Demand Assigned Capacity Management (DACM) in IP Over Optical (IPO) Networks

Daniel O. Awduche, *Senior Member, IEEE*, and Bijan Jabbari, *Senior Member, IEEE*

***Abstract*—The demand assigned capacity management (DACM) problem in IP over optical (IPO) network aims at devising efficient bandwidth replenishment schedules from the optical domain conditioned upon traffic evolution processes in the IP domain. A replenishment schedule specifies the location, sizing, and sequencing of link capacity expansions to support the growth of Internet traffic demand in the IP network subject to economic considerations. A major distinction in the approach presented in this paper is the focus of attention on the economics of "excess bandwidth" in the IP domain, which can be viewed as an inventory system that is endowed with fixed and variable costs and depletes with increase in IP traffic demand requiring replenishment from the optical domain. We develop mathematical models to address the DACM problem in IPO networks based on a class of inventory management replenishment methods. We apply the technique to IPO networks that implement capacity adaptive routing in the IP domain and networks without capacity adaptive routing. We analyze the performance characteristics under both scenarios, in terms of minimizing cumulative replenishment cost over an interval of time. For the non-capacity adaptive routing scenario, we consider a shortest path approach in the IP domain, specifically OSPF. For the capacity adaptive scenario, we use an online constraint-based routing scheme. This study represents an application of integrated traffic engineering which concerns collaborative decision making targeted towards network performance improvement that takes into consideration traffic demands, control capabilities, and network assets at different levels in the network hierarchy.**

***Index Terms*—ASON, bandwidth replenishment, capacity management, demand assigned capacity management, GMPLS, integrated traffic engineering, inventory management, IP over optical networks, IPO, MPLS, network performance optimization, traffic engineering.**

I. Introduction

IP over Optical (IPO) networking is the concept of transporting IP traffic over a switched optical infrastructure without superfluous intervening layers (see, e.g., [1]). The underlying optical network may utilize an IP-based control plane technology, such as GMPLS, for connection management [2], [3]. Discussions of the computational techniques associated with optical channel routing and wavelength assignment (RWA) in relation to capacity activation in the optical domain is well documented in the literature [4]–[6]. A considerable amount of prior studies is devoted to the associated survivability problems [7]–[10]. Some recent papers attempt to jointly perform path selection in the IP domain and wavelength routing in the optical domain using an integrated algorithmic approach [11]. Yet, despite all this attention, the linkage between the rate of growth of Internet traffic demand in the IP domain and capacity replenishment strategies from the optical domain is an issue that has not been given sufficient analysis in prior research on IPO networks. This paper provides a conceptualization of this problem premised on the concept of demand assigned capacity management which has a theoretical basis in industrial production planning and inventory control.

The demand assigned capacity management (DACM) problem in IPO networks concerns efficient management of network bandwidth in the IP and optical domains. The capacity expansion component of DACM aims at replenishing link bandwidth in the IP domain conditioned on Internet traffic demand evolution processes, leveraging capacity activation technologies and policies in the underlying optical network. DACM comes under the subject of integrated traffic engineering which advocates a broader view of network performance optimization relative to traditional traffic engineering [12]. The definition of traditional Internet traffic engineering centers around the performance optimization and performance evaluation of operational IP networks [13]–[16]. In practice, it boils down to mapping IP traffic onto an existing infrastructure in the most effective way. Integrated traffic engineering, on the other hand, concerns collaborative decision making targeted towards network performance improvement that takes into consideration traffic demands, control capabilities, and network assets at different levels in the network hierarchy. In practice, it entails determining the current state of traffic in the network and how it is likely to evolve, and devising improvements to the existing infrastructure, at various levels in the network value chain echelon, to support the traffic. Such improvements can be implemented by modifying the behavioral characteristics of the network (for example, using traditional traffic engineering techniques such as routing optimization), or by augmenting network capacity.

There are two general techniques to IP network capacity expansion, which can be applied independently or in combination. First, the capacity of existing links in the network can be increased. Second, the topology of the network can be restructured by adding new links and possibly deleting some existing links. The later is the well known topology design problem whose optimal solution applies mixed integer multicommodity flow techniques (see, e.g., [17]).

This paper studies the capacity expansion aspects of the DACM problem in IPO networks. The motivation for this study is the need to improve network economics by reflecting and leveraging the enhanced flexibility and agility offered by

Manuscript received July 8, 2003; revised October 31, 2004.

D. O. Awduche is with Verizon Business, Ashburn, VA 20147 USA (e-mail: awduche@ awduche.com).

B. Jabbari is with the Department of Electrical and Computer Engineering, George Mason University, Fairfax, VA 22030 USA (e-mail: bjabbari@gmu.edu).

Digital Object Identifier 10.1109/TNET.2006.879700

technological advances in switched optical networks [18], [2], [19]–[21]. We first present a conceptual framework for DACM in IPO networks. We then describe a class of inventory-based replenishment methods to address the capacity expansion problem. These methods adopt an inventory management oriented view of capacity replenishment in IPO networks.

A basic tenet underlying our theory of DACM in IPO networks is the orientation of attention on the cost of excess bandwidth in the IP domain. This viewpoint implies the DACM problem can be transformed into an inventory control problem, where excess bandwidth in the IP domain is the inventory of interest. In this environment, traffic in the IP domain is growing according to a specified process over a given interval of time resulting in depletion of excess bandwidth. The inventory management approach is essentially the issue of reducing the cost of procuring and holding excess network capacity in the IP domain relative to the cost of congestion caused by insufficient bandwidth to accommodate Internet traffic as the demand grows over time. The key decision problems center on the location, sizing, sequencing, and economics of capacity procurement activities that occur between the IP domain and underlying optical network. The objective is to minimize the long term cost of maintaining and replenishing excess capacity in the IP domain subject to prescribed service level requirements. The decision variables depend very much on the evolutionary characteristics of IP traffic demand, the attributes of the existing technologies underpinning the IPO network, the need to maintain proper service quality, and pertinent economic considerations which center on the inter-play between fixed and variable costs. In general, the system can be parameterized by: 1) the characteristics of the rate of growth of traffic in the IP domain; 2) the characteristics of the lead time to replenish capacity from the optical domain; 3) a set of network attributes; and 4) a set of economic variables. We analyze the performance characteristics of the inventory-based replenishment methods under capacity adaptive and non-capacity adaptive routing strategies in the IP domain.

The remainder of this paper is organized as follows. Section II describes the network context and presents an architecture and a process model for inventory-based DACM in IPO networks. Section III provides a mathematical formulation of the problem and considers the inventory-based DACM problem with deterministic and constant IP traffic growth rates. Subsequently, we consider the inventory-based DACM scheme with dynamic traffic growth rates in Section IV. This is followed by numerical results in Section V. Finally, Section VI contains the concluding remarks.

II. Conceptual Framework for DACM in IPO Networks

The network context considered in this paper consists of an IP infrastructure which is overlaid over an optical network. In this arrangement, routers within the IP domain are interconnected by optical channels provided by the underlying optical transport network. This network model is shown in Fig. 1.

The optical network consists of optical switches interconnected among themselves in a mesh or ring topology using DWDM technology. The optical infrastructure may implement

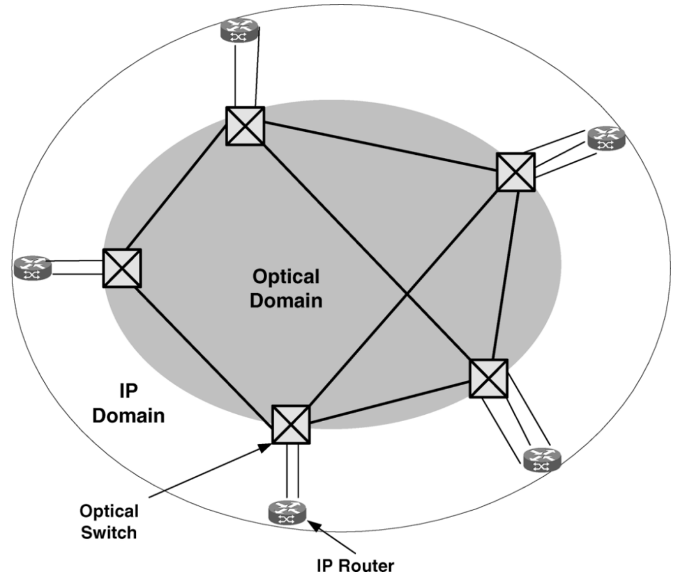

Fig. 1. Physical configuration of IPO networks.

a control plane (e.g., using GMPLS [2], [19]) for capacity activation and dynamic connection management. The optical domain may apply an appropriate computational procedure for optical channel routing and wavelength assignment (see, e.g., [4], [6], and the references therein). We do not make strong assumptions about the control interconnection model between the IP routers and optical switches, except to say that some mechanism exists to communicate capacity replenishment requests from the IP domain to the optical domain. This mechanism may exist within the network itself (e.g., through a UNI, through GMPLS signaling, or through an NNI), or outside the network (e.g., through interactions between Network Management Systems, or through an operations support system, or through a WEB-based customer network management portal).

There are three primary modes of capacity replenishment in the IPO network from the perspective of the IP domain.

The first approach is to increase the capacity of an existing link between two routers. For instance, in the case of SONET, it may be feasible to upgrade an existing link from OC-48 to OC-192 capacity if the underlying technologies and facilities support this type of transaction. It may also be possible to fractionally increment the link capacity in smaller granularities using virtual concatenation in conjunction with a link capacity adjustment scheme (e.g., the ITU-T LCAS [22]).

The second approach to capacity expansion in the IP domain is to establish new optical connections which share the same endpoints as an existing link in the IP domain and use the concept of "bundling" to aggregate the circuits. Link bundling is an abstraction that allows the logical consolidation of multiple parallel links between two routers into one virtual link whose capacity is the sum of the capacities of the component links (see, e.g., [2], [19]).

The third approach to capacity replenishment is to add a new link between two routers that do not already have a link between them. This third approach requires a redesign of the IP network topology. In this paper, we do not address the topology design problem. Instead, we shall be concerned only with the first two modes of capacity expansion mentioned above.

A. Categories of DACM Systems

DACM systems in IPO networks can broadly be classified into three general categories: 1) push systems; 2) pull systems; and 3) synchronous systems. The basic idea behind push DACM systems is the philosophy that asserts that first, excess capacity is installed in the IP network, and then effort is made to stimulate demand. In such systems, procured bandwidth generally exceeds near term traffic requirements. Hence, the network tends to be underutilized, incidents of congestion are rare, and service levels are very high. The disadvantage of push DACM methods is that capacity is procured which may never be used, resulting in inefficient allocation of capital. On the other hand, in pull DACM systems, an attempt is made to match capacity to demand as closely as possible, so that bandwidth is procured to serve existing demand or demand that will become available in the short term. Generally, pull systems are adaptive and cost efficient, but can result in severe congestion if the traffic fluctuates over time.

A synchronous DACM system is a hybrid between a pull system and a push system. The basic concept is to use a mixture of empirical measurements and objective traffic forecasts to make replenishment decisions. In this way, the traffic demand is matched with capacity as closely as possible and adequate service levels are maintained. Synchronous systems may utilize *hedge bandwidth* in some instances to mitigate against uncertainties in capacity supply and traffic demand.

B. Environmental Characteristics

In formulating the DACM problem, different types of assumptions can be made about the trends governing the rate of growth of Internet traffic as a function of time within the IPO network. There are four prominent types of traffic growth trend characteristics that can be envisaged.

1) Deterministic and stationary IP traffic growth rate. Under this assumption, the traffic grows at a constant known rate over time.
2) Deterministic and dynamic IP traffic growth rate. Under this assumption, the traffic grows at a deterministic rate which changes with time.
3) Stochastic IP traffic growth rate with known distribution. Under this assumption, the traffic grows at a stochastic rate with known distribution. That is, the traffic growth rate can be characterized by a known statistical distribution function.
4) Stochastic IP traffic growth rate with unknown distribution. Under this assumption, the traffic grows at a stochastic rate with unknown distribution.

Similarly, a number of assumptions can be made about the lead time for capacity augmentation from the optical domain. Generally, the lead time is a function of the degree of automation within the optical network and the cost of preinstalling cards on optical network elements. The major types of assumptions that can be made about optical connection lead time are as follows:

1) instantaneous capacity activation in which the capacity activation lead time is assumed to be zero;
2) deterministic lead time with constant value, where the lead time to activate additional capacity from the optical domain is a positive constant;
3) stochastic lead time with known distribution, where the lead time to activate capacity is governed by some stochastic process with known distribution;
4) stochastic lead time with unknown distribution, where the lead time to activate additional capacity is governed by a stochastic process with unknown distribution.

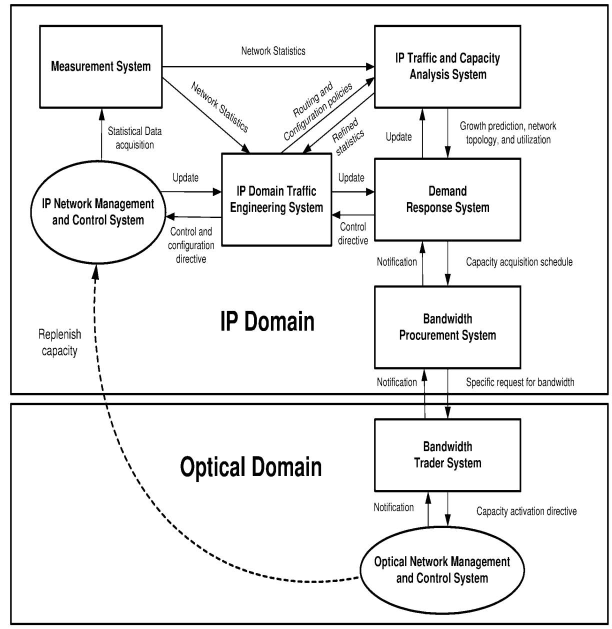

Fig. 2. Architecture for DACM in IPO networks.

C. Architecture for DACM in IPO Networks

A high level architecture for a synchronous DACM system is shown in Fig. 2. The high level architecture provides structural and functional views of pertinent system components. The architecture consists of the following major components: 1) IP domain traffic engineering system; 2) measurement system; 3) traffic and capacity analysis system; 4) demand response system; 5) bandwidth procurement system; 6) optical bandwidth trader system; and 7) optical network management and control system.

The IP domain traffic engineering system consists of the traffic engineering capabilities that exist within the IP domain. This system performs the traditional Internet traffic engineering functions, namely, performance optimization of the IP domain by mapping IP traffic onto the existing IP connectivity infrastructure in the most effective way. Some components of this system may exist inside the IP network, while others may exist outside the network. The IP domain traffic engineering system may utilize MPLS or some other appropriate technology to perform its functions.

The measurement system acquires traffic statistics from the IP layer and dispatches the data to other systems for storage and analysis. The measurement system may perform filtration, compression, and data reduction on the raw statistics.

The traffic analysis system utilizes empirical statistics derived from the measurement system and pertinent external influencing factors to perform various network analysis and forecasting functions. An important function performed by this system is the prediction of the rate of growth of IP traffic between different source-destination nodes as a function of time. The forecasting process may apply trend analysis to historical traffic data stored in a repository.

The demand response system engages in a decision process within the IP domain to determine if, where, when, and how much additional capacity is required in the IP network. First, it determines whether the existing network can accommodate the demand in conjunction with the traffic engineering system. If the demand cannot be accommodated, it initiates a process to acquire additional bandwidth from the optical domain. This involves determining the endpoints, bandwidth requirements, and timing of optical capacity replenishments to accommodate existing and predicted IP traffic demand, subject to economic considerations and network policies.

The bandwidth procurement system within the IP domain is responsible for initiating capacity order requests from the optical network based on input from the demand response system.

The bandwidth trader system resides in the optical domain and is responsible for accepting capacity order requests from clients in the IP domain. Each instance of the bandwidth trader system may be associated with multiple clients. The bandwidth trader system performs preliminary admission control functions to determine whether a capacity order request can be successfully provisioned.

Finally, the optical network management and control system is responsible for optical channel layer bandwidth management, dynamic provisioning of optical capacity, and associated connection management functions.

A process model corresponding to this architecture is shown in Fig. 3. This process model consists of activities that occur within the IP/MPLS domain and activities that occur within the optical domain.

The high level operation of the demand replenishment process is as follows. First, within the IP/MPLS domain, the existing and predicted IP traffic is analyzed to ascertain the demand characteristics and the accompanying bandwidth requirements.

Next, a decision process is invoked in the IP domain (involving the demand response and traffic engineering systems) to determine whether the traffic demand can be successfully mapped onto the existing IP connectivity infrastructure. This set of actions can be performed in one of two ways: 1) using *stable non-capacity adaptive* routing and 2) using *dynamic capacity adaptive routing*.

In DACM systems employing stable non-capacity adaptive routing in the IP domain, traffic is mapped onto the underlying connectivity infrastructure at the commencement of the process using a stable routing method. Once the initial routing is done, the routes are pinned, and dynamic rerouting does not occur except under exceptional conditions. Once routes are selected and pinned, congestion levels on links are monitored continuously. Replenishment capacity is added to existing links in the IP network based on monitored and predicted congestion levels.

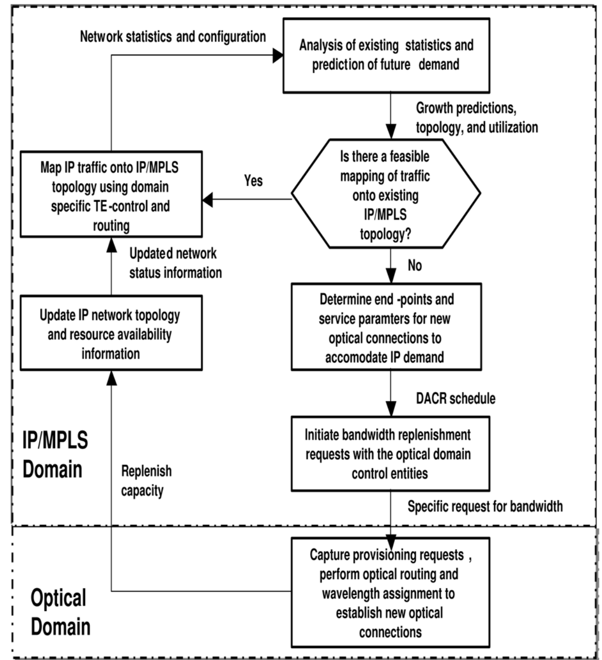

Fig. 3. Bandwidth replenishment process model.

On the other hand, DACM systems employing dynamic capacity adaptive routing in the IP domain have the ability to reroute traffic to take advantage of alternate less congested routes. In this case, the decision process essentially involves ascertaining whether alternate routes exist within the IP network with sufficient capacity to handle the demand. If alternate paths with sufficient capacity exist, then some of the demand is routed through the alternate paths. If this condition is satisfied, then the IP traffic is mapped onto the existing infrastructure, and there is no need to activate additional capacity from the optical domain. However, if the demand cannot be accommodated within the existing IP infrastructure, then action is initiated to acquire additional bandwidth resources from the optical infrastructure.

During the capacity acquisition phase, a bandwidth replenishment request is communicated to the optical domain to activate the required capacity along the desired set of links.

D. *Cost Considerations for DACM*

The cost parameters associated with DACM in IPO networks include: 1) fixed replenishment setup costs; 2) capacity dependent variable initial cost; 3) bandwidth carrying costs; and 4) bandwidth stockout costs. Fixed replenishment setup cost is the fixed cost of ordering optical capacity (e.g., a group of optical channels), which is independent of the capacity ordered. This includes administrative costs and other overhead associated with each group of bandwidth replenishment transaction. Capacity dependent variable initial cost is the variable cost of procuring capacity from the optical domain which is dependent on the amount of capacity procured. Bandwidth carrying cost (also called holding cost or excess circuit cost) is the recurrent cost for excess capacity within the IP domain per unit time. This cost may depend on the geographical area in which the optical capacity is situated, the circuit distance, and the circuit connection attributes (e.g., priority, survivability, etc). Bandwidth stockout cost is the penalty associated with not having sufficient bandwidth in the IP domain to accommodate offered or predicted Internet traffic. Several components may be associated with this cost. One component is the possibility of congestion which degrades service quality resulting in loss of customer goodwill. Another component is the possibility of foregone revenue because capacity does not exist to accommodate new service requests. Finally, there is the possibility of defection of customers to competing providers due to poor service quality.

III. Mathematical Models for Inventory-Based DACM

In a network context, an inventory-based DACM system can be characterized by the following tuple:

$$\Omega = \langle \mathcal{G}, \mathcal{T}, \chi, \mathcal{C}, \mathcal{U} \rangle \tag{1}$$

where Ω denotes the universe of discourse for the DACM network problem. The parameter $\mathcal{G} = \langle V, E \rangle$ is the graph of the network topology consisting of a set of nodes (V) and a set of edges ($E = \{e_{ij} : i, j = 1, \ldots, n\}$), where n is the cardinality of V. Each link has a number of attributes including an initial capacity and a link metric. The parameter $\mathcal{T} = \langle \Lambda, \mathcal{D} \rangle$ denotes the traffic oriented attributes characterizing the demand imposed on the network. In this display, Λ is the initial traffic matrix, so that $\lambda_{uv} \in \Lambda$ is the effective bandwidth of the initial traffic originating from node u and terminating on node v. More generally, $\lambda_{uv}(t)$ is the effective bandwidth of the traffic between node u and v at time t. The component $\mathcal{D}$ is the set of growth rates associated with Λ, such that $d_{uv}(t) \in \mathcal{D}$ is the rate of growth of $\lambda_{uv} \in \Lambda$ at time t (i.e., the rate of growth of traffic between nodes u and v at time t). It should be noted that we are interested in the *rate of growth of traffic*. Otherwise, if traffic is not growing as a function of time, then there is no need to replenish network capacity over time. The parameter $d_{uv}(t)$ also corresponds to the rate of depletion of available capacity along the path traversed by the traffic in the IP domain. The reader should note that $d_{uv}(t)$ represents the trend component of a time series decomposition of demand evolution between nodes u and v. We assume $d_{vu}(t)$ is a non-decreasing function of time, namely $d_{vu}(t+1) \geq d_{vu}(t)$. Initially, we shall assume that $d_{vu}(t)$ is constant (that is, IP traffic grows at a constant rate) over the time horizon of interest, so that $d_{vu}(t) = d_{vu}$, for some constant d_{vu}. This assumption will be relaxed in a subsequent formulation of the problem.

The parameter $\mathcal{C} = \langle \mathcal{A}, \alpha, \beta \rangle$ captures all the cost components of the DACM network problem. These cost parameters can be viewed as link attributes and pertain to the time horizon of interest. Here, $A_{ij} \in \mathcal{A}$ is the fixed replenishment cost for each capacity procurement event across link e_{ij} from the optical domain irrespective of the amount of capacity procured. The cost component $\alpha_{ij} \in \alpha$ is the initial bandwidth-dependent unit cost of procuring bandwidth from the optical domain across link e_{ij}. Thus, α_{ij} is an initial cost which depends on the amount of capacity procured from the optical domain across e_{ij}, that is, if Q_{ij} units of bandwidth are procured, then the initial cost contributed by the α_{ij} parameter is the product of Q_{ij} and α_{ij}. Let $C^1_{ij}(Q_{ij})$ be the total initial cost of procuring Q_{ij} units of bandwidth from the optical domain across link e_{ij}. This represents the contributions of A_{ij} and α_{ij} and is given by

$$C^1_{ij}(Q_{ij}) = \begin{cases} A_{ij} + Q_{ij}\alpha_{ij}, & \text{if } Q_{ij} > 0 \\ 0, & \text{if } Q_{ij} = 0. \end{cases} \tag{2}$$

Lastly, the cost parameter $\gamma_{ij} \in \gamma$ is the recurrent holding cost per unit time per unit of excess bandwidth across link e_{ij} in the IP domain (i.e., available link capacity that is not used to carry IP traffic). We shall assume that $\gamma_{ij} > 0$, in recognition of the fact that costs accrue when excess capacity exists in production networks.

Let $C_{\alpha ij}$ be the normalized total initial cost of procuring bandwidth from the optical domain per unit time over the time horizon of interest. This represents the initial costs per replenishment event on link e_{ij}. Let $C_{\gamma ij}$ be the total recurrent holding cost for excess bandwidth per unit time. Thus, $C_{\gamma ij}$ is the average holding cost for available bandwidth during a replenishment cycle across link e_{ij}. Let C_{Qij} be the total cost (fixed and variable holding costs) of procuring Q_{ij} units of bandwidth from the optical domain.

The parameter χ captures the environmental attributes surrounding a particular instance of the problem. These environmental attributes include the service level specifications that define the threshold policy for capacity replenishment and the routing system that establishes a mapping between the traffic matrix and the set of links in the network.

Finally, the parameter $\mathcal{U}$ from (1) represents the objective function which has to do with minimizing replenishing costs in the long run. That is, the objective is to derive link replenishment capacities, $Q^*_{ij}(t)$, and associated replenishment times that minimize the quantity $\sum_i \sum_j C_{Qij}$.

Let $I_{ijuv}(t)$ be a link-traffic indicator function induced by the routing system in the IP network and defined as follows:

$$I_{ijuv}(t) = \begin{cases} 1, & \text{if traffic from } u \text{ to } v \text{ at time } t \text{ traverse link } e_{ij} \\ 0, & \text{otherwise.} \end{cases} \tag{3}$$

Let $D_{ij}(t)$ be the aggregate rate of growth of all traffic traversing link e_{ij} at time t. It is clear that the following holds:

$$D_{ij}(t) = \sum_u \sum_v d_{uv}(t) I_{ijuv}. \tag{4}$$

Similarly, let $\Theta_{ij}(t)$ be the aggregate effective bandwidth of all traffic traversing link e_{ij} at time t. Then, the following holds:

$$\Theta_{ij}(t) = \sum_u \sum_v \lambda_{uv}(t) I_{ijuv}. \tag{5}$$

Let $B_{ij}(t)$ be the available bandwidth across link e_{ij} at time t. In practice, $B_{ij}(t)$ is the residual capacity across the link which represents the available bandwidth for carrying incremental traffic across it. It can be seen that $B_{ij}(t)$ decreases at a rate of $D_{ij}(t)$ units of bandwidth per unit time, where $D_{ij}(t)$ is the aggregate rate of growth of all traffic traversing link e_{ij} as noted earlier. If $Z_{ij}(t)$ is the capacity of link e_{ij} at time t, then certainly

$$B_{ij}(t) = Z_{ij}(t) - \Theta_{ij}(t). \tag{6}$$

Let $Q_{ij}(t)$ be the amount of capacity ordered across link e_{ij} at a particular replenishment point to offset depletion of residual capacity across the link. Generally, we can suppose $Q_{ij}(t)$ will be multiples of a base unit of capacity, for example OC-12. Let R be the minimum acceptable threshold value of available bandwidth across a link in the IP network when bandwidth replenishment requests are issued to the optical domain. That is, a bandwidth replenishment decision is initiated when $B_{ij}(t) \leq R$. It should be noted that this is a necessary, but not sufficient condition for initiating bandwidth replenishment. For example, in DACM systems with capacity adaptive routing, it might be possible to reroute some traffic within the IP domain along alternate underutilized routes to balance the utilization. The parameter R represents an indication of the minimum service level (or the congestion level) within the IP domain prior to invoking bandwidth replenishment requests from the optical domain. In real operational contexts, capacity replenishment decisions are usually initiated when utilization exceeds 60%. Networks with very aggressive cost control measures may prefer to initiate replenishment requests when utilization exceeds 80%. The parameter R thus corresponds to the unutilized available bandwidth at the instant of making replenishment request which is equal to

$$R = (1 - \rho) * Z_{ij}(t) \tag{7}$$

where ρ is the maximum acceptable link utilization. Let L be the connection activation lead time, which is the amount of time it takes to activate additional capacity from the optical domain. This may correspond to the time required to instantiate a new optical connection or to augment the capacity of an existing link. Let D_{Lij} be the demand during a particular lead time across link e_{ij}. Thus, D_{Lij} represents the amount of available bandwidth consumed by increasing traffic within the IP domain during the lead time for optical connection activation.

A. DACM With Constant Traffic Growth Rate: Single Link Case

We first consider the DACM problem for a single link in isolation, in which the rate of growth of traffic is constant. We then show how the results can be extended to network environments with stable non-capacity adaptive routing and in network scenarios with dynamic capacity adaptive routing. Our interest here is the following: we want to determine how much replenishment capacity to procure from the optical domain to support IP traffic growth, and when to procure the capacity. The results derived in this subsection are analogous to concepts from economic order quantity of inventory theory. We shall make the following assumptions.

Assumption 3.1: For each u and v, the rate of growth of traffic, $d_{uv}(t)$, is deterministic and constant. This implies that $d_{uv}(t) = d_{uv}$ for some constant d_{uv}.

Assumption 3.2: The lead time L for capacity activation from the optical domain is zero.

Assumption 3.3: The DACM capacity replenishment process occurs over an infinite time frame.

For these particular set of assumptions, it can be seen that bandwidth replenishment decisions should be made precisely at those times when the available bandwidth is exactly R. Let T_{ij} be time between replenishments (i.e., the replenishment cycle) across link e_{ij}. Recall that D_{ij} is the aggregate growth rate of all traffic traversing link e_{ij}. It is clear the following stability condition must hold in order to prevent persistent congestion from occurring in the network:

$$\overline{\frac{Q_{ij}}{T_{ij}}} \geq \overline{D_{ij}}. \tag{8}$$

Equation (8) simply asserts that on the average, the rate of increase in capacity must equal or exceed the rate of increase in traffic demand. In what follows, the expressions can be simplified by using R as a baseline value of bandwidth prior to replenishment. So we can assume the available bandwidth is zero for purposes of replenishment when the available bandwidth is equal to R. Suppose we start the inventory-based DACM replenishment process with R units of bandwidth in excess of existing traffic demand and we subsequently use a replenishment quantity of Q_{ij} units of bandwidth. Then, the time to the next replenishment is given by $T_{ij} = Q_{ij}/D_{ij}$. Hence, the number of capacity replenishments per unit time is given by $1/T_{ij} = D_{ij}/Q_{ij}$.

Each replenishment of Q_{ij} units of bandwidth from the optical domain has an associated initial cost of $A_{ij} + \alpha_{ij}Q_{ij}$. Hence, the initial capacity replenishment costs per unit time is given by

$$C_{\alpha ij} = (A_{ij} + Q_{ij}\alpha)\frac{D_{ij}}{Q_{ij}} = A_{ij}\frac{D_{ij}}{Q_{ij}} + \alpha_{ij}D_{ij}. \tag{9}$$

The second term $(\alpha_{ij}D_{ij})$ in the above relation is independent of Q_{ij}, hence it plays no role in determining the optimum amount of bandwidth to procure per unit time from the optical domain.

Notice that if Q_{ij} is a constant amount of bandwidth procured at each replenishment point and if the demand growth rate is a constant, then the average available bandwidth $(\overline{B_{ij}}(t))$ level between capacity replenishments is simply $Q_{ij}/2$. The average recurrent cost $(C_{\gamma ij})$ of *holding* available bandwidth per unit time during the replenishment cycle is now given as follows:

$$C_{\gamma ij} = \frac{\gamma_{ij}Q_{ij}}{2}. \tag{10}$$

The total replenishment cost per unit time for Q_{ij} units of bandwidth is simply [26]

$$C_{Qij} = C_{\alpha ij} + C_{\gamma ij} = \left(A_{ij}\frac{D_{ij}}{Q_{ij}} + \alpha_{ij}D_{ij}\right) + \frac{\gamma_{ij}Q_{ij}}{2}. \tag{11}$$

It can be seen that this total replenishment cost consists of two components ($C_{\alpha ij}$ and $C_{\gamma ij}$). One component increases linearly with increasing values of acquired capacity Q_{ij}, while the other component decreases inversely with increasing values of Q_{ij}.

Fig. 4 shows the graphs of $C_{\alpha ij}$, $C_{\gamma ij}$, and C_{Qij} as a function of Q_{ij}. The graph was plotted with the following values: $A_{ij} = 6$, $\gamma_{ij} = 2$, $D_{ij} = 3$, and $\alpha_{ij} = 0$. It can be seen that the total cost (C_{Qij}) is minimized when $A_{ij}D_{ij}/Q_{ij} = \gamma_{ij}Q_{ij}/2$.

The optimum value of Q_{ij}, denoted Q^*_{ij}, occurs when $d(C_{Qij})/d(Q_{ij}) = 0$. Hence,

$$\frac{d(C_{Qij})}{d(Q_{ij})} = -A_{ij}\frac{D_{ij}}{Q^2_{ij}} + \frac{\gamma_{ij}}{2} = 0. \tag{12}$$

So that

$$Q^*_{ij} = \sqrt{\frac{2D_{ij}A_{ij}}{\gamma_{ij}}}. \tag{13}$$

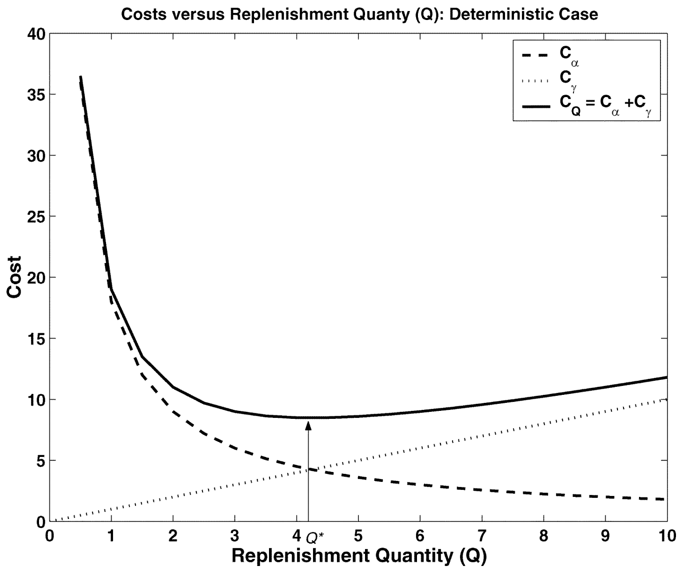

Fig. 4. Replenishment cost as a function of procured bandwidth.

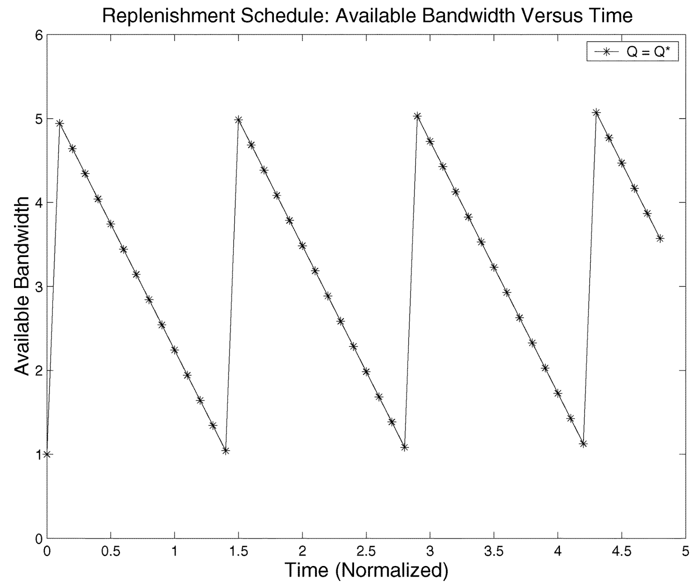

Fig. 5. Link replenishment schedule with constant traffic growth rate.

Equation (13) is interesting because it asserts that under a constant rate of traffic growth, the capacity replenishment required to *economically* support the growth of traffic is directly proportional to the square root of the rate of growth.

We observe that (13) is analogous to the expression for *Economic Order Quantity (EOQ)* commonly encountered in inventory theory under the assumption of constant deterministic demand (see, e.g., [24]).

In practice, the value of Q_{ij}^* will have to be scaled to appropriate integer units of bandwidth (e.g., units of OC-12 bandwidth) for operational applications in IPO networks. So, let B_s be the reference minimum unit of capacity (e.g., OC-12) that can be provisioned from the optical domain. Then the actual amount of bandwidth to be procured is given by $Q_{ij}^{'*}$, where $Q_{ij}^{'*}$ is scaled to the nearest number of units of B_s, so that

$$Q_{ij}^{'*} = \left\lfloor \left(\frac{Q_{ij}^*}{B_s} + 0.5 \right) \right\rfloor. \tag{14}$$

The following is an algorithm for bandwidth replenishment for a link where the rate of growth of bandwidth across the link is constant.

Inventory-Based Single Link DACM Algorithm with Constant Deterministic Link Traffic Growth Rate

Begin
 If $A_{ij} = 0$
  $Q_{ij}^* = D_{ij}$
 Elseif $A_{ij} > 0$ **And** $\gamma_{ij} > 0$
  $Q_{ij}^* = \sqrt{\frac{2D_{ij}A_{ij}}{\gamma_{ij}}}$
 Repeat
  Monitor $B_{ij}(t)$ *across link* (i,j)
  If $B_{ij}(t) \leq R$
   Procure Q_{ij}^* *bandwidth from optical domain*
 Until Stopped by Network Operator

The above algorithm does not take into account survivability considerations, which would mandate procuring additional redundant capacity from the optical domain along topologically diverse alternate routes to provide the ability to restore traffic in the event of network failures. Thus, depending on the degree of survivability required, it may be necessary to procure extra capacity from the optical domain in addition to $Q_{ij}^{'*}$.

Fig. 5 depicts an example of an inventory-based DACM replenishment schedule using the above algorithm for bandwidth replenishment with a constant and deterministic rate of increase in traffic demand. The parameters used in this graph are: $A_{ij} = 6$, $\gamma_{ij} = 2$, $\alpha_{ij} = 0$, $R_{ij} = 1$, and $D_{ij} = 3$. The value of Q_{ij}^* derived from the algorithm is 4.2426 and the corresponding value of cycle time T_{ij} is 1.4142, while the aggregate replenishment cost per cycle time is $C_{Qij} = 8.4853$.

In the previous example, we assumed that the lead time to provision capacity from the optical domain was zero. Here, we relax this assumption and consider the effects of constant positive lead-time (L) on the replenishment schedule. Suppose it takes L units of time to provision capacity from the optical domain following a bandwidth activation request. The total increment in demand during the lead time is given by D_L. To address this scenario, the replenishment policy should be modified so that a capacity replenishment requisition for $Q_{ij}*$ units of bandwidth is issued when $B_{ij}(t) \leq (R + D_L)$. The effect of lead time essentially leaves the sample path of the replenishment schedule unaltered, but shifts the capacity procurement times to the left of the time scale by D_L time units.

B. Sensitivity Analysis

We now illustrate the effects of variations in initial fixed costs by comparing the behavior of the optimal replenishment schedule ($Q_{ij} = Q_{ij}^*$ and $T_{ij} = Q_{ij}^*/D_{ij}$) with the behavior of a myopic schedule ($Q_{ij} = D_{ij}$ and $T_{ij} = 1$).

Fig. 6 shows the replenishment cost as a function of fixed initial cost (A_{ij}).

The schedule $Q_{ij} = D_{ij}$ is the situation where the capacity is replenished at the same rate as the rate of growth in demand. Fig. 6 shows that the cost of the inventory-based DACM scheme using the optimal replenishment capacity (Q_{ij}^*) exhibits lower

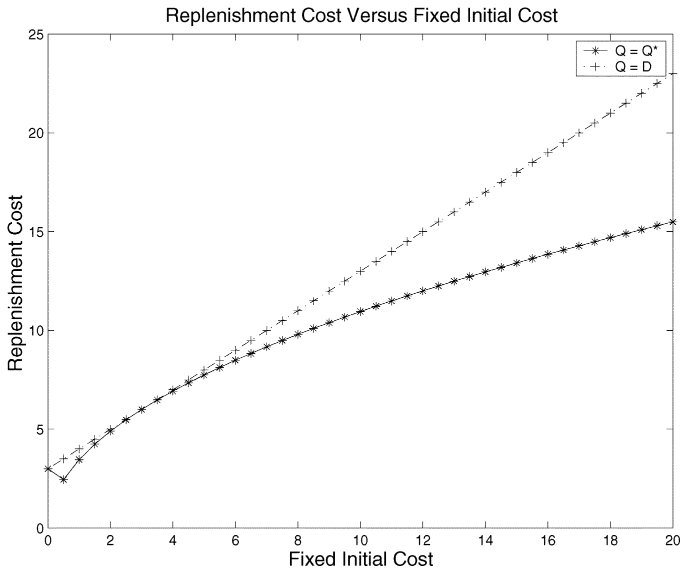

Fig. 6. Total replenishment cost as a function of fixed initial cost.

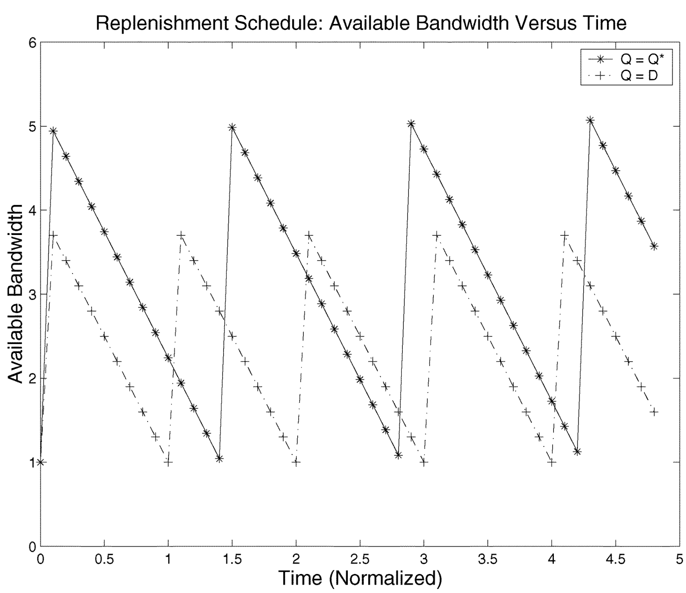

Fig. 7. Effects of varying procured capacity on replenishment schedule.

replenishment cost than the value obtained using D_{ij}. The difference in replenishment cost between the two schedules increases as the value of initial fixed cost increases.

Fig. 7 illustrates the effects of varying the replenishment size (Q_{ij}) on the replenishment schedule for the situation $Q_{ij} = Q^*_{ij}$ and $Q_{ij} = D_{ij}$.

In the following, we explore the sensitivity of the basic deterministic model to perturbations in the estimated traffic demand growth parameter D_{ij}, which centers around the robustness of the deterministic model to perturbations in estimates of the traffic demand growth rate. That is, if there is a change from D_{ij} to say D^1_{ij}, what is the effect on Q^*_{ij}, the optimal replenishment bandwidth? Let Qij^{1*} be the new optimal replenishment bandwidth corresponding to D^1_{ij}. Then, clearly, from (13) the following condition holds:

$$\frac{Q^{1*}_{ij}}{Q^*_{ij}} = \sqrt{\frac{D^1_{ij}}{D_{ij}}}. \tag{15}$$

This means the relative change in the amount of replenishment bandwidth required to support the new demand changes only as a function of the square root of the relative change in traffic growth rate. Thus, the optimal replenishment bandwidth based on the model of this subsection is quite robust to perturbations in traffic growth rate.

The sensitivity analysis can also be approached using the concept of *percentage cost penalty*. Let p denote a measure of the deviation of Q^{1*}_{ij} from Q^*_{ij}, where $p \in [-0.5, 0.5]$. Specifically, let $100p$ represent the percentage deviation of Q^{1*}_{ij} from Q^*_{ij}. Then, we have

$$Q^{1*}_{ij} = (1+p)Q^*_{ij}. \tag{16}$$

In order to derive the percentage cost penalty (PCP) incurred in using Q^{1*}_{ij} instead of Q^*_{ij}, let us first define the total relevant cost (TRC) per unit time of using a particular replenishment quantity. The TRC is simply the total of the cost components which are functions of Q_{ij}. Clearly, $\mathrm{TRC}(Q_{ij})$ is the aggregate of the cost components that are directly influenced by a particular selection of Q_{ij} and is given by

$$\mathrm{TRC}(Q_{ij}) = \frac{A_{ij}D_{ij}}{Q_{ij}} + \frac{Q_{ij}\gamma_{ij}}{2}. \tag{17}$$

The PCP using Q^{1*}_{ij} instead of Q^*_{ij} is then given by the following expression:

$$\mathrm{PCP} = \frac{\mathrm{TRC}(Q^{1*}_{ij}) - \mathrm{TRC}(Q^*_{ij})}{\mathrm{TRC}(Q^*_{ij})} \times 100. \tag{18}$$

Applying (16)–(18), we obtain the following expression for the PCP:

$$\mathrm{PCP} = \frac{p^2}{1+p} \times 50. \tag{19}$$

Sensitivity analysis using percentage cost penalty [based on (19)] is shown in Fig. 8, which depicts variations of PCP as a function of p.

C. DACM With Constant Traffic Growth in a General Network Topology

In what follows, we discuss the applications of the inventory-based DACM technique with constant traffic growth rate to more general topological contexts. We first describe the procedure for applying it in a network with non-capacity adaptive routing. In such networks, the routing system does not change with time except when major events occur, such as outages or explicit configuration control actions by the network operator. This is the situation commonly encountered in IP networks without MPLS where the routing system is generally based on unconstrained shortest path (SPF) algorithms (e.g., OSPF and ISIS).

We assume the topology of the network remains static during the entire period of capacity replenishment. Thus, we focus on

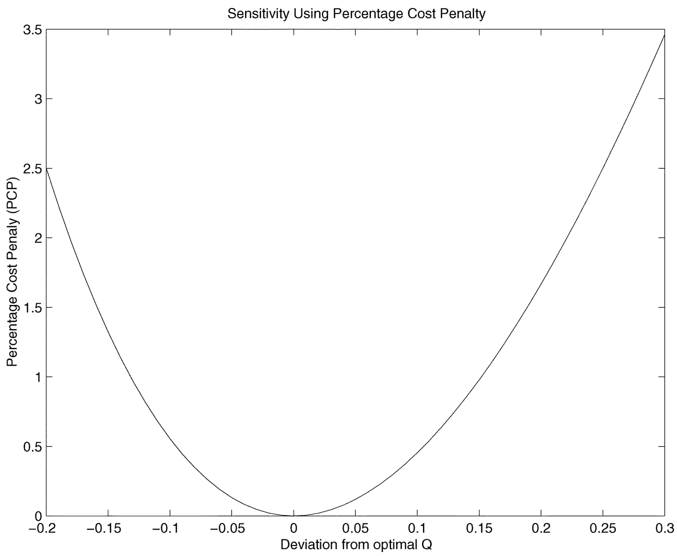

Fig. 8. Sensitivity curve based on percentage cost penalty.

the problem of replenishing capacity within the existing network topology and circumvent the general topology design problem, which is outside the scope of this paper. The methodology presented here is based on decomposition of the problem into a set of link-level capacity replenishment decisions. The approach is summarized in the following algorithm:

Inventory-Based DACM Algorithm: Network Context With Non-Capacity Adaptive Routing

inputs DACM(network) $=< \mathcal{G}, \mathcal{T}, \chi, \mathcal{C}, \mathcal{U} >$

Phase 1: Link-Traffic Assignment

for $u = 1$ to n:
  for $v = 1$ to n:
    Path(u, v) = Route(G, u, v)
    $I(u, v)$ = GenerateIndicator(Path(u, v)).

Phase 2: Link Capacity Replenishment

for t = 1 to TotalTime
  for $i = 1$ to n:
    for $j = 1$ to n:
      if $i \neq j$ and link(i, j).metric $\neq$ inf
        Par(t, i, j) = GenPar($G(t)$, I_{uvij})
        Res(t, i, j) = Comp(Par(t, i, j))
        $G(t+1)$ =Update($G(t)$,Res(t, i, j))

The operation of the algorithm can be summarized as follows:

• First determine the paths between each ingress and egress node in the network using an appropriate non-capacity adaptive routing protocol (e.g., Dijkstra's shortest path algorithm for the OSPF and ISIS protocols).

• Determine the link-traffic indicator functions, namely the set $\{I_{uvij} : \forall u, v, e_{ij}\}$.

• For all links in the network, determine the aggregate utilization of each link based on the demands that traverse the link.

• Determine the rate of growth of traffic on each link based on the rate of growth of the demands that traverse the link.

• Apply the constant demand growth inventory-based DACM algorithm to each link in the network.

• Update the network attributes for the next iteration of the algorithm.

In the above algorithm, the function $\mathrm{Route}(G, \mathrm{u}, \mathrm{v})$ computes a route between node u and v using an appropriate non-capacity adaptive routing algorithm (e.g., a shortest path algorithm). The function $\mathrm{GenerateIndicator}(\mathrm{Path}(u, v))$ generates the link-traffic indicator functions (I_{uvij}) associated with $\mathrm{Path}(i, j)$.

The function $\mathrm{GenPar}(G(t), I_{uvij})$ generates the dynamic link parameters that will be used for the inventory-based DACM computation, such as current link utilization and growth rate of traffic traversing the link. The function $\mathrm{Comp}(\mathrm{Par}(t, i, j))$ computes the inventory-based DACM algorithm for $\mathrm{link}(i, j)$ at time t. The associated return variable ($\mathrm{Res}(t, i, j)$) is simply a data structure that returns various results and performance parameters associated with the inventory-based DACM algorithm at time t. One of the key results is the amount of replenishment capacity for each link at time t (Q^*_{ij}). The function $\mathrm{Comp}(\mathrm{Par}(t, i, j))$ has the following form:

Function: Comp(Par(t, i, j))

Begin (Function Comp(Par(t, i, j)))
  If $A_{ij} = 0$
    $Q^*_{ij} = D_{ij}$
  Elseif $A_{ij} > 0$ **and** $\gamma_{ij} > 0$
    $Q^*_{ij} = \sqrt{\frac{2D_{ij}A_{ij}}{\gamma_{ij}}}$
  If $B_{ij}(t) \leq R$
    Procure Q^*_{ij} *units of bandwidth for link* e_{ij}

Lastly, the function $\mathrm{Update}(G(t), \mathrm{Res}(t, i, j))$ updates the graph attributes for the next iteration.

D. Network DACM With Capacity Adaptive Routing

This subsection considers the network DACM problem with constant traffic growth rates and capacity adaptive routing. The main difference here relative to the prior subsection is that the routing system in the IP domain adapts to take advantage of capacity existing within the network. This means that routes for each source-destination pair in the IP domain may change as traffic between the nodes grow and as capacity is replenished.

In systems with capacity adaptive routing, the routing system within the IP domain attempts to exploit existing capacity in the network as much as possible. Depending on the characteristics of the capacity adaptive routing scheme, this may have the effect of elongating the interval between capacity replenishments and reducing the effective amount of capacity required to operate the network. However, our experiments suggest that this is not always the case and the outcome is contingent on the type of routing scheme. In systems with capacity adaptive routing, the routing function is invoked by the inventory-based DACM algorithm during each replenishment iteration, whereas routing is done just once in systems without capacity adaptive routing.

There are many ways to implement a capacity adaptive routing system within the IP domain. In the following, we depict a capacity adaptive scheme described in [15] which is widely implemented in commercial MPLS label switching routers. This capacity adaptive routing algorithm computes shortest paths between each source-destination pair that satisfies the bandwidth requirements of the traffic between the nodes. A network inventory-based DACM scheme utilizing the simple capacity adaptive routing scheme is shown below.

Network DACM with CSPF Capacity Adaptive Routing

inputs
DACM(network) $= < \mathcal{G}, \mathcal{T}, \chi, \mathcal{C}, \mathcal{U} >$

for t = 1 to TotalTime
  for $u = 1$ to n:
    for $v = 1$ to n:
      PrunedGraph = Prune($G(t)$, Traffic(u, v)
      Path(u, v) = ModDijkstra(PrunedGraph)
      $G(t')$ = AllocBw(Path(u, v), $G(t)$, λ_{uv})
      $I(u, v)$ = GenerateIndicator(Path(u, v)).
  for $i = 1$ to n:
    for $j = 1$ to n:
      if $i \neq j$ and link(i, j).metric $\neq$ inf
        Par(t, i, j) = GenPar($G(t)$, I_{ijuv})
        Res(t, i, j) = Comp(Par(t, i, j))
        $G(t+1)$ =Update($G(t)$,Res(t, i, j))

In the above display, the function $\mathrm{Prune}(G(t), \lambda_{uv}))$ has an input set consisting of the current graph $(G(t))$ at time t containing all pertinent graph attributes (such as the current capacity and metrics of each link) and the current traffic at time t between nodes u and v (λ_{uv}). Using these inputs, the Prune function provides an output consisting of a residual graph $(\mathrm{PrunedGraph})$ in which all links whose capacity is less than λ_{uv} have been removed. The function $\mathrm{ModDijkstra}(\mathrm{PrunedGraph})$ runs Dijkstra's algorithm on the pruned residual graph to determine the minimum cost path between nodes u and v. The function $\mathrm{AllocBw}(\mathrm{Path}(u, v), G(t), \lambda_{uv})$ allocates the traffic between nodes u and v along the path $(\mathrm{Path}(u, v))$. This is accomplished by subtracting the λ_{uv} parameter from the current link capacity of each link along $\mathrm{Path}(u, v)$. The output of this routine is the updated graph $G(t')$. The function $\mathrm{GenerateIndicator}(\mathrm{Path}(u, v))$ generates the link-traffic indicator variables associated with the path between nodes u and v (i.e., I_{ijuv}).

Once the link-traffic indicator variables are derived for all u, v, i, j at each time step, the inventory-based DACM algorithm can then be applied to each link to determine the capacity replenishment schedule.

IV. DACM With Dynamic Traffic Growth Rate

Now, let us consider the situation where traffic growth in the IP domain is deterministic, but the rate of growth changes with time. That is, the assumption that $d_{uv}(t) = d_{uv}$ (a constant) no longer holds. In this case, $d_{uv}(t)$ can be represented by an appropriate piece-wise linear function that captures the desired changes of traffic growth rate with time over the horizon of interest.

For this problem, we shall focus attention on the net effects of the traffic growth variations on a link in relation to bandwidth replenishment across the link. That is, we use the fact that the rate of growth of traffic changes with time to capture the net effects imposed by the aggregate traffic traversing the link. We first characterize the structure of an optimal solution to the dynamic DACM problem and then we define the structure of an optimal solution using a dynamic programming recursion.

For this particular context, we make the replenishment bandwidth a variable, so that $Q_{ij}(t)$ is the amount of bandwidth procured from the optical domain at time t. To simplify the problem, let us assume a discrete time model, so that time has been quantized into a sequence of periods, $t = 0, 1, 2, \ldots, T$, within the horizon of interest, say $[0, T]$, and replenishment decisions occur at the specific discrete times, $t = 0, 1, 2, \ldots, T$.

The problem is to determine how much capacity to procure from the optical domain during each period, $t = 0, 1, 2, \ldots, T$, so that the total cost within the horizon $[0, T]$ is minimized. We need additional definitions in order to formulate the problem.

Let $\delta(y)$ be an indicator function, with the following property:

$$\delta(y) = \begin{cases} 1, & \text{if } y > 0 \\ 0, & \text{if } y \leq 0. \end{cases} \tag{20}$$

With the aforementioned conditions, we can now characterize the structure of an optimal inventory-based DACM solution to the dynamic DACM problem. Given the above hypothesis, the dynamic DACM problem reduces to the *Wagner–Whitin* model of inventory theory, which can be formulated in the following fashion (see, e.g., [24], [25]):

Minimize:

$$\sum_{t=0}^{T-1} [A_{ij}\delta(Q_{ij}(t)) + \alpha_{ij}Q_{ij}(t)] + \sum_{t=1}^{T} \gamma_{ij}B_{ij}(t). \tag{21}$$

Subject to:

$$B_{ij}(0) = 0 \tag{22}$$

$$B_{ij}(t) \geq 0 \quad t = 1, \ldots T \tag{23}$$

$$Q_{ij}(t) \geq 0 \quad t = 0, \ldots T-1 \tag{24}$$

$$B_{ij}(t+1) = B_{ij}(t) + Q_{ij}(t) - D_{ij}(t) \quad t = 0, \ldots T-1. \tag{25}$$

The above formulation of the DACM problem in IPO networks is essentially a *dynamic control problem*. In this formulation, the set of variables $B_{ij}(t)\ \forall i, j$ (available bandwidth at time t across link e_{ij}) represent the state of the system at time t. The variable $D_{ij}(t)$ (rate of growth of traffic at time t across link e_{ij}) represents the external input into the system. Finally, the variable $Q_{ij}(t)$ (the replenishment capacity procured from the optical domain to augment link e_{ij} at time t) represents the control variable which can be adjusted by the network operator to achieve some optimization criteria in the network.

Thus, the DACM problem in IPO networks can be transformed into a dynamic control problem when the traffic demand growth rate is deterministic, but changes with time. Equation (25) is the conservation balance equation, which states that the available bandwidth at the start of period $t+1$ (denoted $B_{ij}(t+1)$) is equal to the available at the start of period t (denoted $B_{ij}(t)$) minus the growth in bandwidth demand during period t (denoted $D_{ij}(t)$) plus the replenishment capacity that was added to the link during period t (denoted $Q_{ij}(t)$).

It is well known in inventory theory that the system of (21), (23)–(25) can be converted into an equivalent dynamic programming formulation. For this purpose, let us simplify the notation.

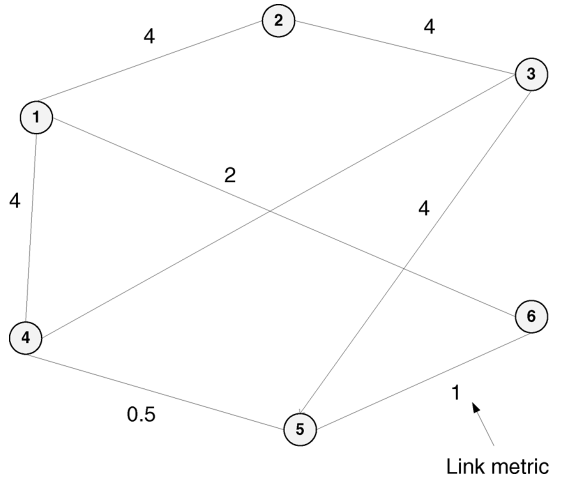

Fig. 9. Sample network for validation of inventory-based DACM with stable routing.

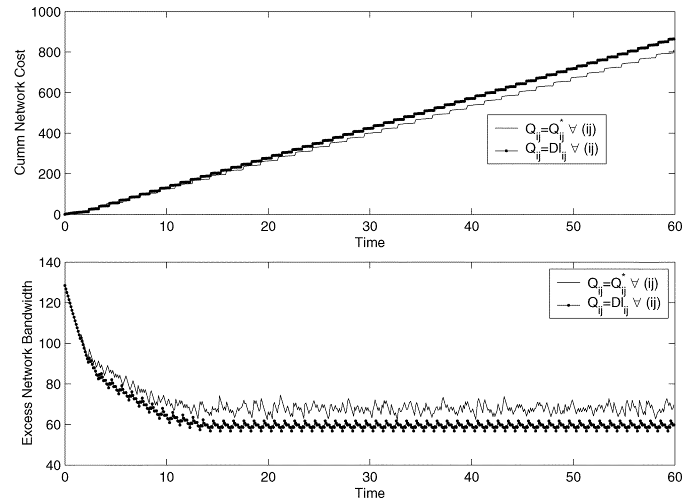

Fig. 10. Network inventory-based DACM with non-capacity adaptive routing: Case 1: Q_{ij}^* and $Q_{ij} = D_{ij}$.

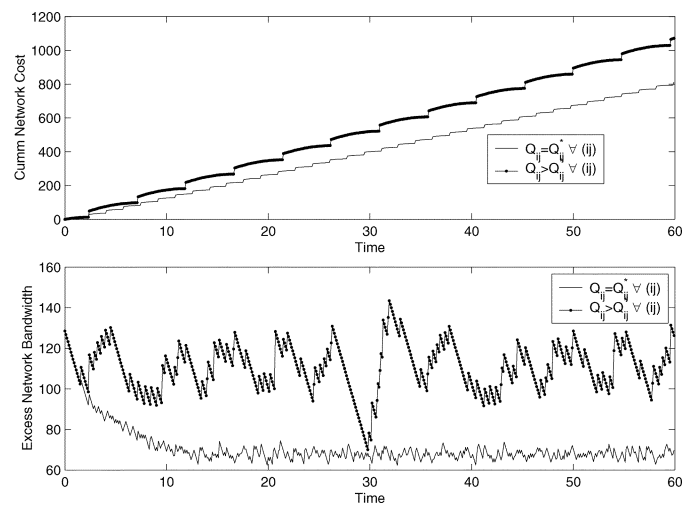

Fig. 11. Network inventory-based DACM with non-capacity adaptive routing: Case 2.

Let b_t be the available bandwidth across the target link at the start of period t (i.e., b_t is equivalent to $B_{ij}(t)$). Let $V_t(b_{t+1})$ be the minimum total replenishment cost incurred for periods $1, 2, \ldots, t$ given that the available capacity across the target link at the end of the period is b_{t+1}. Let Q_t be the amount of bandwidth procured at the start of period t (i.e., Q_t is equivalent to $Q_{ij}(t)$). Let D_t be the demand on the target link during period t. Let $C_t(Q_t)$ be the total cost of procuring Q_t units of capacity in period t. The key step in the dynamic programming solution of the dynamic DACM problem is to express $V_t()$ as a function of $V_{t-1}()$. The dynamic programming recursions are given by the following:

$$V_1(b_2) = \min_{0 \le Q_1 \le D_1 + b_2} \{C_1(Q_1) + \gamma_1 b_2\} \tag{26}$$

$$V_k(b_{k+1}) = \min_{0 \le Q_k \le D_k + b_{k+1}} \{C_k(Q_k) + \gamma_k b_{k+1} \tag{27}$$

$$+ V_{k-1}(b_{k+1} + D_k - b_k)\}. \tag{28}$$

V. Numerical Results and Discussions

In order to evaluate the inventory-based DACM algorithms, we implemented the algorithms and evaluated their performance using different network topologies, under capacity adaptive and non-capacity adaptive routing in the IP domain. This section summarizes some of the results. The sample network of Fig. 9 was used to derive the results described in the first part of this section.

The following performance measures were used to evaluate the effectiveness of the procedures:

1) the cumulative replenishment cost across all links in the network during a well-defined time horizon;
2) the aggregate excess bandwidth in the network as a function of time;
3) the total installed network capacity as a function of time.

Some results relating to inventory-based DACM with non-capacity adaptive routing are illustrated in Figs. 10–13. Each figure depicts two graphs: 1) the cumulative network cost as a function of time and 2) the aggregate excess bandwidth in the network as a function of time. In each case, results derived from the approach described in this paper is compared with results from alternative approaches in which the replenishment capacity used during each iteration differs from the values based on our approach.

Fig. 10 compares the inventory-based DACM replenishment schedule (where $Q_{ij} = Q_{ij}^* \; \forall e_{ij}$) against the situation where $Q_{ij} = D_{ij} \; \forall e_{ij}$. The scenario where $Q_{ij} = D_{ij}$ occurs when the replenishment capacity on each link is equal to the rate of growth of traffic on the link. It can be seen that the inventory-based DACM approach generates cumulative costs that are less than the alternative approach, even though the alternative scheme produces lower aggregate excess bandwidth.

The remaining figures (Figs. 11–13) depict various scenarios where alternative replenishment capacities are used which deviate from the inventory-based DACM quantities. In all cases, the inventory-based DACM approach is seen to be superior in terms of minimizing cumulative network costs associated with excess bandwidth.

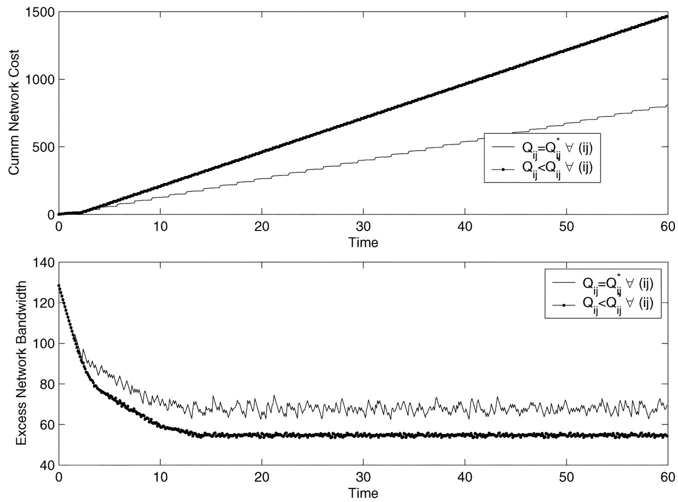

Fig. 12. Network inventory-based DACM with non-capacity adaptive routing: Case 3.

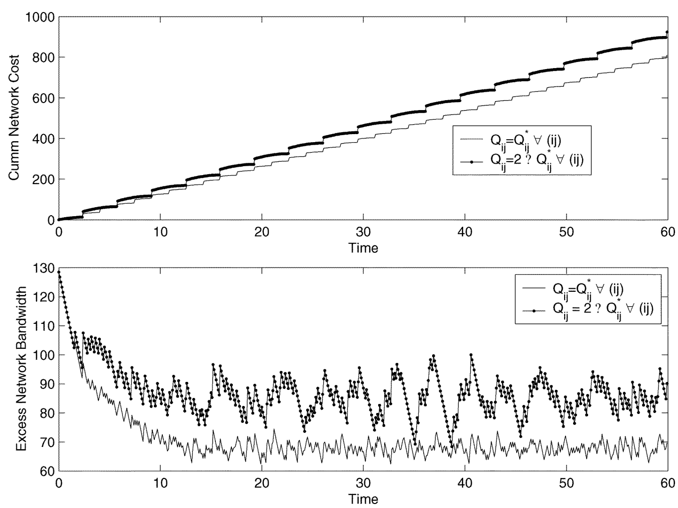

Fig. 13. Network inventory-based DACM with non-capacity adaptive routing: Case 4.

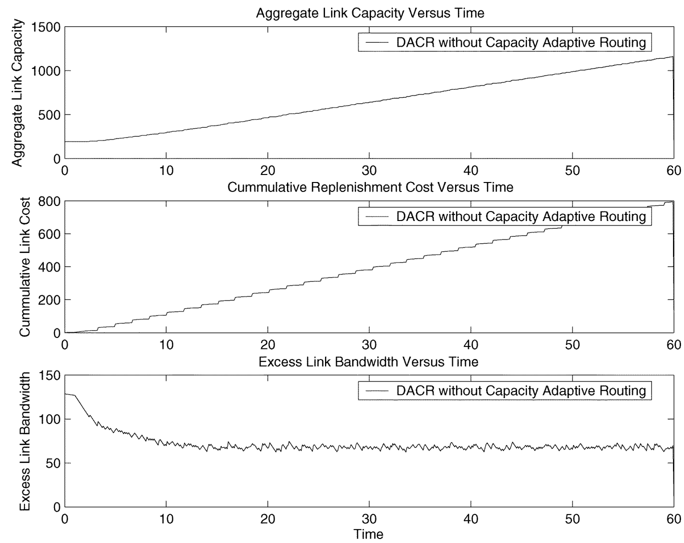

Fig. 14. Network inventory-based DACM without capacity adaptive routing.

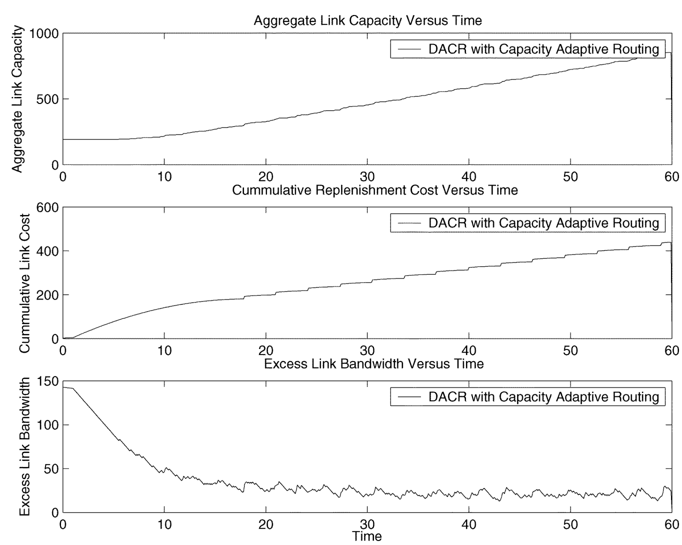

Fig. 15. Network inventory-based DACM with capacity adaptive routing.

In the following, we compare the results for the inventory-based DACM using capacity adaptive routing against results for systems with non-capacity adaptive routing. In systems with capacity adaptive routing, the routing system attempts to exploit existing capacity in the network as much as possible. Depending on the characteristics of the capacity adaptive routing system, this may elongate the interval between capacity replenishments and reduce the frequency of capacity replenishments relative to systems without capacity adaptive routing.

We implemented the capacity adaptive routing scheme based on the pruning procedure described previously and compared the results against results obtained using non-capacity adaptive routing. The performance measures are: 1) aggregate network capacity; 2) cumulative network costs associated with excess capacity; and 3) excess bandwidth.

Fig. 14 shows pertinent results for an inventory-based DACM system with non-capacity adaptive routine. Fig. 15 shows the results for the same network using inventory-based DACM with capacity adaptive routing. For this particular network configuration, Figs. 14 and 15 demonstrate that the inventory-based DACM system with capacity adaptive routing produced superior results relative to the inventory-based DACM system without capacity adaptive routing. However, the superiority of the simple capacity adaptive scheme does not hold in general.

We also conducted numerical experiments on a larger network scenario depicted in Fig. 16, in which the capacity adaptive approach did not produce superior results relative to the non-capacity adaptive case. Fig. 17 depicts pertinent results for the large network using inventory-based DACM with non-capacity adaptive routing. Fig. 18 shows the results for the same network using inventory-based DACM with capacity adaptive routing.

We now proceed to discuss the results presented in this section. The graphs show how variations in Q_{ij} and corresponding replenishment schedules affect cumulative network costs over an interval of time. The results suggest that cumulative network costs derived from the inventory-based DACM replenishment schedules ($Q_{ij} = Q^*_{ij}\ \forall ij$) were smaller than corresponding costs derived from alternative schedules (i.e., when

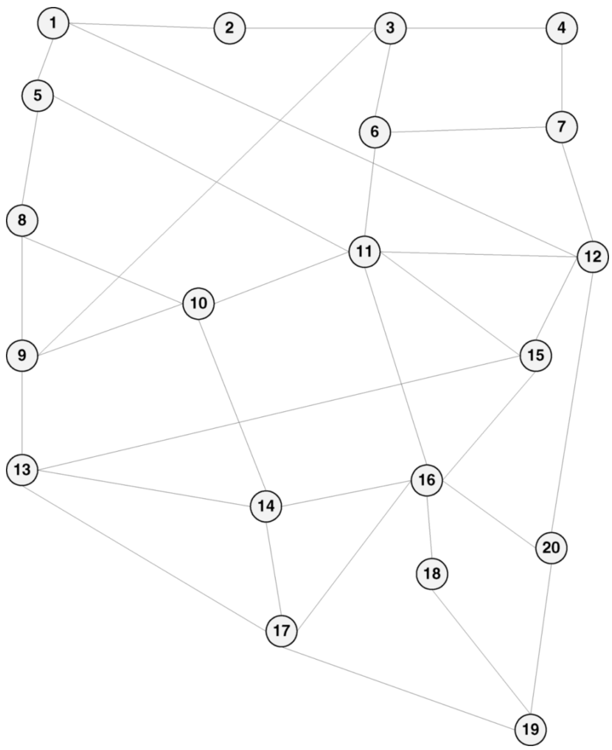

Fig. 16. Sample network for validation of inventory-based DACM with stable routing.

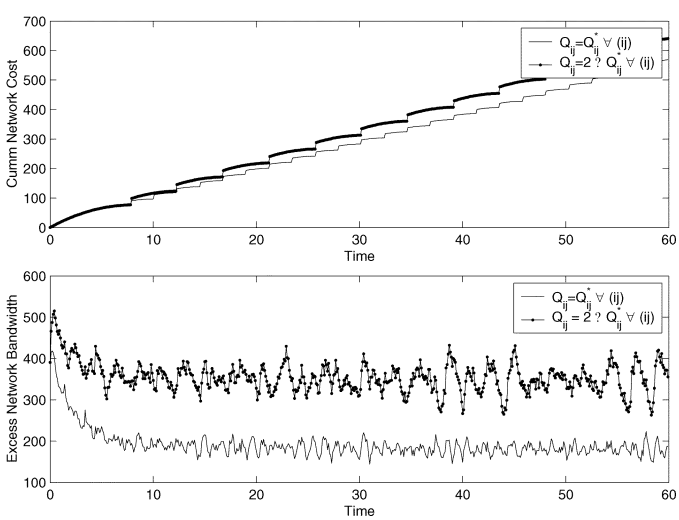

Fig. 17. Large network inventory-based DACM with non-capacity adaptive routing.

$Q_{ij} \neq Q_{ij}^{*}\ \forall ij$). We investigated numerous scenarios in which the replenishment quantities were less than or more than the inventory-based DACM replenishment quantities. In all instances, the cumulative costs from the inventory-based DACM schedule were smaller than costs from alternative approaches.

The figures also illustrate the variations of excess network bandwidth with variations in replenishment quantities. Figs. 10 and 12 demonstrate these variations for scenarios where the replenishment quantities were less than the inventory-based DACM values. In such instances, the aggregate excess bandwidth were lower than corresponding values from

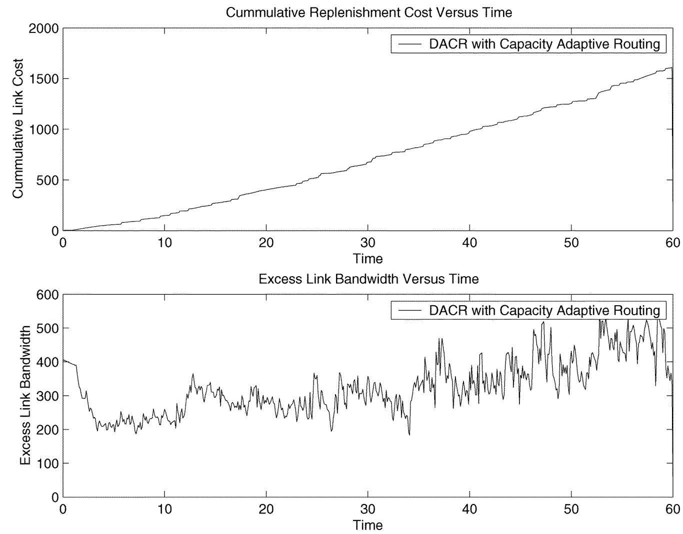

Fig. 18. Network inventory-based DACM with capacity adaptive routing.

the inventory-based DACM schedule. This is clearly to be expected since less capacity is procured during each replenishment event. However, cumulative network costs were higher because capacity replenishments had to occur more frequently resulting in higher fixed replenishment costs. On the other hand, Figs. 11 and 13 illustrate the alternative scenarios where the replenishment quantities were more than the inventory-based DACM values. The graphs indicate that the excess network bandwidth and cumulative costs from these schedules were higher than corresponding values from the inventory-based DACM schedules.

VI. Conclusions

This paper presented a framework for DACM in IPO networks and described a class of inventory-based DACM algorithms for efficient replenishment of network capacity conditioned on the evolutionary processes of Internet traffic. The objective is to improve network economics by leveraging the improved flexibility provided by advances in optical networking technologies. These algorithms take the growth characteristics of IP traffic into account to compute a replenishment schedule which determines when, where, and how much additional capacity to procure from the optical domain. The optimization criteria was to minimize the cost of excess bandwidth in the IP domain over the time horizon of interest. We analyzed the performance characteristics of the inventory-based DACM methods under capacity adaptive and non-capacity adaptive routing strategies. When the rate of growth of traffic is constant, the inventory-based DACM strategy was found to be more cost effective than alternative strategies for non-capacity adaptive systems. In the case of capacity adaptive routing schemes, the routing method in the IP domain had a significant effect on the results.

The numerical study was conducted within a deterministic setting, in which the rate of growth of traffic was assumed to be deterministic, even though it may change with time. However, the framework and architectural concepts for DACM presented here are applicable to more general scenarios. For stochastic settings, the computational machinery needs to be modified to account for the fact that the rate of growth of traffic in the IP domain and the lead time to replenish capacity from the optical domain are random variables. The cases presented here are likely to be representative of those that can be used in future operational IPO infrastructures.

REFERENCES

[1] B. Rajagopalan, J. Luciani, and D. Awduche, "IP over Optical Networks: A Framework," RFC-3717, Mar. 2004.
[2] E. Mannie, "GMPLS Architecture," RFC-3945, Oct. 2004.
[3] D. Awduche, Z. Mansourati, and A. Vernon, "Signaling and routing enhancements to support dynamic wavelength management," presented at the National Fiber Optic Engineers Conf. (NFOEC-2001) Baltimore, MD, Jul. 2001.
[4] H. Zang, J. P. Jue, and B. Mukherjee, "A review of routing and wavelength assignment approaches for wavelength-routed optical WDM networks," *Opt. Netw. Mag.*, vol. 1, no. 1, pp. 47–60, Jan. 2000.
[5] S. Chaudhuri and E. Goldstein, "On the value of optical layer reconfigurability in IP-over-WDM lightwave networks," *Photon. Technol. Lett.*, vol. 12, no. 8, pp. 1097–1099, Aug. 2000.
[6] R. Ramaswami and K. N. Sivarajan, "Design of logical topologies for wavelength-routed optical networks," *IEEE J. Sel. Areas Commun.*, vol. 14, no. 5, pp. 40–51, Jun. 1996.
[7] J.-F. Labourdette, E. Bouillet, and S. Chaudhuri, "Role of optical network and spare router strategy in resilient IP backbone architecture," in *Proc. 4th Int. Workshop on Design of Reliable Communications Networks (DRCN)*, Oct. 2003, pp. 244–253.
[8] A. L. Chiu and J. Strand, "An agile optical layer restoration method for router failures," *IEEE Network Mag.*, vol. 17, no. 2, pp. 38–42, Mar./Apr. 2003.
[9] S. Philips, R. Doverspike, and N. Reingold, "Network studies in IP/optical layer restoration," in *Proc. Optical Fiber Communication Conf. (OFC 2002)*, 2002, pp. 425–427.
[10] O. Hauser, M. Kodialam, and T. V. Lakshman, "Capacity design of fast path restorable optical networks," in *Proc. IEEE INFOCOM*, 2002, pp. 817–826.
[11] M. Kodialam and T. V. Lakshman, "Integrated dynamic IP and wavelength routing in IP over WDM networks," in *Proc. IEEE INFOCOM*, 2001, pp. 358–366.
[12] D. Awduche, "Integrated traffic engineering in IP over optical networks: General considerations," presented at the National Fiber Optic Engineers Conf. (NFOEC-2002) Dallas, TX, Sep. 2002.
[13] D. Awduche and B. Jabbari, "Internet traffic engineering using multiprotocol label switching (MPLS)," *J. Comput. Netw.*, vol. 40, no. 1, Sep. 2002.
[14] D. Awduche, A. Chiu, A. Elwalid, I. Widjaja, and X. Xiao, "Overview and Principles of Internet Traffic Engineering," RFC 3272, May 2002.
[15] D. Awduche, J. Malcolm, J. Agogbua, J. McManus, and M. O'Dell, "Requirements for Traffic Engineering over MPLS," RFC 2702, Sep. 1999.
[16] D. Awduche, "MPLS and traffic engineering in IP networks," *IEEE Commun. Mag.*, vol. 37, no. 12, pp. 42–47, Dec. 1999.
[17] H. Hoc, " Topological optimization of networks: A nonlinear mixed integer model employing generalized Benders decomposition," *IEEE Trans. Autom. Contr.*, vol. 27, no. 1, pp. 164–169, Feb. 1982.
[18] D. Awduche and Y. Rekhter, "Multiprotocol lambda switching: combining MPLS traffic engineering control with optical crossconnects," *IEEE Commun. Mag.*, vol. 39, no. 3, pp. 111–116, Mar. 2001.
[19] L. Berger, "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description," RFC 3471, Jan. 2003.
[20] "Requirements for the Automatic Switched Transport Network (ASTN)," ITU-T Rec. G.8070/Y.1301, 2001.
[21] "Architecture of the Automatic Switched Optical Network (ASON)," ITU-T Rec. G.8080/Y.1304, 2001.
[22] "Link Capacity Adjustment Scheme," ITU-T Rec. G.7042/Y.1305, 2001.
[23] D. Awduche, "Integrated traffic engineering in IP over optical networks: Fundamental concepts," presented at the SPIE ITCOM-2002 Boston, MA, Jul. 2002.
[24] P. Zipkin, *Foundations for Inventory Theory*. New York: McGraw-Hill, 2000.
[25] J. Bramel and D. Simchi-Levi, *The Logic of Logistics*. New York: Springer-Verlag, 1997.
[26] D. Awduche and B. Jabbari, "A framework for demand assigned capacity management (DACM) in IP over optical networks," in *IEEE Int. Conf. Communications*, Jun. 2004, pp. 1589–1593.

Daniel O. Awduche (M'88–SM'03) received the M.S. degree in computer systems engineering from the University of Massachusetts, Amherst, and the Ph.D. degree in information technology from George Mason University, Fairfax, VA. He also received the M.B.A. degree from George Mason University.

He leads the Verizon Business Internet Engineering organization in Ashburn, VA, responsible for network engineering of its public IP network (previously known as UUNET), which is one of the largest Internet backbone networks in the world. He is notable for championing the development and standardization of MPLS traffic engineering technology and pioneering the concept of multiprotocol lambda switching, which is now called Generalized MPLS (GMPLS).

Dr. Awduche has served as a Guest Editor of the *Proceedings of the IEEE* (September 2002, Special Issue on Internet Technology and the Convergence of Communications Services); and a Guest Editor of the IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS (October 2004, Special Issue on Advances in Metropolitan Optical Networks).

Bijan Jabbari (S'80–M'81–SM'85) received the Ph.D. degree in electrical engineering from Stanford University, Stanford, CA.

He is a Professor of electrical engineering at George Mason University, Fairfax, VA, and an affiliated faculty with ENST-Paris, France.

Prof. Jabbari is an International Division Editor for the *Journal of Communications and Networks*, an Editor for the IEEE TRANSACTIONS ON COMMUNICATIONS, and was on the editorial board of *Proceedings of the IEEE*. He is the past chairman of the IEEE Communications Society technical committee on Communications Switching and Routing. He is a recipient of the IEEE Millennium Medal in 2000 and the Washington DC Metropolitan Area Engineer of the Year Award, in 2003. He continues research on multi-access communications and high performance networking.