor obtained by the variation of Approach 1 and the default ISSI routing

Network ID	Group Addressed Traffic (%)	Lower Bounds	Upper Bounds	Percentage Difference (%)	CPU Time (sec)
OCT	0	0.570	0.590	3.5	322.08
	5	0.572	0.587	2.6	323.68
ARPA2	0	0.400	0.410	2.5	176.95
	5	0.403	0.417	3.5	177.83
NORDIC	0	0.100	0.120	20.0	51.16
	5	0.101	0.122	20.8	51.41
SWIFT	0	0.220	0.260	18.2	69.55
	5	0.222	0.266	19.8	69.90
SITA	0	0.030	0.030	0.0	29.23
	5	0.031	0.042	35.5	29.38
PSS	0	0.130	0.130	0.0	59.01
	5	0.135	0.137	1.5	59.30
GTE	0	0.060	0.070	16.7	42.42
	5	0.059	0.072	22.0	42.63
TRANSPAC	0	0.120	0.130	8.3	41.69
	5	0.118	0.134	13.6	41.89

Table 3. Summary of computational results of Approach 2

Network ID	Group Addressed Traffic (%)	Z_p^h (packets/sec)	Z^{ISSI} (packets/sec)	$\frac{Z^{ISSI} - Z_p^h}{Z_p^h}$ (%)
OCT	0	59.00	80.00	35.59
	5	58.70	80.65	37.39
ARPA2	0	41.00	47.00	14.63
	5	41.70	47.40	13.67
NORDIC	0	12.00	17.00	41.67
	5	12.15	17.25	41.98
SWIFT	0	26.00	26.00	0.00
	5	26.55	26.55	0.00
SITA	0	3.00	7.00	133.33
	5	4.20	7.35	75.00
PSS	0	13.00	18.00	38.46
	5	13.65	18.30	34.07
GTE	0	7.00	12.00	71.43
	5	7.15	12.35	72.73
TRANSPAC	0	13.00	17.00	30.76
	5	13.35	17.40	30.34

Table 4. Comparison of the maximal link utilization factor obtained by Approach 2 and the default ISSI routing