Telecommunication Systems 3(1995)397–413

An analytical approach to the dynamic topology problem★

Chelsea C. White, III
Department of Industrial and Operations Engineering, University of Michigan, Ann Arbor, MI 48109-2117, USA

Edward A. Sykes
Department of Systems Engineering, University of Virginia, Thornton Hall, Charlottesville, VA 22903-2442, USA

Janet A. Morrow
U.S. Army Foreign Science and Technology Center, 220 Seventh Street NE, Charlottesville, VA 22903, USA

Received June 1992; revised August 1994

Currently, it is possible to modify (say, hourly) the topology of a data communications network by adding or deleting network links and/or by increasing or decreasing bandwidth on existing links in response to changing traffic loads and/or projected network conditions. The intent of this paper is to study a Markov decision process (MDP) model of the dynamic topology problem (DTP), the problem of activating and/or deleting links, as a function of the current traffic in the network and of the most recent network topology design. We present a decomposition of this model and structural results for the decomposition. The decomposition and structural results enhance the tractability of procedures for determining optimal link control policies. A numerical example is used to illustrate these results.

1. Introduction

With the advent of data networking technology in the late 1960's and early 1970's came a paradigm for their design and construction based entirely on the process for acquiring transmission media. Originally comprised entirely of various voice grade telephone lines with a maximum line speed not exceeding 19.2 kilobits/second

★Research supported by ARO Contract No. DAAG-29-85-K0089, NSF Grant No. ECS-8708183, and DCA Contract No. DCA 100-89-C-0031.

(kb/s), transmission technology has evolved considerably, through wideband analog circuits with a maximum speed of 50 kb/s and into the world of digital circuits of speeds ranging from 56 kb/s to 1544 kb/s (T1) and higher. Yet the basic mechanism for the acquisition of transmission circuits remains unchanged: Circuits traditionally have been leased on a long-term basis, charged monthly, from the various vendors (AT&T, MCI, etc.), and there is often a significant time delay in between the time a circuit is ordered and the time it is in place. This process continues today along with the static design paradigm it engendered (discussed below).

The technology of digital communications, however, continues to evolve. It is becoming increasingly possible (using digital cross-connect technology) for data circuits to be acquired on a per-use basis (i.e. in the same way that a person can place a call of brief duration, a packet switch can now establish a circuit to another for a brief period of time). This possibility has a number of implications for the operation and design of data networks. Most notably, if we presume for the moment that data networks can adapt their topologies effectively to changing traffic requirements and/or network performance (in essence, redesign themselves in near real-time), what currently constitutes design may radically change.

The status quo in data network design is the design of a static network topology driven by a fixed (i.e. expected) traffic matrix. The traffic matrix is typically constructed using measured or estimated peak traffic rates between node pairs across the network (where different node pairs may reach their peak traffic rates at different times of the day), so that the final traffic matrix (an input to the design process) represents an upper bound on traffic at any one point in time. Yet in the case of a network geographically distributed across the United States, one can easily imagine that traffic will peak in the East, then the Midwest, then the West, according to normal business hours. Depending on the cost trade-offs between temporary and permanent (leased) circuits, a dynamic response in the network topology to the traffic variations it measures in real time may eliminate the need for a large number of the circuits which would have been present in the static design, at least for long periods of time. It may also provide circuits not anticipated in the static design due to unforeseen traffic requirements. Indeed, with the advent of many distributed computing applications, high bandwidth services will increasingly be required, not to create economies of scale, but rather to support individual computing requirements of relatively short duration. In short, dynamic topologies may produce both significant cost and performance gains.

What constitutes network design in the context of dynamic topologies is at best speculation until the mechanisms underlying the topologies and the costs and benefits associated with dynamic topologies have been determined. It is clear, however, that at least the potential exists for: (i) entirely dynamic topologies, or (ii) dynamic augmentation of a skeletal permanent backbone topology, depending on the relative prevalence of permanent and temporary circuits and the purpose for which each type of circuit is used.

This paper is organized as follows. A brief literature review is provided in section 2. The MDP model of the DTP is presented in section 3. We present standard results for the MDP in section 4, and the numerical implications of those results are discussed in section 5. These implications indicate the need to examine special cases of the DTP in order to enhance the tractability of the MDP model. Section 6 lists assumptions that lead to an MDP decomposition. Structural policy results for the decomposition are presented in section 7. Section 8 presents a numerical example that illustrates the decomposition and the structured results. Conclusions and future work are listed in section 9.

2. Literature review

Investigation of approaches to optimizing network topologies in response to changes in traffic demands is relatively new. Related modeling and optimization work is fairly abundant for self-organizing networks, principally radio networks (see Noakes et al. [15] and Shor and Robertazzi [18] for two examples and further literature in that area). Some additional related work in modeling and optimization of dynamic hierarchical networks is also available (see Moose [12], Nance and Moose [14], Moose and Nance [13] for key literature reviews of work by these and other authors). To our knowledge, however, the first and, to date, only paper to directly address the subject of this paper (i.e. the DTP for packet/router networks) is LeBlanc [9], which did so in a static, mathematical programming framework (see also related work by LeBlanc and Harder [8] and LeBlanc and Simmons [10]). It is instructive to note that LeBlanc's literature review also contained no references to prior work on this problem. Our consideration of the DTP for packet/router networks is novel, therefore, and our MDP approach to is appears to be unique.

Recent applications of MDPs to other communications networking problems are relatively abundant, however. Rouskas and Ammar [17] used an MDP approach in dynamically configuring paths in multihop wave division multiplexing lightwave networks. Hwang [6] developed an MDP approach to state-dependent virtual path routing in high-speed networks. Ross and Tsang [16] and Gopal and Stern [22] provided MDP approaches to the access port design problem for integrated services networks, the former paper employing policy iteration methods, and the latter value iteration and linear programming methods (the three approaches together make up the usual range of MDP solution techniques). MDPs have also been applied in such areas as: optimal vacation control for an $M/G/1$ queue [1], a problem of interest in local area network modeling and optimization; server allocation [11], and flow control [3,4].

3. An MDP model of the DTP

Let N represent the set of *nodes* in the network. We assume there are two types of *links*:

1. *Permanent* links, $A_P \subseteq N \times N$.
2. *Switchable* links, i.e. links that can be activated and deactivated, $A_S \subseteq N \times N$.

We assume $A_P \cap A_S = \varnothing$, the null set.

Let $a_l(t) = 1$ $(=0)$ if link $l \in A_S$ is activated (is not activated) during the time period $[t, t+1)$. Let $a(t) = \{a_1(t), \ldots, a_L(t)\}$, where $L = |A_S|$, the cardinality of the set A_S. The vector $a(t)$ represents the decision variable at time t.

Let $\tau_{ij}(t)$ be the number of message units per unit time at time t that originate at node $i \in N$ and have node $j \in N$ as their destination, and let $\tau(t) = \{\tau_{ij}(t)\}$. We assume that $\{\tau(t), t = 0, 1, \ldots\}$ is a time and action invariant Markov chain having known transition probabilities $P[\tau(t+1) \mid \tau(t)]$. For simplicity in this study, we have assumed that these transition probabilities are time invariant. In general, it may be more realistic to assume that they are periodic (with a period, say, of 24 hours).

Let $s(t) = \{\tau(t), a(t-1\}$, which we will call the *state* at time t. Note that $\{s(t), t = 0, 1, \ldots\}$ is a controlled Markov process. Let $c(s, a)$ be the expected cost to be accrued over the period $[t, t+1)$, assuming $s(t) = s$ and $a(t) = a$. We assume $c(s, a)$ is of the form

$$c(s,a) = c_{\$}(s,a) + c_{D}(s,a),$$

where $c_{\$}(s, a) = wC_{\$}(s, a)$, $c_D(s, a) = (1-w)C_D(s, a)$, w is a given trade-off constant, $C_{\$}(s, a)$ is the operating cost of the network, $C_D(s, a)$ is the expected cost associated with expected delays in the network. Let c_l^1, c_l^2 and c_l^3 be the nonnegative, time-invariant costs of activating link l, deactivating link l, and having link l activated during the period $[t, t+1)$, respectively. We assume that

$$C_{\$}(s,a) = \sum_{l} C_{\$l}(s,a),$$

where

$$\begin{aligned} C_{\$l}[a_l(t-1), a_l(t)] &= c_l^1 \max\{a_l(t) - a_l(t-1), 0\} \\ &\quad + c_l^2 \max\{a_l(t-1) - a_l(t), 0\} + c_l^3 a_l(t). \end{aligned}$$

Let $z_{ij}(s, a)$ be a measure of the expected delay of a message unit originating at node i and having node j as its destination during the period $[t, t+1)$. Then,

$$C_D(s,a) = b \sum_{i \neq j} \tau_{ij}(t) z_{ij}(s,a),$$

where $s(t) = \{\tau(t), a(t-1)\} = s$, $a(t) = a$, and b is the per unit cost of delay. Specific definitions of $z_{ij}(s, a)$ can vary according to network type and routing. As an example of a specific definition, consider: the expected message delay for a message whose origin is node i and whose destination is node j in seconds. This is the traditional definition – see Kleinrock [7, p. 317] and note the intentional similarity of notation there and here; note also that $C_D(s, a)$ under this definition of z is a cost-commensurated average packet delay. A second example definition is: the average total message delay from node i to node j in hops, as will be illustrated in a numerical example below.

We note that $z_{ij}(s, a)$ in the above expression must be dependent (albeit implicitly) on the routing protocol. Further note that the entirety of our model's dependency on routing is captured in the set $\{z_{ij}(s, a)\}$. The nature and mechanism of this dependence itself depends on the type of network and routing protocol(s) employed. For example, consider the case of dynamic delay-based routing in a packet-switched network, where the routing pattern adapts to changes in traffic. In this case, it may be nearly impossible to compute $z_{ij}(s, a)$ exactly using known analytical means and hence we must resort to approximations to obtain $\{z_{ij}(s, a)\}$. At the opposite extreme, in the case of high-speed circuit-switched routing using paths precomputed and stored in a table, the $\{z_{ij}(s, a)\}$ are deterministic.

A *policy* at time t is a mapping that selects an action on the basis of the current state; i.e. the decision to activate or deactivate any switchable link at time t assumes knowledge of $s(t) = \{\tau(t) = a(t-1)\}$. Let $A = \{0, 1\}^L$ be the set of all possible actions, and let S be the (finite) state space, i.e. the set of all possible values of $s(t)$ for any t. Both S and A are assumed to be time invariant. Thus, a policy at time t is a function $\delta_i : S \to A$. Let $\tilde{\Delta}$ be the set of all policies. A *strategy* is an infinite sequence of time-dependent policies; i.e. $\pi = \{\delta_1, \delta_2, \dots\}$. Our *objective* is to find a strategy that minimizes the following *criterion*:

$$E\left\{\sum_{t=0}^{\infty} \beta^t c[s(t), a(t)] \mid s(0)\right\},$$

where $\beta \in [0, 1)$ is the given *discount factor* and $E\{\cdot \mid s(0)\}$ is the expectation operator, conditioned on initial state $s(0)$.

We have assumed that the action set A equals $\{0, 1\}^L$. In general, however, it may be a strict subset of $\{0, 1\}^L$. This would occur if (N, A_P) is not sufficiently connected and some of the links in A_S must be activated in order to guarantee the required level of connectivity. Under these circumstances, determination of A may not be a simple task and/or the description of A may be difficult to deal with numerically. In seems clear, however, that $a \in A$ and $a \le a'$ would imply $a' \in A$.

4. Optimality equations

We now present optimality equations for the MDP model of the DTP. Define the operators $\tilde{H}_\delta$ and $\tilde{H}$ as

$$[\tilde{H}_\delta v](s) = c[s, \delta(s)] + \beta \sum_{s'} P[s' \mid s, \delta(s)] v(s')$$

and

$$\tilde{H}v = \inf_{\delta} \tilde{H}_{\delta} v,$$

respectively, where $\delta \in \tilde{\Delta}$ and where $P[s' | s, \delta(s)]$ is $P[s(t+1) = s' | s(t) = s]$, indexed by δ. Then, standard results (fould and explained in detail, for example, in Bertsekas [2, section 5.3] imply that

1. $\tilde{H}$ has a unique fixed point $\tilde{v}$; i.e. $\tilde{v} = \tilde{H}\tilde{v}$.
2. If v_0 is such that $\|v_0\| = \max\{|v_0(i)| : i \in S\}$ is bounded, then $\lim_{n\to\infty} \|v_n - \tilde{v}\| = 0$, where $v_{n+1} = \tilde{H} v_n$ (which is the successive approximations algorithm).
3. If $\tilde{\delta} \in \tilde{\Delta}$ is such that $\tilde{v} = \tilde{H}_{\tilde{\delta}}\tilde{v} = \tilde{H}\tilde{v}$, then the time invariant strategy $\pi = \{\tilde{\delta}, \tilde{\delta}, \dots\}$ is optimal.

These results provide a theoretical basis for the numerical determination of an optimal strategy for the DTP.

5. Numerical considerations

We now investigate the numerical implications of using successive approximations to determine $\tilde{v}$ and $\tilde{\delta}$. We note that $|S||A||\bar{S}|$ multiplications (and $|S||A||\bar{S}|$ additions) are required to determine v_{n+1} from v_n, where $\bar{S}$ is the set of all non-zero values of $P(s' | s, a)$ for each $s \in S$ and $a \in A$. Let σ be the number of possible values $\tau_{ij}(t)$ can take, for any (i, j) pair. Then $|S| = 2^L \sigma^{|N|(|N|-1)}$ and $|\bar{S}| = \sigma^{|N|(|N|-1)}$. Clearly, $|A| = 2^L$. Thus, $2^{2L}\sigma^{2|N|(|N|-1)}$ multiplications (and additions) are required to determine v_{n+1} from v_n. We note that this implies that even for modestly sized networks, the standard MDP approach to solution determination will be numerically infeasible. We now investigate a decomposition procedure that can significantly enhance the tractability of the MDP approach to optimal policy determination.

6. Structural assumptions

It is clear from earlier discussions that

$$P[\underline{s} = (\underline{\tau}, \underline{a}) | s' = (\tau', a'), a] = P(\underline{\tau} | \tau') P(\underline{a} | a),$$

where $P(\underline{a} | a) = 1 (= 0)$ if $\underline{a} = a (\neq \underline{a})$.

Assume the set $\mathcal{P}_l \subseteq N \times N$, associated with link $l \in A_S$, is such that $(i, j) \in \mathcal{P}_l$ if and only if $z_{ij}(s, a)$ depends on only $s_l = \{s_{ij}, (i, j) \in \mathcal{P}_l\}$ and a_l. That is, $z_{ij}(s, a) = z_{ij}(s_l, a_l)$ if and only if $(i, j) \in \mathcal{P}_l$. Thus, $(i, j) \in \mathcal{P}_l$ indicates that the expected delay of a message unit originating at node i and having destination j is affected only by the activation or deactivation of link l and by the message rates from i to j for only pairs $(i, j) \in \mathcal{P}_l$.

Assume $\mathcal{P}_l \cap \mathcal{P}_k = \emptyset$, for all $l \neq k$. Consideration of the case where $\mathcal{P}_l \cap \mathcal{P}_k$ is not necessarily null, for $l \neq k$, is a topic for future research. Our efforts so far, however, indicate that the case is treatable and is likely to require the results in White and Schlussel [19], who assume a decentralized information structure in order to reduce computational complexity (and often more accurately mirror reality).

A general comment can be made at this time, however, which applies to all cases. It is instructive to note that the *choice* of the sets of permanent and switchable links (A_P and A_S, respectively), along the routing procedure, will determine the sets $\{\mathcal{P}_l: l \in A_S\}$ and correspondingly whether the condition $\mathcal{P}_l \cap \mathcal{P}_k = \emptyset$, for all $l \neq k$, holds. Indeed, while this paper is largely directed at determining topology control policies given A_P, A_S, and routing, a broader issue is that of design, i.e. selection of A_P and A_S, given the routing procedure and the nature of control strategy. Hypothetically (and perhaps rather unlikely, given White and Schlussel), if it turned out to be computationally infeasible or impractical to determine a topology control policy when $\mathcal{P}_l \cap \mathcal{P}_k = \emptyset$, for all $l \neq k$, one could place the requirement that $\mathcal{P}_l \cap \mathcal{P}_k = \emptyset$, for all $l \neq k$ on the process that selects A_P and A_S.

The assumption that $\mathcal{P}_l \cap \mathcal{P}_k = \emptyset$, for all $l \neq k$, permits the DTP to be decomposed into L independent MDPs, the lth of which has the following optimality equation:

$$v_l(\tau_l, \bar{a}_l) = \min \begin{Bmatrix} \tau_l z_l(\tau_l, \bar{a}_l, 0) + \beta \sum P_l(\tau_l' \mid \tau_l) v_l(\tau_l', 0), \\ B(\bar{a}_l) + \tau_l z_l(\tau_l, \bar{a}_l, 1) + \beta \sum P_l(\tau_l' \mid \tau_l) v_l(\tau_l', 1), \end{Bmatrix},$$

where the summations are over all τ_l', where

$$\tau_l = \{\tau_{ij}: (i,j) \in \mathcal{P}_l\},$$

$$\tau_l z_l(\tau_l, \bar{a}_l, a) = \sum_{(i,j) \in \mathcal{P}_l} \tau_{ij} z_{ij}(\tau_l, \bar{a}_l, a),$$

$$B(\bar{a}_l) = c_l^1 \max\{1 - \bar{a}_l, 0\} + c_l^2 \max\{\bar{a}_l - 1, 0\} + c_l^3;$$

$\bar{a}_l$ represents the status of link l one time unit in the past, where we assume that $P(\tau' \mid \tau) = P_l(\tau_l' \mid \tau_l) P_{\bar{l}}(\tau_{\bar{l}}' \mid \tau_{\bar{l}})$ for $\tau = \{\tau_l, \tau_{\bar{l}}\}$ and $l \neq \bar{l}$, and where $P_l(\tau_l' \mid \tau_l)$ is the conditional probability associated with the lth MDP. We now examine the structural properties implied by the decomposed optimality equation which, for notational simplicity, we define

$$[Hv](\tau, \bar{a}) = \min \begin{Bmatrix} \tau z(\tau, \bar{a}, 0) + \beta \sum_{\tau'} P(\tau' \mid \tau) v(\tau', 0), \\ B(\bar{a}) + \tau z(\tau, \bar{a}, 1) + \beta \sum_{\tau'} P(\tau' \mid \tau) v(\tau', 1) \end{Bmatrix}.$$

Determination of Hv, given v, requires $4^{2|\mathcal{P}_l|}$ multiplications, which for reasonably sized σ and $|\mathcal{P}_l|$ represents a tractable problem. Note that the number of multiplications per iteration of successive approximations for the decomposed MDP is an exponential reduction, relative to the number of multiplications per iteration of successive approximations, from the original MDP model.

7. Structural results

We now give a reasonable condition which implies that the optimal policy for the decomposed MDP is isotone (monotonically nondecreasing) in both τ and $\bar{a}$. We begin with an important preliminary result.

LEMMA

Assume that:

(i) $\sum_{\tau' \in \mathcal{K}} P(\tau' | \tau)$ is isotone in τ for each set $\mathcal{K}$ such that $\tau \in \mathcal{K}$ and $\tau \le \tau'$ imply that $\tau' \in \mathcal{K}$.

(ii) The function $z(\tau, 0, a) - z(\tau, 1, a)$ is nonnegative and isotone in τ and a.

(iii) The function $z(\tau, 0, 0) - z(\tau, 1, 1)$ is nonnegative and isotone in τ.

Then, if the real-valued function v is such that $v(\tau, 0) - v(\tau, 1)$ is isotone in τ:

(a) There exists a policy δ that achieves the minimum in Hv and is isotone in τ and $\bar{a}$.

(b) $[Hv](\tau, 0) - [Hv](\tau, 1)$ is isotone in τ.

Assumption (i) indicates that the transition matrix describing $\{\tau(t), t = 0, 1, \ldots\}$ has, in the terminology used by Derman [5] in determining structured policy results for machine replacement problems, an *increasing failure rate* (IFR). Restricted to the scalar case and using the terminology here, assumption (i) implies that the probability that the level of traffic in the system will make transition into level of traffic k or above is never decreasing as a function of the current level of traffic, and that this condition holds for all k. This appears to be reasonable assumption for many communications networks.

Assumptions (ii) and (iii) concern properties of certain differences in expected delays, where the expected delay is assumed to be a function of traffic (τ), link status one unit of time in the past ($\bar{a}$), and current link status (a). These assumptions are stated in the above form for complete generality. In most cases of practical interest, however, $z(\tau, \bar{a}, a) = z(\tau, a)$, i.e. expected delay depends only on traffic and current link status. In our travails thus far, we have considered three such cases of the $z(\tau, \bar{a}, a)$, given below in order of increasing generality. We present these cases now, and then turn to proving the lemma.

Case 1

Assume that $z(\tau, \bar{a}, a) = z(a)$. Then, assumption (ii) is trivially satisfied and the condition $z(0) - z(1) \geq 0$ satisfies assumption (iii). This case arises when expected delay is traffic independent, i.e. is affected only by link status. A network where the delay from node i to node j is assumed to be a constant times the number of links traversed from i to j is an example of this case.

Case 2

$(i,j) \in \mathcal{P}_l$ implies that $z_{ij}(\tau, \bar{a}_l, a_l) = z_{ij}(\tau_{ij}, a_l) = h_{ij}(a_l)/(\mu - \tau_{ij})$, where assumption (ii) is trivially satisfied, and the condition $h_{ij}(0) > h_{ij}(1)$ satisfies assumption (iii) for all $\tau_{ij} < \mu$. If $h_{ij}(a_l)$ is the number of links traversed from node i to node j and μ is the average link transmission rate in messages per unit time (note: μ may be dependent on i and j), $z_{ij}(\tau, \bar{a}_l, a_l)$ can be interpreted as the expected delay accrued along a sequence of $h_{ij}(a_l)$ independent $M/M/1$ queues in tandem seeing only the traffic τ_{ij}. This case has arisen in a real-time computing test-bed, where a full-period backbone is to be augmented in a cost-effective manner using temporary circuits to handle large, short-duration traffic requirements between certain node pairs and there is little other traffic.

Case 3

$(i,j) \in \mathcal{P}_l$ implies that $z_{ij}(\tau, \bar{a}_l, a_l) = z_{ij}(\tau_l, a_l) = \sum_{r \in R(i,j,a_l)} T_r(\tau_l, a_l)$, where $R(i,j,a_l)$ = set of indices of circuits traversed in the route from node i to node j, given link status a_l, and $T_r(\tau_l, a_l) = 1/[\mu_r - \lambda_r(\tau_l, a_l)]$ is the expected delay on circuit r, modeled as an $M/M/1$ queue with mean service rate μ_r and mean arrival rate $\lambda_r(\tau_l, a_l)$. Note that $\lambda_r(\tau_l, a_l) = \sum_{(i,j) \in \mathcal{P}_l} x_{ijr}(a_l)\tau_{ij}$, where $x_{ijr}(a_l) = 1[=0]$ when $r \in R(i,j,a_l)[\notin R(i,j,a_l]$ for all i, j, r, and a_l (i.e. $\{x_{ijr}(a_l)\}$ is a set of indicator variables specifying the circuits used by each node pair $(i,j) \in \mathcal{P}_l$ as a function of a_l). For this case, assumption (ii) holds trivially, as before. Assumption (iii) can be satisfied by appropriately selecting two routing patterns $\{R(i,j,0), (i,j) \in \mathcal{P}_l\}$ and $\{R(i,j,1), (i,j) \in \mathcal{P}_l\}$, using some standard non-bifurcating routing procedure (e.g. minimum hop routing, minimum delay routing on the unloaded network, etc.), as long as $\lambda_r(\tau_l, a_l) < \mu_r$ for all r, τ_l, and a_l.

Case 3 is considerably more general than either case 1 or 2, but as such requires more careful attention to the determination of $\mathcal{P}_l$. A subtlety in determining $\mathcal{P}_l$ for pairs (i,j) whose intra-pair *routes* do not change as a result of activating or deactivating link l is that they still can have their *delays* affected by its presence or lack thereof. As an illustration, consider the network fragment given in fig. 1, assume that all message rates are zero save τ_{97}, τ_{73}, and τ_{93}, and for simplicity assume minimum hop routing. If the switchable link is deactivated (off), all message flows are zero save $\lambda_{97} = \tau_{97} + \tau_{93}$ and $\lambda_{73} = \tau_{73} + \tau_{93}$. If the switchable link is activated

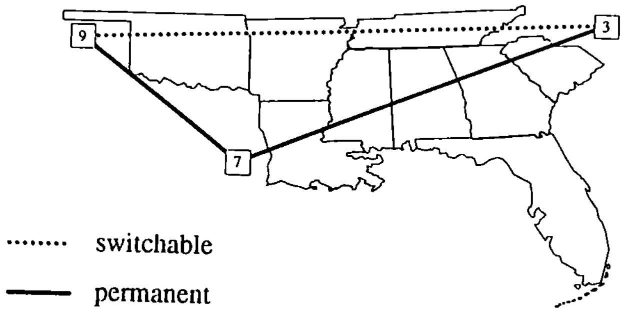

Fig. 1. Network fragment.

(on), all message flows are zero except $\lambda_{97} = \tau_{97}$, $\lambda_{73} = \tau_{73}$, and $\lambda_{93} = \tau_{93}$. In this case, $\mathcal{P}_{93}$ must include at least the pairs (9,7), (7,3), and (9,3). An an aside, note that if the circuits here are modeled as *M/M/*1 queues, assumption (iii) is satisfied.

We now conclude our discussion of the lemma with its proof.

Proof of the lemma

The isotonicity of δ in $\bar{a}$ follows directly from the fact that $c^1 \geq 0$. With respect to the isotonicity of δ in τ, note that $\delta(\tau, \bar{a}) = 0$ if and only if

$$\tau z(\tau, \bar{a}, 0) + \beta \sum_{\tau'} P(\tau' | \tau) v(\tau', 0) \leq B(\bar{a}) + \tau z(\tau, \bar{a}, 1) + \beta \sum_{\tau'} P(\tau' | \tau) v(\tau', 1)$$

or equivalently

$$\tau[z(\tau, \bar{a}, 0) - z(\tau, \bar{a}, 1)] + \beta \sum_{\tau'} P(\tau' | \tau)[v(\tau', 0) - v(\tau', 1)] \leq B(\bar{a}).$$

It follows from results in Derman that the left-hand side of the latter inequality is isotone in τ, which implies that $\delta(\tau, \bar{a})$ is isotone in τ.

We now show that $[Hv](\tau, 0) - [Hv](\tau, 1)$ is isotone in τ. Let $\tau \leq \tau'$. There are six possible cases to consider (see table 1). All other possibilities are ruled out because of the isotonicity of δ. Let

$$\Delta H = \{[Hv](\tau', 0) - [Hv](\tau', 1)\} - \{[Hv](\tau, 0) - [Hv](\tau, 1)\}.$$

The values of ΔH for the six cases are given in table 2.

Cases 1, 3, and 6 clearly imply the result. The result follows for case 4 by results in Derman, as mentioned previously. For case 2, it was assumed that $\delta(\tau', 1) = 1$, and hence from the definition of H,

$$c^3 + \tau' z(\tau', 1, 1) + \beta \sum_{\tau''} P(\tau'' | \tau') v(\tau'', 1) \leq \tau' z(\tau', 1, 0) + \beta \sum_{\tau''} P(\tau'' | \tau') v(\tau'', 0),$$

which implies the result. Similarly, for case 5, the assumption that $\delta(\tau, 0) = 0$ implies the result. $\square$

Table 1

Case no.	$\delta(\tau, 0)$	$\delta(\tau, 1)$	$\delta(\tau', 0)$	$\delta(\tau', 1)$
1	0	0	0	0
2	0	0	0	1
3	0	0	1	1
4	0	1	0	1
5	0	1	1	1
6	1	1	1	1

Table 2

Case no.	ΔH
1	$\tau'[z(\tau', 0, 0) - z(\tau', 1, 0)] - \tau[z(\tau, 0, 0) - z(\tau, 1, 0)]$
2	$\tau'[z(\tau', 0, 0) - z(\tau', 1, 0)] - \tau[z(\tau, 0, 0) - z(\tau, 1, 0)] - c^3 + \beta \sum_{\tau''} P(\tau'' \mid \tau')[v(\tau'', 0) - v(\tau'', 1)]$
3	$c^1 + \tau'[z(\tau', 0, 1) - z(\tau', 1, 1)] - \tau[z(\tau, 0, 0) - z(\tau, 1, 0)]$
4	$\tau'[z(\tau', 0, 0) - z(\tau', 1, 1)] - \tau[z(\tau, 0, 0) - z(\tau, 1, 1)] + \beta \sum_{\tau''} [P(\tau'' \mid \tau') - P(\tau'' \mid \tau)][v(\tau'', 0) - v(\tau'', 1)]$
5	$c^1 + c^3 + \tau'[z(\tau', 0, 1) - z(\tau', 1, 1)] - \tau[z(\tau, 0, 0) - z(\tau, 1, 1)] + \beta \sum_{\tau''} P(\tau'' \mid \tau)[v(\tau'', 0) - v(\tau'', 1)]$
6	$\tau'[z(\tau', 0, 1) - z(\tau', 1, 1)] - \tau[z(\tau, 0, 1) - z(\tau, 1, 1)]$

We now present our main result.

THEOREM

Assume that v^* is the fixed point of H, and that assumptions (i), (ii), and (iii) in the lemma hold. Then:

(a) There is an optimal policy that is isotone in τ and $\bar{a}$.

(b) $v^*(\tau, 0) \geq v^*(\tau, 1) \geq 0$ for all τ.

(c) $v^*(\tau, 0) - v^*(\tau, 1)$ is isotone in τ.

(d) $v^*(\tau, \bar{a})$ is isotone in τ for each $\bar{a}$.

Proof of the theorem

Consider the sequence $\{v_n\}$, where $v_0 = 0$ and $v_{n+1} = Hv_n$. As stated earlier, $\lim_{n\to\infty} \| v_n - v^* \| = 0$. Note that v_0 satisfies conditions (b), (c), and (d); assume v_{n+1} also satisfies these conditions. The isotonicity of Hv in v implies that v_n satisfies (c). A standard induction argument implies that (b), (c), and (d) hold. Application of the lemma implies that (a) holds. □

There are (at least) two advantages that can result from a structural optimal policy. First, such policies are often more easily implemented. Second, they often suggest reduced numerical effort. In the current context, in calculating Hv_{n-1}, if $\delta(\tau, \bar{a}) = 1$, where δ is such that $H_\delta v_{n-1} = Hv_{n-1}$, then we know $\delta(\tau', \bar{a}') = 1$ for all $(\tau', \bar{a}')$ such that $\tau' \leq \tau$ and $\bar{a} \leq \bar{a}'$ and hence it is unnecessary to calculate

$$\tau' z(\tau', \bar{a}', 0) + \beta \sum_{\tau''} P(\tau'' \mid \tau') v_{n-1}(\tau'', 0).$$

The above results have assumed that $\{\tau_l(t), t = 0, 1, \ldots\}$ is conditionally independent of $\{\tau_{\bar{l}}(t), t = 0, 1, \ldots\}$, where for all t, $\tau(t) = \{\tau_l(t), \tau_{\bar{l}}(t)\}$. It seems reasonable to assume that

$$P(\tau' \mid \tau) = \prod_{i<j} P(\tau'_{ij}, \tau'_{ji} \mid \tau_{ij}, \tau_{ji})$$

and hence

$$P_l(\tau'_l \mid \tau_l) = \prod_{\substack{i<j \\ (i,j) \in \mathcal{P}_1}} P(\tau'_{ij}, \tau'_{ji} \mid \tau_{ij}, \tau_{ji}).$$

It also seems reasonable to assume that the process $\{\{\tau_{ij}(t), \tau_{ji}(t)\}, t = 0, 1, \ldots\}$ has an IFR, for all $(i, j) \in N \times N$. Unfortunately, it does not necessarily follow that $\{\tau_l(t), t = 0, 1, \ldots\}$ is IFR if the $\{\{\tau_{ij}(t), \tau_{ji}(t)\}, t = 0, 1, \ldots\}$ are IFR, for all $(i, j) \in \mathcal{P}_l$, and hence the optimal policy may not be isotone if the $\{\{\tau_{ij}(t), \tau_{ji}(t)\}, t = 0, 1, \ldots\}$ are IFR, for all $(i, j) \in \mathcal{P}_l$. See White and Schlussel [20] for a counterexample. However, it seems reasonable to expect that if the $\{\{\tau_{ij}(t), \tau_{ji}(t)\}, t = 0, 1, \ldots\}$ are IFR, for all $(i, j) \in \mathcal{P}_l$, then there exist high quality suboptimal designs within the class of isotone policies.

8. A numerical example

Consider the problem of determining when to activate and deactivate two links, one between nodes 4 and 7 (which we will denote as 4–7) and the other between nodes 3 and 9 (3–9), in the network given in fig. 2. Assume $z_{ij}(s, a) = z_{ij}(a_l)$ equals the minimum number of links to be traversed in going from node i to node j (i.e. the minimum hop count from i to j). This is case 1 from our previous discussion

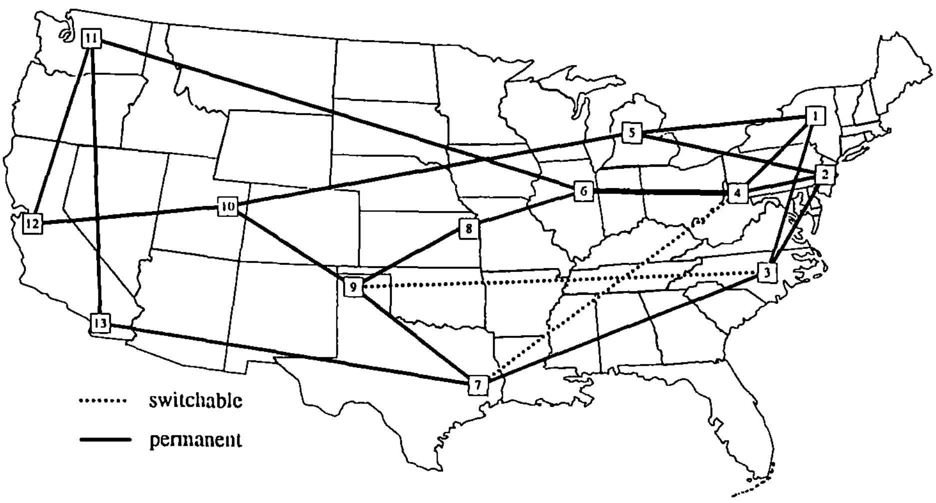

Fig. 2. Network for the numerical example.

Table 3

Hop counts: $\mathcal{P}_{4-7}$ and $\mathcal{P}_{3-9}$.

(a) Link 4–7 min-hop information

Node pair ($\mathcal{P}_{4-7}$)	Hops without link 4–7 ($\bar{a}_{4-7}=0$)	Hops with link 4–7 ($\bar{a}_{4-7}=1$)
(4,7)	3	1
(4,9)	3	2
(4,13)	3	2
(6,7)	3	2

(b) Link 3–9 min-hop information

Node pair ($\mathcal{P}_{3-9}$)	Hops without link 3–9 ($\bar{a}_{3-9}=0$)	Hops with link 3–9 ($\bar{a}_{3-9}=1$)
(1,9)	3	2
(2,9)	3	2
(3,8)	3	2
(3,9)	2	1
(3,10)	3	2
(3,12)	4	3

of the lemma and is used here as an example for clarity, since $\mathcal{P}_{4-7}$ and $\mathcal{P}_{3-9}$ can be identified by inspection. Table 3 indicates that four node pairs are affected by link 4–7 and that six node pairs are affected by link 3–9. Note that $\mathcal{P}_{4-7} \cap \mathcal{P}_{3-9} = \varnothing$.

We assume that for each $(i, j) \in \mathcal{P}_l$, $l \in \{4-7, 3-9\}$, $\{\tau_{ij}, \tau_{ji}\}$ can be in one of four states and has the IFR transition matrix:

$$\begin{bmatrix} 0.40 & 0.30 & 0.20 & 0.10 \\ 0.25 & 0.35 & 0.25 & 0.15 \\ 0.10 & 0.25 & 0.40 & 0.25 \\ 0.05 & 0.20 & 0.35 & 0.40 \end{bmatrix}.$$

Other parameter values are: $w = 0.5$, $\beta = 0.9$, $c^1_{4-7} = c^1_{3-9} = 5$ units of cost, and $c^2_{4-7} = c^2_{3-9} = 3$ units of cost. We assume that the amount of traffic between any pair (i, j) is symmetric at 10, 20, 40, and 70 packets per second when $\{\tau_{ji}, \tau_{ji}\}$ is in state 1, 2, 3, and 4, respectively. We also assume that the cost of delay per packet-hop (parameter b) equals 0.025.

The MDP decomposition for link 4–7 (3–9) required 0.25 seconds (38.55 seconds) of CPU time per iteration of the standard successive approximations algorithm on a Sun SPARC station™ 2GX for a total of 3.28 seconds (578.23 seconds). Thirteen iterations (fifteen iterations) were needed for the link 4–7 (3–9) problem to satisfy the stopping criterion $\max\{|v_{n+1}(s) - v_n(s)| : s \in S\} \leq 0.5$. All state transition matrices were calculated from the individual $\{\tau_{ij}, \tau_{ji}\}$ transition matrices, for each state, action, and iteration. The above reported CPU times could have been reduced significantly had these matrices been precomputed. A variety of numerical procedures (see White and White [21]) could have been used to further reduce CPU time and the number of iterations until convergence. We have indicated earlier that problems like this numerical example do not satisfy the IFR assumption in the theorem. It is worthy to note, however, that the policies that resulted from the two problems are isotone.

The state for each MDP decomposition is $s_l(t) = \{(\tau_{ij}(t), (i, j) \in \mathcal{P}_l); a(t-1)\}$, so there are $(4^4)(2) = 512$ possible states for the link 4–7 problem and $(4^6)(2) = 8{,}192$ possible states for the link 3–9 problem. The policies resulting from the successive approximation solution of the two problems simply map each possible state into one of the two possible actions $a(t) = 0$ or 1. Table 4 gives a portion of the optimal policy for the link 4–7 problem. For clarity in illustration, the indices of the traffic levels are used rather than the numerical levels themselves, i.e. level 1 corresponds with 10 packets per second, 2 with 20, 3 with 40, and 4 with 70. Also, action 1 corresponds with link 4–7 deactivated (off) and action 2 with activated (on). Table 4(a) demonstrates the isotonicity of the policy in traffic and table 4(b) demonstrates its isotonicity in $\bar{a}_{4-7}$. In both cases, the traffic-$\bar{a}_{4-7}$ combination that first triggered activating link 4–7 is delineated by double thin lines.

Table 4

Partial policy for the link 4–7 problem.

(a) Illustration of policy isotonicity in traffic

State {Traffic level $\{\tau_{ji}, \tau_{ji}\}$ for $(i,j) \in \mathcal{P}_{4-7}$;				$\bar{a}_{4-7}\}$	Action a_{4-7}
Pair (4,7)	Pair (4,9)	Pair (4,13)	Pair (6,7)		
1	3	4	4	1	1
1	4	3	4	1	1
1	4	4	3	1	1
1	4	4	4	1	2
2	4	4	4	1	2

(b) Illustration of policy isotonicity in $\bar{a}_{4-7}$

State {Traffic level $\{\tau_{ji}, \tau_{ji}\}$ for $(i,j) \in \mathcal{P}_{4-7}$;				$\bar{a}_{4-7}\}$	Action a_{4-7}
Pair (4,7)	Pair (4,9)	Pair (4,13)	Pair (6,7)		
1	1	1	3	2	1
1	1	1	4	1	1
1	1	1	4	2	2
2	1	1	4	2	2
1	2	1	4	2	2
1	1	2	4	2	2

9. Conclusions

We have presented an MDP model of the DTP, a decomposition of this model, and structural results for the decomposition. The decomposition and structural results enhance the potential tractability of procedures for determining optimal and good suboptimal policies for activating and deactivating links as a function of current traffic in the network and past control decisions. A numerical example illustrated these procedures.

Several assumptions were made in order to obtain the (optimal) results presented in this paper. These assumptions were:

1. $z_{ij}(s, a) = z_{ij}(\tau_{ij}, a_l)$ for $(i,j) \in \mathcal{P}_l$,
2. $\mathcal{P}_l \cap \mathcal{P}_k = \varnothing$ if $l \neq k$,

3. the IFR assumption,
4. the dependence of the cost structure on the current traffic matrix and the current and most recent control decisions.
5. $A = \{0, 1\}^L$.

We (implicitly) assumed that the MDP decomposition was tractable. It was also assumed that each link was either activated or deactivated; but, if $|A| \geq 2^L$, then the opitmal dynamic allocation of bandwidth could be studied. Removal of any or all of these assumptions provides a direction for future research. Use of numerical procedures to reduce CPU time to convergence is also a topic for future research.

References

[1] E. Altman and P. Nain, Optimal control of the *M/G/*1 queue with repeated vacations of the server, IEEE Trans. Auto. Control 38(1993)1766–1775.
[2] D. Bertsekas, *Dynamic Programming: Deterministic and Stochastic Models* (Prentice–Hall, Englewood Cliffs, 1987).
[3] R.M. Bournas, F.J. Butler and D. Teneketzis, Optimal flow control allocation policies in communications networks with multiple message classes, IEE Trans. Auto. Control 38(1993) 390–403.
[4] R.M. Bournas, F.J. Butler and D. Teneketzis, Properties of optimal hop-by-hop allocation policies in networks with multiple transmitters and linear equal holding costs, IEEE Trans. Auto. Control 36(1991)1450–1463.
[5] C. Derman, On optimal replacement rules when changes of state are Markovian, in: *Mathematical Optimization Techniques*, ed. R. Bellman (University of California Press, Berkeley, CA, 1963).
[6] R.-H. Hwang, Routing in high-speed networks, Ph.D. Dissertation, Department of Computer Science, University of Massachusetts, Amherst (1993).
[7] L. Kleinrock, *Queueing Systems Volume II: Computer Applications* (Wiley, 1976).
[8] L. LeBlanc and J.H. Harder, Optimization models for design and routing in wide-area data-communication networks, Stud. Locational Anal. 6(1994)3–18.
[9] L. LeBlanc, Design and operation of packet-switched networks with uncertain message requirements, IEEE Trans. Commun. 38(1990)1223–1230.
[10] L. LeBlanc and R. Simmons, Continuous models for capacity design of large packet-switched telecommunications networks, ORSA J. Comp. 1(1989)271–286.
[11] A.M. Makowski and A. Shwartz, On constrained optimization of the Klimov network and related Markov decision processes, IEEE Trans. Auto. Control 38(1993)354–359.
[12] R.L. Moose, Modeling networks with dynamic topologies, ORSA J. Comp. 1(1988)223–231.
[13] R.L. Moose and R.E. Nance, Link models for networks with dynamic topologies, Technical Report SRC 87-006, Systems Research Center, Virginia Polyrechnic Institute and State University (1987).
[14] R.E. Nance and R.L. Moose, Link capacity assignment in dynamic hierarchical networks, Comp. Networks and ISDN Syst. 15(1988)189–202.
[15] M.D. Noakes, J.B. Cain, J.W. Nieto and E.L. Althouse, An adaptive link assignment algorithm for dynamically changing topologies, IEEE Trans. Commun. 41(1993)694–706.
[16] K.W. Ross and D.H.K. Tsang, Optimal circuit access policies in an ISDN environment: A Markov decision approach, IEEE Trans. Commun.37(1989)934–939.
[17] G.N. Rouskas and M.H. Ammar, Dynamic reconfiguration in multihop WDM networks, *Proc. Photonics*, 1993.

[18] J. Shor and T.G. Robertazzi, Traffic sensitive algorithms and performance measures for the generation of self-organizing radio network schedules, IEEE Trans. Commun. 41(1993)16–21.
[19] C.C. White and K. Schlussel, Suboptimal design for large-scale multi-module systems, Oper. Res. 29(1981)865–875.
[20] C.C. White and K. Schlussel, Optimal replacement policies for completely observed multi-component systems, *Proc. Southeastcon*, Williamsburg, VA, 1977.
[21] C.C. White and D.J. White, Markov decision processes, Euro. J. Oper. Res. 39(1989)1–16.
[22] I.S. Gopal and T.E. Stern, Optimal call blocking policies in an integrated services environment, *Conf. Inform. Sci. Syst.*, Johns Hopkins University, 1983, pp. 383–388.