Managing Capacity for Telecommunications Networks Under Uncertainty

Yann d'Halluin, Peter A. Forsyth, and Kenneth R. Vetzal

***Abstract*—The existing telecommunications infrastructure in most of the world is adequate to deliver voice and text applications, but demand for broadband services such as streaming video and large file transfer (e.g., movies) is accelerating. The explosion in Internet use has created a huge demand for telecommunications capacity. However, this demand is extremely volatile, making network planning difficult. In this paper, modern financial option pricing methods are applied to the problem of network investment decision timing. In particular, we study the optimal decision problem of building new network capacity in the presence of stochastic demand for services. Adding new capacity requires a capital investment, which must be balanced by uncertain future revenues. We study the underlying risk factor in the bandwidth market and then apply real options theory to the upgrade decision problem. We notice that sometimes it is optimal to wait until the maximum capacity of a line is nearly reached before upgrading directly to the line with the highest known transmission rate (skipping the intermediate lines). It appears that past upgrade practice underestimates the conflicting effects of growth and volatility. This explains the current overcapacity in available bandwidth. To the best of our knowledge, this real options approach has not been used previously in the area of network capacity planning. Consequently, we believe that this methodology can offer insights for network management.**

***Index Terms*—Network planning, real options, uncertain demand for capacity.**

I. Introduction

IN THE PAST, bandwidth was traded infrequently, with deals taking months to complete. Performance, which can be monitored in terms of packet losses and/or response times (ping times), was rarely considered. However, in May 1999, Enron proposed the development of a global bandwidth commodity market. The concept of a pooling point where bandwidth market players could settle contracts in a matter of seconds was introduced. Today, long-term contracts (e.g., indefeasible rights of usage or IRUs) are being replaced by shorter term contracts, and bandwidth is moving toward being effectively traded on demand. With a forecasted notional size exceeding 1 trillion dollars annually [1], the bandwidth market is expected to become similar in size to large commodity markets.

The bandwidth market is still in its infancy and high-quality detailed market data do not exist yet. The volatility present in the demand market for capacity requires the development of risk management and investment decision systems. As in most corporate investment decisions, timing is crucial in a competitive market. The traditional approach of valuing an investment, which consists of using the net present value (NPV) rule to decide whether to invest, ignores the opportunity of waiting for new information before committing to the investment. As an alternative, the real options approach [2], [3] can be used to effectively model investment flexibility.[1] A number of publications discuss the use of real options theory for optimal investment timing (e.g., [2]–[4] and references therein), but researchers in network planning and management do not appear to have used these concepts [5], [6].

In this paper, we apply a real options framework to the problem of the optimal timing investment into new capacity. Given a set of lines with different characteristics, we aim to find the percentage (in terms of the maximum transmission rate of a line) at which it is optimal to upgrade. To the best of our knowledge, this approach has not been used previously. Consequently, we believe that this methodology can offer insights for the telecommunications industry. The outline of this paper is as follows. Section II describes the modeling framework; Section III presents the mathematical model and introduces an upgrade decision algorithm; Section IV presents the estimated model parameters; and Section V contains different simulated results. The conclusions are provided in Section VI.

II. Background

In some cases, the fundamental factor driving profitability is the amount which can be sold, as opposed to the price received per unit. The owner of a fiber optic network faces this type of situation. New wavelength services allow the user to purchase limited capacity for days or even hours. Effectively, users pay only for the bandwidth that they use. Consequently, the revenue to the owner of the network is determined by the prevailing price and the amount used (demand for capacity). A study of the bandwidth market reveals some interesting facts. As will be discussed later in Section V, the demand for bandwidth is highly volatile. Our estimate of volatility is about 95%

Manuscript received May 19, 2001; revised February 25, 2002; approved by IEEE/ACM Transactions on Networking Editor S. Low. This work was supported by Bell University Laboratories at the University of Waterloo, Waterloo, ON, Canada.

Y. d'Halluin and P. A. Forsyth are with the Department of Computer Science, University of Waterloo, Waterloo, ON N2L 3G1, Canada (e-mail: ydhallui@elora.uwaterloo.ca; paforsyt@elora.uwaterloo.ca).

K. R. Vetzal is with the Centre for Advanced Studies in Finance, University of Waterloo, Waterloo, ON N2L 3G1, Canada (e-mail: kvetzal@watarts.uwaterloo.ca).

Publisher Item Identifier 10.1109/TNET.2002.801416.

[1]Note that the distinction between the real options approach and the NPV approach is more one of degree than one of kind. The NPV approach, as traditionally applied, fails to sufficiently account for much of the optionality inherent in corporate investments. The real options approach focuses on this optionality. Provided that the values of all options are appropriately included in the analysis, the NPV approach will give the same conclusions as the real options approach.

TABLE I
PRICES FOR DIFFERENT TRANSMISSION RATES. PRICES ARE IN \$/YEAR/DS-0 MILE. A DS-0 LINE HAS A MAXIMUM TRANSMISSION RATE OF 64 kb/s

	Q1 1999	Q1 2000	Q1 2001	Q1 2002
OC-3	0.013	0.011	0.0082	0.0055
OC-12	0.012	0.01	0.0066	0.0045
OC-48	0.01	0.0095	0.0054	0.0035

(see Appendix B). This can be compared with a volatility of 20%–40% for most major stock market indexes. In contrast, the price per bit per second is falling exponentially (see Table I). This phenomenon is in fact typical of computer components. Relentless technological development decrease costs while demand increases exponentially. An example of this effect can be found in storage media. The price per megabyte of disk drives has decreased exponentially. However, a disk manufacturing plant can be profitable as long as demand increases rapidly enough to offset falling prices due to technological improvements.

This situation is in contrast with traditional financial markets where price, and not demand, is the dominating factor (in other words, in financial markets it is almost always assumed that demand curves are perfectly elastic). The current inefficiencies in the bandwidth market can be explained by the fact that deregulation is recent and consumers are paying for capacity rather than paying for consumption. For example, most Internet Service Providers (ISPs) have flat base fees for their clients, but usage is very different across consumers. In an efficient market, price should reflect the actual amount of bandwidth used by a consumer. Already, some companies have started to charge their customers based on usage (e.g., Wavelength Service contracts). This is similar to telephone companies, which charge long distance phone calls on a minute usage basis, or electricity companies, which charge according to consumption. Consequently, we believe that as the bandwidth market becomes more efficient, contracts will be based on bandwidth spot prices. However, during the current period of transition, we do not believe that there is a liquid enough spot market to value contracts and investments based on these prices. Demand for capacity is, in our opinion, the largest uncertainty in bandwidth investment. Consequently, investment decisions should be valued based on demand. However, we must emphasize that this situation will be only temporary, i.e., for the next few years at most. We expect that, as bandwidth market inefficiencies disappear, price and demand may both be determining risky factors. We note that a model for the bandwidth market which includes both price and quantity effects has been described in [7]. Nevertheless, we should point out that demand may always be the main underlying source of risk. To go back to our example of disk storage, the price per megabyte has decreased fairly smoothly over the last few years, while demand for storage media has been more uncertain. Technology-based industries seem to be "demand pushed," i.e., technology keeps shifting the supply curve [2]. For example, at today's level of demand for storage, most users could not use a one terabyte disk drive, even if it was very cheap.

III. MATHEMATICAL MODEL

A. Background on Option Theory

Before proceeding further, let us review the basics of options. A financial *option* gives its holder the right, but not the obligation, to trade (buy or sell) at a future time for a specified price. For example, an investor may own an option to purchase a particular stock for \$50 on June 1. If the price of that stock on June 1 is higher than \$50, e.g., \$60, then the option has turned out to be valuable since the investor can use it to purchase the stock for \$50 and then sell the stock immediately for a profit of \$10. Conversely, if the stock price on June 1 is lower than \$50, the option becomes worthless: the investor would not use it to pay \$50 to own the stock when it could be purchased at a lower price on a financial exchange. The field of *option pricing* is largely concerned with determining the fair price to pay for options. Over the past three decades, this has been a very active area of research in finance (see [8] or [9] for an overview).

The related field of *real options* extends the basic ideas of option pricing to corporate investment decisions. For example, a network planning manager may buy an option on a bundle of dark fiber lines. This gives the manager the right, but not the obligation, to buy this bundle of lines in six months at a certain price X. Six months later, the manager will decide if it is still necessary to buy the dark fiber lines. If so, the manager will compare the exercise price of the option X to the current price Y of the bundle. If $X < Y$, the option will be exercised to buy the bundle of dark fiber lines for X. On the other hand, if $X > Y$ the option will not be exercised. The bundle will simply be purchased at the current price of Y.

Every option has some defining characteristics. Two key ones which we need to consider for our application are:

- an **expiry date** (e.g., six months): the last date at which the option can be exercised. Note that some options have multiple possible exercise dates;
- a **strike price** (X): the price at which one party has agreed to pay the other party should the option be exercised. This is also known as the exercise price.

In addition, in our real options setting we must define:

- an **investment horizon**: the time at which all equipment is written off, and the value of any contracts/rights is converted to cash.

An underlying random factor determines the value of an option contract as an investment. In financial markets, this is usually the price of the asset for which the option can be used. However, in our case, we will consider usage to be the underlying factor during the transition from regulated to deregulated markets. In order to use this factor, we need to show how we can model the demand for capacity over time.

B. Mathematical Model

Let $\mathcal{Q}$ represent the demand for capacity (measured in megabits). The different possible paths followed by the demand can be modeled as a stochastic process given by

$$d\mathcal{Q} = \mu\mathcal{Q}\,dt + \sigma\mathcal{Q}\,d\mathcal{Z} \quad (1)$$

where μ is the drift rate or growth rate, σ is the volatility, and $d\mathcal{Z}$ is the increment of a Wiener process (readers unfamiliar with these ideas should consult [8], [2], and [9] for a simple introduction).

Based on hedging arguments (see [8], [9], and Appendix A), a partial differential equation for the value of an investment $\mathcal{V}(t, \mathcal{Q})$ is found to be

$$\frac{\partial \mathcal{V}}{\partial t}+\frac{1}{2}\sigma^2\mathcal{Q}^2\frac{\partial^2\mathcal{V}}{\partial \mathcal{Q}^2}+(\mu-\kappa\sigma)\mathcal{Q}\frac{\partial \mathcal{V}}{\partial \mathcal{Q}}-r\mathcal{V}+\mathcal{R}(\mathcal{Q}, t)=0 \quad (2)$$

where $\mathcal{V}$ is the value of the investment in dollars ($), $\mathcal{R}(\mathcal{Q}, t)$ is the revenue term in $/time, r is the risk free interest rate, and κ is the market price of risk. This latter variable is a function of $\mathcal{Q}$. It will hereafter be referred to as the "telecom market price of risk." Essentially, κ captures the tradeoffs between risk and return for investments that are dependent on $\mathcal{Q}$. A complete derivation of the equation is given in Appendix A.

In the pricing of financial options, the price of the option at the expiry date is known (as a function of the underlying stock price), but the price of the option before expiry is unknown (and that is what we would be attempting to determine). In our case, we consider an investment horizon T. Mathematically, we then have

$$\mathcal{V}(T, \mathcal{Q}) = f(\mathcal{Q}).$$

Although the methods discussed in this paper can be used with arbitrary $f(\mathcal{Q})$, for simplicity we will restrict attention to the case where the value of all capital investment at T is assumed to be zero, i.e., $f(\mathcal{Q}) = 0$. We will take this investment horizon to be $T = 5$ years. This may be somewhat pessimistic, but in the rapidly changing field of telecommunications, this is perhaps an appropriate length of time to consider.

Since the value of the investment is known at $t = T$, the forward equation (2) is transformed into a backward equation by substituting $\tau = T - t$ (which evolves from the future investment horizon date T back to the present date) to give

$$\frac{\partial \mathcal{V}}{\partial \tau}=\frac{1}{2}\sigma^2\mathcal{Q}^2\frac{\partial^2\mathcal{V}}{\partial \mathcal{Q}^2}+(\mu-\kappa\sigma)\frac{\partial \mathcal{V}}{\partial \mathcal{Q}}-r\mathcal{V}+\mathcal{R}(\mathcal{Q}, \tau). \quad (3)$$

In the following, we describe the factors that need to be taken into account when determining the optimal decision to upgrade and their effects on our initial modeling equation (3).

C. Payment

We assume that the owner of the line receives continuous payments. For line i of maximum transmission rate $\overline{\mathcal{Q}_i}$ (in megabits), we have

$$\mathcal{R}_i(\mathcal{Q}, \tau) = \min(\mathcal{Q}, \overline{\mathcal{Q}_i})\mathcal{D}\mathcal{P}(\tau) \quad (4)$$

where $\mathcal{D}$ is the length of the line (in miles) and $\mathcal{P}(\tau)$ is the spot price at time τ for bandwidth (in $/year/mile/megabit). The payment received can be no larger than the maximum transmission rate of the line multiplied by the price and the length of the line. We assume that the price is a known decreasing function of time [1]. Note that this does not create an arbitrage opportunity

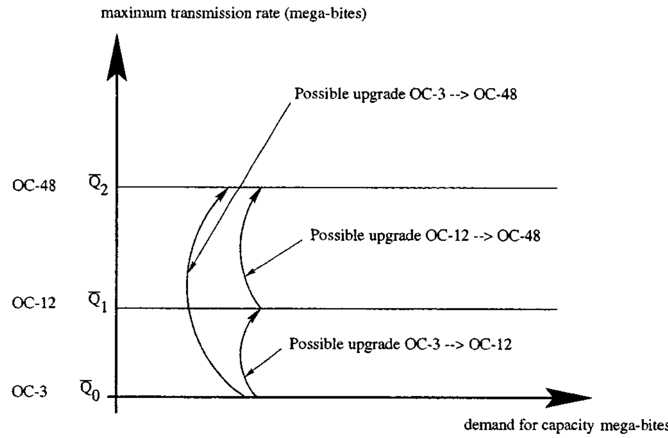

Fig. 1. We consider a set of lines $\mathcal{Q}_i$ with maximum transmission rate $\overline{\mathcal{Q}_i}$. We solve a set of PDEs (7) for each upgrade possibility.

because unused bandwidth cannot be stored for later use. This function is given by

$$\mathcal{P}(\tau) = \mathcal{P}_0 \exp(-\alpha(T-\tau))$$

where $\mathcal{P}_0$ is the current spot price and α is a decay parameter determining the rate at which the spot price decreases. Note that we require that the spot price be the same across the various possible lines at any point so as to avoid arbitrage. For example, if the spot price (in $/year/mile/megabit) of an OC-48 line was less than that of an OC-12, we could buy capacity on an OC-48 line and then immediately sell it at the OC-12 spot price. Consequently, (3) becomes

$$\frac{\partial \mathcal{V}}{\partial \tau}=\frac{1}{2}\sigma^2\mathcal{Q}^2\frac{\partial^2\mathcal{V}}{\partial \mathcal{Q}^2}+(\mu-\kappa\sigma)\frac{\partial \mathcal{V}}{\partial \mathcal{Q}}-r\mathcal{V} + \min\left(\mathcal{Q}, \overline{\mathcal{Q}_i}\right)\mathcal{D}\mathcal{P}(\tau). \quad (5)$$

D. Maintenance Costs

Each line has a different maintenance cost that is constant with time. The maintenance costs are paid at discrete time intervals (i.e., monthly). Let τ^- and τ^+, respectively, denote the times immediately before and after a maintenance payment. Given a line i with maximum transmission rate $\overline{\mathcal{Q}_i}$, we have

$$\mathcal{V}_i(\mathcal{Q}, \tau^+) = \mathcal{V}_i(\mathcal{Q}, \tau^-) - \mathcal{M}_i\mathcal{D}\Delta\tau \quad (6)$$

where $\mathcal{M}_i$ is the maintenance cost of the line in $/year/mile, $\mathcal{D}$ is the length of the line in miles, and $\Delta\tau = 1/12$.

E. Upgrade Decision

Consider a set of lines $\mathcal{Q}_i$ with maximum transmission rate $\overline{\mathcal{Q}_i}$. For example, $\overline{\mathcal{Q}_1}$ could be an OC-12, $\overline{\mathcal{Q}_2}$ an OC-48, and so on (see Fig. 1). Let $\mathcal{V}_i$ be the value of an investment in a line with capacity $\overline{\mathcal{Q}_i}$. We must solve a set of partial differential equation (PDEs) (3) for each upgrade possibility, i.e.,

$$\frac{\partial \mathcal{V}_i}{\partial \tau}=\frac{1}{2}\sigma^2\mathcal{Q}^2\frac{\partial^2\mathcal{V}_i}{\partial \mathcal{Q}^2}+(\mu-\kappa\sigma)\frac{\partial \mathcal{V}_i}{\partial \mathcal{Q}}-r\mathcal{V}_i + \min\left(\mathcal{Q}, \overline{\mathcal{Q}_i}\right)\mathcal{D}\mathcal{P}(\tau). \quad (7)$$

A set of equations (7) must be solved for each possible line capacity $i = 1, \ldots, n$, where n is the maximum number of

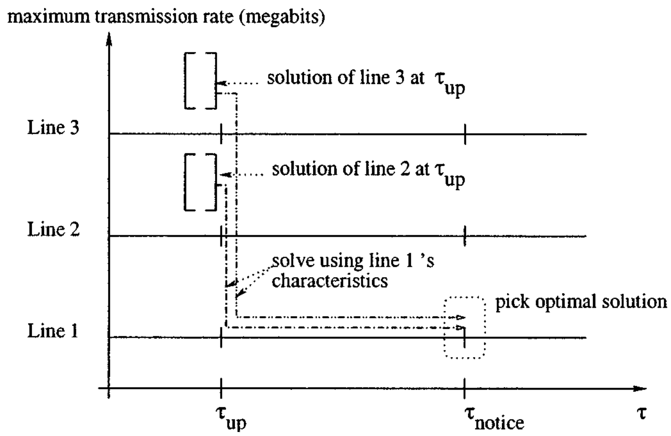

Fig. 2. We consider a set of lines $\mathcal{Q}_i$ with maximum transmission rate $\overline{\mathcal{Q}_i}$. At the notice date τ_{notice}, for each line j with maximum transmission rate higher than $\overline{\mathcal{Q}_i}$, we solve the PDE [(7)] of line i using the solution of line j at time τ_{up}, to obtain $V_i^j(\tau_{\text{notice}})$. We then compare each solution $V_i^k(\tau_{\text{notice}})$ for $k = i, \ldots, n$ to determine the optimal choice.

types of line. This can be done easily using a general numerical PDE solver [10], [11]. This numerical PDE approach involves discretizing equation (7) using a finite volume method [12]. This method has been extensively studied in [10] and [11]. As it is beyond the scope of this paper, we will not present the details of the discretization scheme here. Interested readers should see [13] and [12]. It is worth pointing out that the actual algorithm used shares some characteristics with that used in [13] for pricing a callable bond.

We assume that the upgrade decision is evaluated periodically (i.e., on a quarterly basis). Times at which this evaluation takes place are called notice dates. We use a dynamic programming approach. Essentially, we solve the PDEs (7) backward in time (τ increasing) and determine the optimal decision at each notice date. Consider the lines ordered as $i = 1, \ldots, n$ where $\overline{\mathcal{Q}_{i+1}} > \overline{\mathcal{Q}_i}$ (see Fig. 1). At each notice date τ_{notice}, for each line j with maximum transmission rate higher than $\overline{\mathcal{Q}_i}$, we solve the PDE (7) of line i using the solution of line j at time τ_{up}, to obtain $V_i^j(\tau_{\text{notice}})$. We then compare each solution $V_i^k(\tau_{\text{notice}})$ for $k = i, \ldots, n$ to determine the optimal choice (see Fig. 2). More specifically, our algorithm is:

FOR $i = 1 \cdots n$ **DO**
FOR $j = i + 1 \cdots n$ **DO**
FOR all $\mathcal{Q}$ **DO**

$$\mathcal{V}_i^j(\mathcal{Q}, \tau_{\text{notice}}) = \max(\mathcal{V}_i^j(\mathcal{Q}, \tau_{\text{notice}}) - \mathcal{K}_{i\to j}(\tau_{\text{notice}}), \mathcal{V}_i^i(\mathcal{Q}, \tau_{\text{notice}}))$$

END
END
END

where $\mathcal{K}_{i\to j}(\tau_{\text{notice}})$ is the cost of upgrading from line i of maximum transmission rate $\overline{\mathcal{Q}_i}$ to line j of maximum transmission rate $\overline{\mathcal{Q}_j}$. The upgrade costs are assumed to follow the same decreasing pattern as the spot price per year/mile/megabit. Thus we will use the same decay factor α as for the spot price, i.e.,

$$\mathcal{K}_{i\to j}(\tau) = \hat{\mathcal{K}}_{i\to j} \exp(-\alpha(T - \tau)) \tag{8}$$

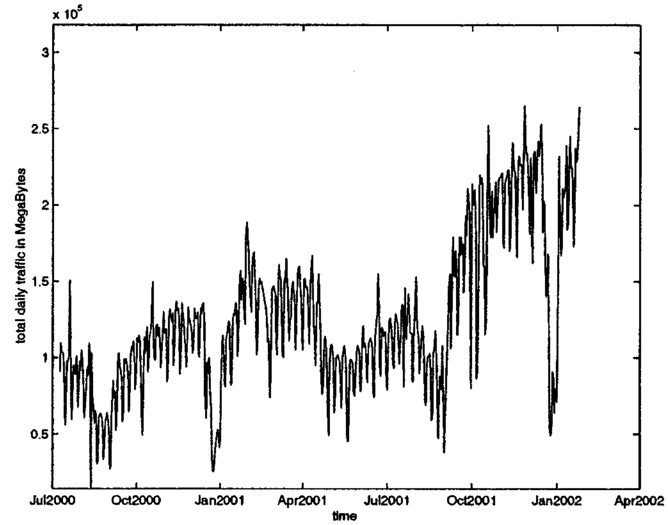

Fig. 3. Total daily network traffic into and out of the University of Waterloo since July 7, 2000.

where $\hat{\mathcal{K}}_{i\to j}$ is the initial upgrade cost from line i to j in \$. Note that the algorithm above allows for the possibility that it may be optimal to jump several levels of capacity when the demand for capacity $\mathcal{Q}$ is very large or the cost $\mathcal{K}(\tau)$ is low.

When the decision to upgrade is made, there is a delay before the new equipment is available (usually three months) [14], [15]. Meanwhile, the line to be upgraded still generates revenue. There is no downtime period during which the stream of revenues is interrupted.

IV. Estimation of the Parameters

Parameter values will obviously have a large impact on our computed results. Parameter estimation is complicated by the fact that the bandwidth market is still in its infancy and past data is limited. This section describes how parameters were estimated for this paper in broad terms. More details are provided in Appendices B–D.

To estimate the uncertainty parameter σ, we use data from the University of Waterloo campus network [16]. Fig. 3 shows the total daily network traffic into and out of the university since July 7, 2000. An initial analysis of the data showed strong autocorrelation of the time-series within each week. This is not surprising, since we expect that there will be repetitive patterns within each week. To filter out this effect, we average the network traffic for the entire period for each day of the week separately and choose the day with the highest average total daily traffic. We then use this same day each week to estimate the effects from week to week. This smoothing is not unusual and should not be considered as a source of inaccuracy. We find that $\sigma = 0.95$. As an aside, note that as σ^2 is measured on a per year basis, σ is in units of year$^{1/2}$. See Appendix B for more details on the estimation of σ.

The growth rate μ is also obtained from University of Waterloo data. The data set used here, however, differs from that used to estimate σ. In particular, to estimate μ, we use summary traffic data for each academic term dating back to 1997. Statistically, this should provide a better estimate of the long

TABLE II
TRANSMISSION RATE CONVERSION

Rate	Max Capacity (Mbps)	Voice Circuit Capacity
OC-12	622	8064
OC-48	2488	32256
OC-192	9952	129024
OC-768	39808	516096

TABLE III
ESTIMATED PARAMETERS

Parameter	Estimated Value
σ	0.95
r	0.05
μ	0.75
κ	0.10
P	0.90
α	1.40

run growth trend (though it would not give a good estimate of volatility, as the time between observations is too long). Our estimate of about $\mu = 0.75$ per year is consistent with the results reported in [17]. (See Appendix B for further details).

We use $r = 0.05$ per year for the risk-free interest rate. The telecom price of risk κ (in units of year$^{1/2}$ is obtained by fitting the price of dark fiber to the cost of an OC-12 investment (see Appendix D for details). The prevailing price of dark fiber when our estimate was made was $1500 per mile between New York and Toronto. We find that $\kappa \approx$ 10%. As a rough check on this estimate, note that the expected return on an investment in the telecom sector should be about $r + \kappa\sigma$ (this assumes that the risk of a typical telecom investment is approximately the same as that of investing in dark fiber). Using $r = 0.05$, $\kappa = 0.10$, and $\sigma = 0.95$ implies an expected return of approximately 15%, consistent with that commonly used in NPV calculations in the industry [18].

In this study, we consider only the fiber optic lines contained in Table II. At the time of writing, the spot price for an OC-12 was approximately 0.069 dollars/year/mile/voice capacity [15]. To obtain the lease price per year per megabit, we multiply the OC-12 price by the voice circuit capacity of an OC-12 line and divide by the OC-12 maximum capacity in megabits. We find that the spot price $P = 0.90$/year/mile/megabit.

The decay factor is obtained from the observation that prices are decreasing by about 50% every six months. We find that $\alpha = 1.4$ per year. These estimated parameters are summarized in Table III.

The costs of upgrading and maintenance are critical to determining when it is optimal to upgrade. We obtained the data in Table IV from the hardware manager of the Mathematics Faculty Computing Facility of the University of Waterloo [14]. These costs are based on the fact that upon upgrading only the switching cards are changed, while the rest of the basic hardware stays the same. We also made this assumption when estimating the telecom market price of risk (see Appendix D). From Table IV, we build the upgrade cost Table V.

TABLE IV
ESTIMATED UPGRADE COSTS FOR BOTH SWITCHES IN $. THESE UPGRADE COSTS ARE BASED ON THE ASSUMPTION THAT THE BASIC HARDWARE INFRASTRUCTURE IS ALREADY IN PLACE AND ONLY SWITCHING CARDS HAVE TO BE CHANGED

Upgrade from line i to line j	Price $
OC-3 $\rightarrow$ OC-12	4,000
OC-12 $\rightarrow$ OC-48	30,000
OC-48 $\rightarrow$ OC-192	80,000
OC-192 $\rightarrow$ OC-768	160,000

TABLE V
UPGRADE COST TABLE CONSTRUCTED FROM THE COST ESTIMATES IN TABLE IV

	OC-12	OC-48	OC-192	OC-768
OC-12	n.a.	$30,000	$80,000	$160,000
OC-48	n.a.	n.a.	$80,000	$160,000
OC-192	n.a.	n.a.	n.a.	$160,000

TABLE VI
MAINTENANCE COSTS IN $/MILE/YEAR FOR DIFFERENT TRANSMISSION RATES

Rate	Cost in $/mile/year
OC-12	2.4
OC-48	18
OC-192	48
OC-768	96

As a rule of thumb, maintenance costs are usually between 3% and 10% of the upgrade capital cost. To simplify matters, we assume that maintenance costs do not include repeater costs. Repeaters in fiber optic lines are used to clean noise from the signal; this noise is intrinsic to the transportation of a signal over a long distance. The details of the maintenance cost computation are contained in Appendix C. The maintenance cost rates can be found in Table VI.

V. SIMULATION RESULTS

In this section we conduct a sensitivity analysis of the optimal decision to upgrade to the different model parameters. Unless specified otherwise, we are solving for a five-year investment, with continuous revenues, monthly maintenance costs, and quarterly investment decisions. We consider the Toronto/New York City pair which is 550 miles apart. We study the optimal investment strategy for four different transmission rates: OC-12, OC-48, OC-192, and OC-768.

A. Base Case

We begin by considering the issues raised by the optimal decision scheme. Anecdotal evidence suggests that lines were upgraded when usage reached about 50% of the maximum transmission rate. In Table VII, we notice that the minimum upgrade percentage found for an OC-12 and OC-48 is 101% and 76%. The minimum percentage for an OC-192 line is found to be 18%. This low percentage can be explained by the fact that the maximum transmission rate considered in these

TABLE VII
UPGRADE PERCENTAGE FOR DIFFERENT TRANSMISSION RATES. BASE CASE PARAMETERS. WE SOLVE FOR A FIVE-YEAR INVESTMENT HORIZON WITH $\mu = 0.75$, $\sigma = 0.95$, $\kappa = 0.1$, AND $r = 0.05$. UPGRADE DECISIONS ARE MADE QUARTERLY. ONLY THE UPGRADE RESULTS FOR THE FIRST TWO YEARS ARE REPORTED

	OC-12		OC-48		OC-192	
Time (year)	Upgrade (%)	Upgrading to	Upgrade (%)	Upgrading to	Upgrade (%)	Upgrading to
0.00	101.28	OC-48	76.36	OC-768	18.00	OC-768
0.25	104.50	OC-48	82.39	OC-768	19.00	OC-768
0.50	109.32	OC-48	88.42	OC-768	20.00	OC-768
0.75	118.97	OC-48	97.26	OC-768	22.00	OC-768
1.00	125.40	OC-48	108.52	OC-768	24.20	OC-768
1.25	144.69	OC-48	120.57	OC-768	27.50	OC-768
1.50	160.77	OC-48	136.65	OC-768	32.00	OC-768
1.75	176.84	OC-48	160.77	OC-768	36.00	OC-768
2.00	209.00	OC-48	184.88	OC-768	40.00	OC-768

TABLE VIII
UPGRADE PERCENTAGE FOR DIFFERENT TRANSMISSION RATES. EFFECTS OF GROWTH RATE AND VOLATILITY. WE SOLVE FOR A FIVE-YEAR INVESTMENT HORIZON WITH $\kappa = 0.1$ AND $r = 0.05$. UPGRADE DECISIONS ARE MADE QUARTERLY. ONLY THE UPGRADE RESULTS FOR THE FIRST TWO YEARS ARE REPORTED

	$\mu = 1.25, \sigma = .95$		$\mu = 0.75, \sigma = .35$	
	OC-48		OC-48	
Time (year)	Upgrade (%)	Upgrading to	Upgrade (%)	Upgrading to
0.00	52.25	OC-768	60.28	OC-768
0.25	56.27	OC-768	36.17	OC-192
0.50	64.30	OC-768	38.18	OC-192
0.75	72.34	OC-768	44.21	OC-192
1.00	80.38	OC-768	48.23	OC-192
1.25	88.42	OC-768	56.27	OC-192
1.50	100.48	OC-768	68.32	OC-192
1.75	116.55	OC-768	80.38	OC-192
2.00	136.65	OC-768	92.44	OC-192

simulations is that of an OC-768 line. Thus, this percentage is not truly representative of what would happen if we had included lines with a higher transmission rate. Note that these minimum percentages are all found at the start of the five-year period: the upgrade percentages rise as time evolves over the first two years. This is because it takes time to recover the costs of upgrading and the higher costs maintenance associated with a higher capacity line. With less time remaining in the five-year horizon, there is less incentive to upgrade.

Furthermore, we see that for an OC-48 line it is better to upgrade to the maximum transmission rate attainable (i.e., OC-768) rather than to upgrade to the intermediate OC-192 line. This appears to contradict the common conception that upon upgrading we should go to the next available transmission rate. This contradiction with practice comes from the high degree of uncertainty in the demand for capacity. Intuitively, this means that it is worthwhile waiting to see how demand evolves and then upgrading to the maximum transmission rate line we can build (e.g., OC-768) at that time. It can be optimal to upgrade an OC-48 line to an OC-192 line, but in this particular case it is optimal to wait until demand is at 76% of an OC-48 transmission capacity and then to upgrade directly to an OC-768. We emphasize that the results indicate that it is better to wait until the maximum capacity of the line is reached and then upgrade to the highest transmission rate possible, rather than to upgrade incrementally at lower usages. This is due to decreasing upgrade costs (with time) and the uncertainty due to volatility in usage.

B. Uncertainty and Growth Rate Sensitivity Analysis

In this scenario, we try to reconcile the results given by our simulations and past industry practice. By inspection, we try to find the volatility and the growth rate that would make our model results agree with industry practice.

The anecdotal rule that lines should be upgraded at 50% of the maximum transmission rate appears to apply only for cases of extremely high growth ($\sigma = 0.95$, $\mu \geq 1.25$) or very low volatility ($\sigma = 0.35$, $\mu = 0.75$) when considering the investment decision for an OC-48 line (see Table VIII). These values are not realistic given today's market conditions. They are indicative of either overoptimism about growth or underestimation of volatility. This would suggest that past industry practice would lead to overcapacity, consistent with today's market.

C. Penalty Due to Network Congestion

Due to the uncertainty in demand, we have observed that it may be optimal to wait until the maximum capacity for a line is reached before upgrading. Essentially, this is because an increase in usage may be a random event and may not be sustained. However, there will be many times when the demand for capacity exceeds the maximum capacity available. Up to now, we have ignored any detrimental effects due to network congestion. We remark that some European operators have contracts where performance is guaranteed, i.e., slow packets are delivered free. As well, a congested network may drive customers to other bandwidth suppliers. We will model these effects in a

TABLE IX
UPGRADE PERCENTAGE FOR DIFFERENT TRANSMISSION RATES. EFFECTS OF NETWORK CONGESTION. WE SOLVE FOR A FIVE-YEAR INVESTMENT HORIZON WITH $\mu = 0.75$, $\sigma = 0.95$, $\kappa = 0.1$, AND $r = 0.05$. UPGRADE DECISIONS ARE MADE QUARTERLY. ONLY THE UPGRADE RESULTS FOR THE FIRST TWO YEARS ARE REPORTED. REVENUES ARE SET TO ZERO IF THE DEMAND IS HIGHER THAN 120% OF THE LINE'S MAXIMUM TRANSMISSION RATE

	OC-12		OC-48	
Time (year)	Upgrade (%)	Upgrading to	Upgrade (%)	Upgrading to
0.00	72.34	OC-48	64.30	OC-768
0.25	75.56	OC-48	68.32	OC-768
0.50	78.77	OC-48	72.34	OC-768
0.75	81.99	OC-48	80.38	OC-768
1.00	86.81	OC-48	88.42	OC-768
1.25	94.85	OC-48	96.46	OC-768
1.50	102.89	OC-48	108.52	OC-768
1.75	115.75	OC-48	120.57	OC-768
2.00	136.65	OC-48	136.65	OC-768

simple way. We assume that if the demand for capacity exceeds the maximum transmission rate available by 20%, revenue is reduced to zero. Again, we could view this as a penalty factor introduced into contracts or as a penalty for producing customer dissatisfaction. The new revenue term is given by [see (7)]

$$\mathcal{R}_i(\mathcal{Q}, \tau) = \begin{cases} \min\left(\mathcal{Q}, \overline{\mathcal{Q}_i}\right) \mathcal{DP}(\tau), & \text{if } \mathcal{Q} \leq 1.2\overline{\mathcal{Q}_i} \\ 0, & \text{otherwise.} \end{cases}$$

As expected, the upgrade decision arises much earlier. See the results reported in Table IX.

VI. CONCLUSION

In this paper, we have considered only a limited number of scenarios. However, our modeling framework allows us to consider numerous possibilities. For example, we could solve for longer investment horizons (e.g., more than five years) and introduce new lines that will only be available in the future (e.g., four, five, and six years from now). Nevertheless, our study allows us to draw some very interesting conclusions.

For our base case parameters, the only situation for which the anecdotal rule of upgrading at a usage level of 50% of maximum capacity is optimal is for an OC-192 line. However, this is an unusual case because there is only one line of higher capacity to which it could be upgraded. For lines of lower capacity, e.g., an OC-48 line, we find that the 50% rule is only optimal if we use apparently unrealistic parameters (either very high growth or very low volatility). Past industry practice apparently reflects either excessive optimism about growth rates or an underestimation of volatility. This is a possible explanation for the current overcapacity in bandwidth.

We also notice that in some cases it is optimal to skip the intermediate lines and go directly to the line with the highest transmission rate. This results with an upgrade decision at about 76% of the current transmission rate when there is the full five years remaining. These results are in general agreement with the conclusions (in a different context) of [4].

Finally, we remark that a certain number of numerical issues arise in our application. Since it was outside the scope of this paper, most of the numerical issues (e.g., oscillations, numerical accuracy) were not discussed. However, for future work it may be worth considering a detailed study of the convergence of the solution and how it is affected by the discontinuities introduced at each notice date. Moreover, we could extend our modeling framework to handle discontinuous jumps in the demand for capacity.

APPENDIX A
DERIVATION OF THE MATHEMATICAL MODEL

Let $\mathcal{Q}$ be the variable representing the demand for capacity. Mathematically, the different paths followed by the demand can be modeled by a stochastic process defined by

$$d\mathcal{Q} = \mu \mathcal{Q}\, dt + \sigma \mathcal{Q}\, d\mathcal{Z} \tag{9}$$

where $\mathcal{Q}$ is in megabits and $d\mathcal{Z}$ is a Wiener process. Let $\mathcal{V}(\mathcal{Q}, t)$ be the value of an investment dependent only on $\mathcal{Q}$ and time t. Using Itô's lemma, the process followed by $\mathcal{V}(\mathcal{Q}, t)$ is

$$d\mathcal{V} = (\beta\mathcal{V} + \mathcal{R})\, dt - \gamma\mathcal{V}\, d\mathcal{Z} \tag{10}$$

where $\mathcal{R}$ represents revenue in \$/year

$$\beta\mathcal{V} = \frac{\partial\mathcal{V}}{\partial t} + \mu\mathcal{Q}\frac{\partial\mathcal{V}}{\partial\mathcal{Q}} + \frac{1}{2}\sigma^2\mathcal{Q}^2\frac{\partial^2\mathcal{V}}{\partial\mathcal{Q}^2} \tag{11}$$

and

$$\gamma\mathcal{V} = -\sigma\mathcal{Q}\frac{\partial\mathcal{V}}{\partial\mathcal{Q}}. \tag{12}$$

Let us pick two investments V_1 and V_2 expiring at some future time ($>t$). From (10), we have

$$\begin{aligned} d\mathcal{V}_1 &= (\beta_1\mathcal{V}_1 + \mathcal{R}_1)\, dt - \gamma_1\mathcal{V}_1\, d\mathcal{Z} \\ d\mathcal{V}_2 &= (\beta_2\mathcal{V}_2 + \mathcal{R}_2)\, dt - \gamma_2\mathcal{V}_2\, d\mathcal{Z}. \end{aligned}$$

Both V_1 and V_2 have the same factor of uncertainty $d\mathcal{Z}$. We can thus construct a portfolio Π composed of V_1 and V_2 such that the return of this portfolio Π is nonstochastic. Let x_1 be the fraction of the amount invested in V_1 and x_2 be the fraction of the amount invested in V_2. Note that $x_1 + x_2 = 1$. The return on the portfolio is given by

$$\begin{aligned} d\Pi &= x_1\frac{d\mathcal{V}_1}{V_1} + x_2\frac{d\mathcal{V}_2}{V_2} \\ &= \left(x_1\beta_1 + x_1\frac{\mathcal{R}_1}{\mathcal{V}_1} + x_2\beta_2 + x_2\frac{\mathcal{R}_2}{\mathcal{V}_2}\right) dt \\ &\quad - (x_1\gamma_1 + x_2\gamma_2)\, d\mathcal{Z}. \end{aligned} \tag{13}$$

Choosing $x_1 = -\gamma_2/(\gamma_1 - \gamma_2)$ and $x_2 = \gamma_1/(\gamma_1 - \gamma_2)$, we have $x_1\gamma_1 + x_2\gamma_2 = 0$. Thus, we have

$$\begin{aligned} d\Pi &= \left(x_1\beta_1 + x_1 \frac{\mathcal{R}_1}{\mathcal{V}_1} + x_2\beta_2 + x_2 \frac{\mathcal{R}_2}{\mathcal{V}_2} \right) dt \\ &= r\Pi dt \end{aligned} \tag{14}$$

where the second equality comes from the fact that, as Π is riskless, it must earn the risk free rate r. It follows that

$$\frac{\beta_1 + \frac{\mathcal{R}_1}{\mathcal{V}_1} - r}{\gamma_1} = \frac{\beta_2 + \frac{\mathcal{R}_2}{\mathcal{V}_2} - r}{\gamma_2}. \tag{15}$$

Define κ as the value of each side of (15), i.e.,

$$\frac{\beta_1 + \frac{\mathcal{R}_1}{\mathcal{V}_1} - r}{\gamma_1} = \frac{\beta_2 + \frac{\mathcal{R}_2}{\mathcal{V}_2} - r}{\gamma_2} = \kappa.$$

Dropping the subscripts, we have shown that if $\mathcal{V}$ is an investment dependent on $\mathcal{Q}$ and t, such that

$$d\mathcal{V} = (\beta\mathcal{V} + \mathcal{R})\, dt + \gamma\mathcal{V}\, d\mathcal{Z}$$

then

$$\beta - r + \frac{\mathcal{R}}{\mathcal{V}} = \gamma\kappa. \tag{16}$$

Substituting β from (11) and γ from (12) into (16), we find

$$\frac{\partial\mathcal{V}}{\partial t} + \frac{1}{2}\sigma^2\mathcal{Q}^2 \frac{\partial^2\mathcal{V}}{\partial\mathcal{Q}^2} + (\mu - \kappa\sigma)\frac{\partial\mathcal{V}}{\partial\mathcal{Q}} - r\mathcal{V} + \mathcal{R} = 0. \tag{17}$$

Appendix B
Estimation of Growth Rate and Volatility Parameters

We begin by considering the growth rate μ. We base our estimate on summary traffic data for each academic term (three per year) from 1997 to the present. We use this data so as to obtain an estimate of long-term growth in network traffic. In particular, we use the following procedure. Given measurements $\mathcal{Q}_1, \mathcal{Q}_2, \ldots, \mathcal{Q}_N$ at times $t_1, t_2, \ldots, t_N$, we estimate μ by least squares:

$$\min_{\mu} \left(\begin{bmatrix} u_1 \\ u_2 \\ \vdots \\ u_{N-1} \end{bmatrix} - \mu \begin{bmatrix} \delta t \\ \delta t \\ \vdots \\ \delta t \end{bmatrix} \right)^2$$

where N is the number of observations and $u_i = \log(\mathcal{Q}_{i+1}/\mathcal{Q}_i)$ for $i = 1, \ldots, N-1$. We find that $\mu = 0.75$ per year.

Our estimate of volatility is based on daily traffic data from mid-2000 until January 2002. In Fig. 3, the time series shows strong autocorrelation within each week. To filter out this effect, we average the network traffic for the entire period for each day of the week separately, and the day with the highest average total daily traffic (Thursday) is chosen. This same day is used each week to estimate week to week effects (see Fig. 4).

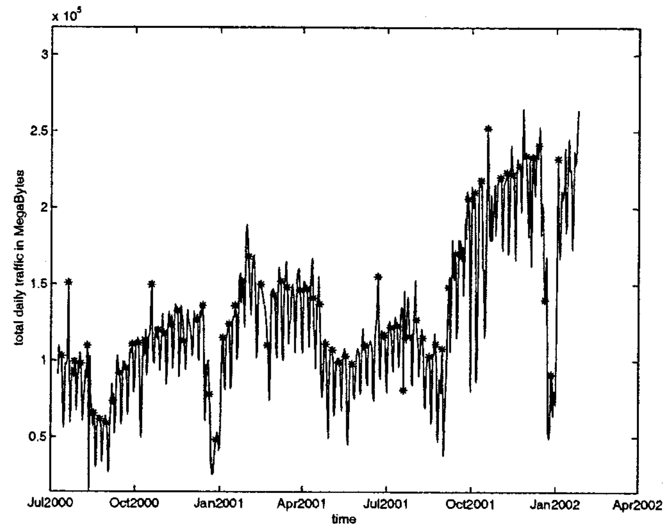

Fig. 4. Total daily network traffic into and out of the University of Waterloo since July 7, 2000. The stars indicate the weekly highest traffic days (Thursdays).

An inspection of Fig. 4 reveals several suspicious large drops in traffic, some for extended periods of time. Basing our volatility estimate on this data would produce a very high value. Some of the short-term declines are simply due to holidays. As these are known events for low network traffic, we should not take them into account when estimating volatility. Consequently, we smooth them out using linear interpolation. As the University of Waterloo has an extended break period in August, we remove this period entirely from the sample. There is also an extended drop around May 2001, which we delete from the sample. One might think that this is due to the start of the summer term, but Waterloo is a co-operative university with three full academic terms per year (the summer term running from May to the first part of August), so the effect of summer on network traffic is probably lower than at most other academic institutions. The final data sample used to estimate volatility is presented in Fig. 5.

Prior to computing the volatility, we first remove the growth trend μ from the data. We then conduct the Ljung-Box Q test [19] as a check that our smoothed and detrended data contains pure white noise. Each part of the time series (see Fig. 5) is studied independently. We can conclude from the p-values reported in Table X that there is no serial correlation in our data.

The volatility is finally estimated under the assumption that relative changes in network traffic are lognormally distributed, as implied by (1). We find that $\sigma = 0.95$ (as noted above in Section IV, this is in units of year$^{1/2}$).

Appendix C
Upgrade and Maintenance Costs

As noted in Section IV, data on upgrade costs were obtained from [14] (see Table IV). These were used to construct Table V.

As a rule of thumb, annual maintenance costs are usually between 3% and 10% of the upgrade capital cost. As mentioned earlier in the paper, maintenance costs do not include repeater costs. We assume that repeaters are situated every 100 miles. The maintenance cost for an OC-12 is computed as follows: we

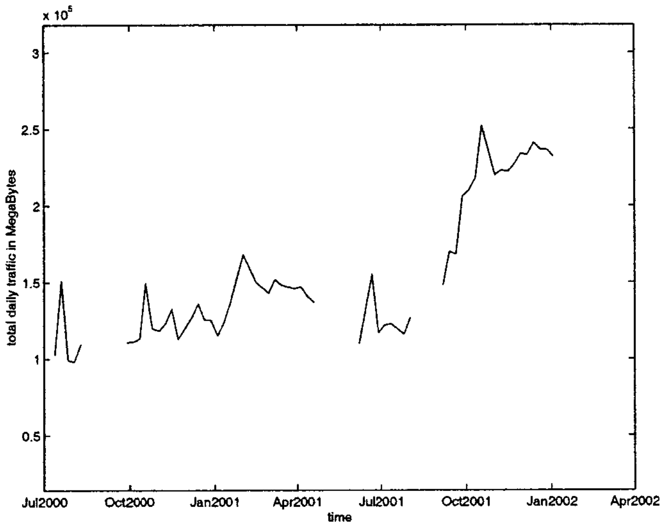

Fig. 5. Total daily network traffic into and out of the University of Waterloo since July 7, 2000. The data has been smoothed using linear interpolation around short term drops due to factors such as holidays. Extended declines in traffic in August 2000, May 2001, and August 2001 have been deleted.

TABLE X
LJUNG-BOX TEST FOR SERIAL CORRELATION. THE NUMBERS REPORTED ARE P-VALUES. THERE IS NO EVIDENCE OF SERIAL CORRELATION AT CONVENTIONAL STATISTICAL SIGNIFICANCE LEVELS. THE PARTS OF THE DATA (PRIOR TO DETRENDING) ARE SHOWN IN FIG. 5

Ljung-Box Test (p-values)				
Lag	1st part	2nd part	3rd part	4th part
1	0.180	0.265	0.168	0.221
2	0.354	0.485	0.386	0.208
3	0.512	0.675	0.592	0.252
4	0.680	0.768	0.753	0.143

take the upgrade cost from an OC-3 to an OC-12 (see Table IV) and divide it by 100 miles (fiber optic distance without any repeaters). We then take 6% of the resulting number. We thus have the following for an OC-12:

$$\frac{4000}{100} = \$40/\text{month/mile}$$
$$\mathcal{M}_{\text{OC}-12} = 40 \times 0.06 = \$2.40/\text{year/mile}.$$

Analogous calculations were made for the other lines in order to construct Table VI.

APPENDIX D
TELECOM MARKET PRICE OF RISK ESTIMATION

The price of dark fiber is estimated to vary between \$1000 and \$1500 per mile [15]. The cost of the initial circuit equipment investment is estimated to be $\approx$\$500 000 [14] for both ends of the line. As a rough estimate, and consistent with [17], we assume that we are using 50% of the available bandwidth. Thus, assuming that the average capacity of a line is equivalent to an OC-48 line (maximum transmission rate: 2488 Mb/s), we estimate that the level of demand is at $\mathcal{Q} = 1244$ Mb/s.

Although virtually all of the discussion in the paper relates to the timing of the upgrade decision, note that this is only one of the outputs of our model. The total value of the investment $\mathcal{V}$ is also determined in our computations. To calculate the telecom market price of risk, we calibrate the investment value given by our model (at $\mathcal{Q} = 1244$ Mb/s) to the investment cost in the market. The investment cost is obtained by multiplying the dark fiber price by the distance in miles between two city pairs and adding the initial investment. In our simulations, we considered the Toronto/New York City pair (550 miles apart). We obtain the telecom market price of risk κ by fitting the investment value given by our model to $550 \times \$1000 + \$500\,000 = \$1\,050\,000$. We find that $\kappa \approx 1$.

We choose to estimate κ by calibrating our model to observed market prices. This approach is standard in financial applications such as the pricing of bond derivatives [8]. This approach makes no assumption about market equilibrium. Alternatively, economic general equilibrium theory shows that the market price of risk is determined by the covariance between demand for network bandwidth and changes in aggregate wealth in the economy. As this covariance is difficult to estimate, we prefer the calibration approach as described above.

ACKNOWLEDGMENT

The authors would like to thank G. Kealey for many useful discussions.

REFERENCES

[1] Risk Books, *The Telecoms Trading Revolution*. London, U.K.: Risk Pubs., 2000.
[2] A. K. Dixit and R. S. Pindyck, *Investment Under Uncertainty*. Princeton, NJ: Princeton Univ. Press, 1994.
[3] L. Trigeorgis, *Real Options: Managerial Flexibility and Strategy in Resource Allocation*. Cambridge, MA: MIT Press, 1996.
[4] T. Dangl, "Investment and capacity choice under uncertain demand," *Eur. J. Oper. Res.*, vol. 117, pp. 415–428, 1999.
[5] H. Yaïche, R. R. Mazumdar, and C. Rosenberg, "A game theoretic framework for bandwidth allocation and pricing in broadband networks," *IEEE/ACM Trans. Networking*, vol. 8, pp. 667–678, Oct. 2000.
[6] S. H. Low, "Equilibrium bandwidth and buffer allocations for elastic traffics," *IEEE/ACM Trans. Networking*, vol. 8, pp. 373–383, June 2000.
[7] C. Kenyon and G. Cheliotis, "Stochastic models for telecom commodity prices," *Comput. Networks*, vol. 36, pp. 533–555, 2001.
[8] J. Hull, *Options, Futures, and Other Derivatives*, 4th ed. Upper Saddle River, NJ: Prentice-Hall, 2001.
[9] P. Wilmott, *Derivatives: The Theory and Practice of Financial Engineering*. West Sussex, U.K.: Wiley, 1998.
[10] R. Zvan, P. A. Forsyth, and K. R. Vetzal, "Robust numerical methods for PDE models of Asian options," *J. Comput. Finance*, vol. 1, pp. 39–78, Winter 1998.
[11] R. Zvan, P. A. Forsyth, and K. R. Vetzal, "A finite volume approach for contingent claims valuation," *IMA J. Numer. Anal.*, vol. 21, pp. 703–731, 2001.
[12] D. Kröner, *Numerical Schemes for Conservation Laws*. Chichester, U.K.: Wiley, 1997.
[13] Y. d'Halluin, P. A. Forsyth, K. R. Vetzal, and G. Labahn, "A numerical PDE approach for pricing callable bonds," *Appl. Math. Finance*, vol. 8, pp. 49–77, 2001.
[14] K. Martin, Hardware Manager, Mathematics Faculty Computing Facility, Univ. Waterloo, private communication, 2001.
[15] G. Kealey, Managing Consultant, Bell Nexxia, private communication, 2001.
[16] R. W. Watt. University of Waterloo campus network statistics. [Online]. Available: http://www.ist.uwaterloo.ca/cn/#Stats
[17] G. Stix, "The triumph of the light," *Scientif. Amer.*, vol. 284, no. 1, pp. 80–86, 2001.
[18] G. Williams, BCE Capital, private communication, 2001.
[19] C. Chatfield, *The Analysis of Time Series: An Introduction*, 4th ed. London, U.K.: Chapman and Hall, 1989.

Yann d'Halluin received an Engineering degree in scientific computation from the Engineering School of Science and Technology, Lyon Institute, Lyon, France in 1998 and the M.Math. degree in computer science from the University of Waterloo, Waterloo, ON, Canada, in 2000. He is currently working toward the Ph.D. degree in the Department of Computer Science of the University of Waterloo.

His research interests are primarily in the area of numerical valuation of complex real and financial options.

Peter A. Forsyth received the B.Sc. degree from the University of Western Ontario, London, ON, Canada, in 1975, the M.Sc. degree from the Australian National University, Canberra, ACT, Australia, in 1977, and the Ph.D. degree from the University of Western Ontario in 1979.

Beginning in 1979, he was a Senior Simulation Scientist with the Computer Modeling Group (CMG), Calgary, Canada, where he developed petroleum reservoir simulation software. After leaving CMG, he was the founding President of Dynamic Reservoir Systems (DRS), also in Calgary. In 1987, he joined the University of Waterloo, Waterloo, ON, Canada, where he is currently a Professor in the Department of Computer Science. His research interests include numerical algorithms for pricing options, numerical solution of partial differential equations, and iterative methods for sparse matrices.

Kenneth R. Vetzal received the M.A. degree in economics and the Ph.D. degree in finance from the University of Toronto, Toronto, ON, Canada, in 1983 and 1992, respectively.

He is currently an Associate Professor with the Centre for Advanced Studies in Finance, University of Waterloo, Waterloo, ON, Canada. His research interests are primarily in the area of valuation and hedging of complex derivative instruments, capital budgeting, and the term structure of interest rates.