Wireless Network Capacity Investment

Y. d'Halluin[*], P.A. Forsyth[†], and K.R. Vetzal[‡]

Draft Version: February 26, 2003

Abstract

This paper applies modern financial option valuation methods to the problem of new wireless network capacity investment decision timing. In particular, given a cluster of base stations (with a certain traffic capacity per base station), we determine when it is optimal to increase capacity for each of the base stations contained in the cluster. We express this in terms of the fraction of total cluster capacity in use, i.e. we calculate the optimal time to upgrade in terms of the ratio of observed usage to existing capacity. We study the optimal decision problem of adding new capacity in the presence of stochastic wireless traffic for services. We develop a four factor algorithm that captures all of the constraints of wireless network management, based on a real options formulation. We study the upgrade decision for different upgrade decision intervals (e.g. monthly, quarterly, etc.), and we investigate the effect of a safety level (i.e. the maximum allowed capacity used in practice on a daily basis—which differs from the theoretical maximum).

1 Introduction

Wireless telephones are now regarded as essential communication tools, dramatically impacting how people approach personal and business communications. As new network infrastructure is built and competition between wireless carriers increases, digital wireless subscribers are becoming ever more critical of the service and voice quality they receive from network providers. Wireless operators must provide a guaranteed level of service to customers while maximizing profit. The current environment, with decreasing revenue per minute and increasing demands on networks from new features in wireless equipment, places conflicting demands on network managers [10].

Traditional wireless network management is based on experience and heuristics. While these methods often appear to work well in practice, theoretical work is needed to evaluate the generally agreed upon approaches. Previous work on in the related area of bandwidth network management [3] presented some interesting numerical methods and results, but there was little high quality data which could be used to estimate parameters. The lack of high quality data can be partly attributed to the relative infancy of the bandwidth market. On the other hand, wireless networks have been in place for quite some time, and better data can be obtained. In addition, the algorithm developed in [3] could only handle the case where decision date intervals were equivalent to the time period required to order, install and test new equipment.

[*]Y. d'Halluin is a Ph.D. student with the School of Computer Science at the University of Waterloo, Waterloo ON, Canada N2L 3G1 (e-mail: `ydhallui@elora.uwaterloo.ca`).

[†]P.A. Forsyth is a professor with the School of Computer Science at the University of Waterloo, Waterloo ON, Canada N2L 3G1 (e-mail: `paforsyt@elora.uwaterloo.ca`).

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

In this paper we apply a real options framework to the problem of the optimal timing of investment into new capacity and extend the algorithm of [3] to handle arbitrary decision date intervals. Given a cluster of base stations, our objective is to find the percentage (in terms of the ratio of observed usage to existing capacity) at which it is optimal to add capacity by installing one or more carriers to each of the base stations in the cluster. This optimal upgrade decision will maximize the value of the investment to the network operator. The remainder of this paper is organized as follows: Section 2 describes the modeling framework; Section 3 presents the mathematical model and the decision to upgrade algorithm; Section 4 provides the estimated model parameters; and Section 5 contains various simulated results. Conclusions are given in Section 6.

2 Background

As in the bandwidth market [3], the revenue to the owner of a wireless network is determined by the prevailing price per minute and the amount of traffic. Average revenue per wireless user is decreasing with relatively little uncertainty (Figure 1). However, an examination of wireless network traffic reveals some interesting features. Although traffic has obviously been increasing, it has done so in a somewhat uneven way, with apparent randomness. This is shown for an illustrative switch in Figure 2. In Appendix A, we show that not all network traffic movements can be attributed to deterministic drift and cyclical patterns. When cycles, trends, anomalous drops and statutory holidays are removed, large volatility in network traffic remains.

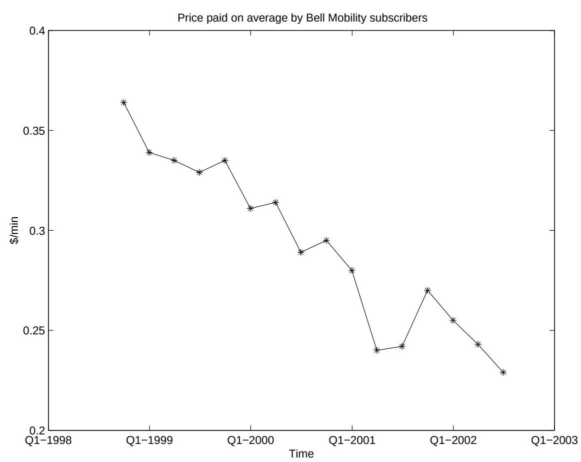

FIGURE 1: *Price paid on average by wireless network subscribers in $ per minute. The price is obtained by dividing the average revenue per user by the usage per subscriber. The price decline can be fit reasonably well by the function* $\mathcal{P}(t) = \mathcal{P}_0 \exp(-\mu t)$, *where* $\mu = -.08/year$. *These data were obtained from Bell Canada quarterly financial reports.*

Traditionally, for a given set of base stations (i.e. a *cluster*) and a desired *grade of service/blocking probability*, the traffic engineer predicts the amount of capacity (or the number of carriers) necessary to satisfy the given demand while maximizing revenue. When there is high traffic, the base stations experience very high *blocking* and new equipment must be installed at the base station level (i.e. *carriers*). Figure 3 provides a representation of a simplified cluster of base stations/cell sites. A typical cluster contains at least 20 cell sites. Each cell site has a certain coverage area that is divided into sectors. A simplistic solution to blocking would be to conduct a traffic study at the cell site level and increase the capacity of the cells that

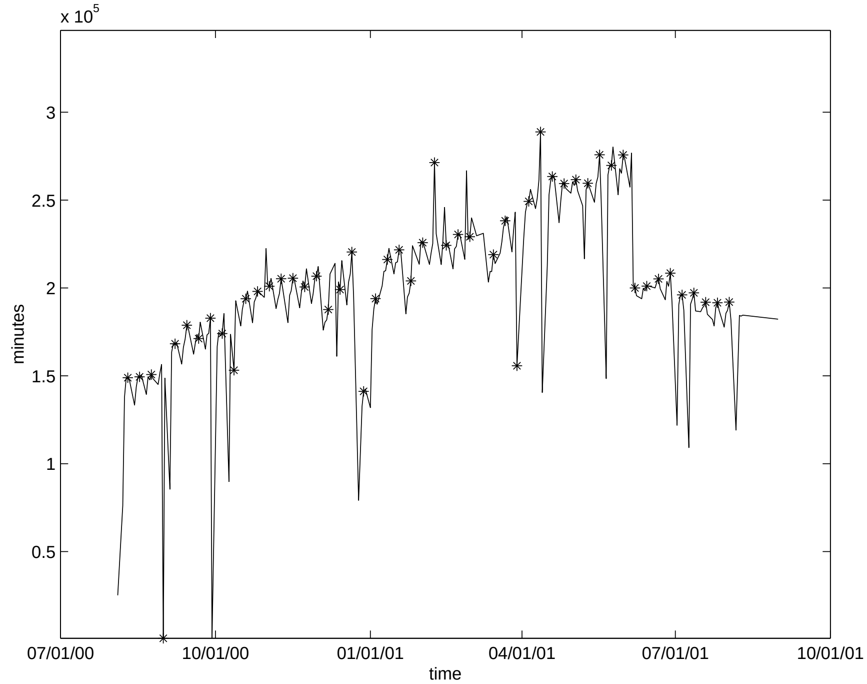

FIGURE 2: *Daily bouncing busy hour data traffic on a representative switch. Data obtained from [9]. To remove weekly cycles, the weekday with the most traffic on average in a year is selected. The stars are used to indicate this day. In our data, the day with highest average traffic turned out to be Thursday, so we use weekly data sampled every Thursday to estimate the volatility* σ *and the growth rate* μ*.*

are experiencing too much blocking.

However, due to the Code Division Multiplexing Access (CDMA) [7] technology that is currently used in leading edge wireless networks, it is not possible to only add carriers to the cell sites where high blocking occurs. A user on a network using CDMA technology may talk simultaneously to many base stations since a principle called *soft hand-off* is used [7]. As such, when there is too much blocking on a particular cell site, all the cell sites within the cluster must have a new carrier added to maintain homogeneity.

When studying the number of traffic minutes before a cell site starts blocking, the following factors should be considered:

- The number of carriers currently deployed.
- Whether efficiencies arising from pooling resources of each carrier can be realized, i.e. the ability to allocate mobile traffic to the least busy carrier.
- Currently, carriers are grouped into blocks of three, meaning traffic can be allocated evenly for three carriers. However, for the fourth carrier, a new pooling group must be formed. For this study, we will assume that the traffic is homogeneous throughout the cluster and that the carriers are a shared resource.

With these factors in mind, we can consider traffic and grade of service/blocking probability. We must also define the unit of measurement for traffic. Traffic is measured in units of *Erlangs* [8]. An Erlang is defined as the average number of simultaneous calls, or equivalently, the total usage during a time interval

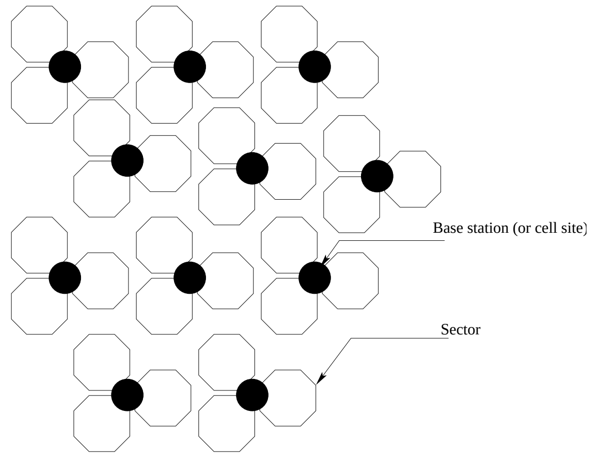

FIGURE 3: *An example of a cluster of cell sites. A typical cluster contains at least 20 base stations. The coverage area or sector is generally not hexagonal, but this diagram gives an idea of how cell sites are deployed.*

divided by the length of that interval. Most network management systems measure usage during a one hour interval. The *blocking probability* is the probability that a call is blocked when there is no channel available. The blocking probability is evaluated for the load during *bouncing busy hours*. In practice, network managers study the load during bouncing busy hours and then decide whether or not new carriers must be deployed. The calculation used to determine the bouncing busy hour is as follows:

1. For each sector of every cell site, determine which hour had the most traffic for a given day. This is called the *busiest bouncing hour*. This hour will most likely be different for each sector. Note that the hour begins at the top of the hour, i.e. the busiest hour is not the 60 consecutive minutes where traffic is the highest.

2. For a given carrier, sum the number of minutes during the busiest hour for every sector of every cell site. This number will be the bouncing busy hour traffic for that switch on that particular day.[1]

Based on the Erlang mathematical models (Erlang-B or Erlang-C [8]), the relationship between blocking probability and demand can be established. See [8] for details.

In this work we assume that a single carrier can handle approximately 20 users. Using Erlang tables, this means 13.2 Erlang of traffic can be handled at 2% blocking. If a second carrier is added, 40 users can be accommodated, meaning 31 Erlang of traffic can be handled at 2% blocking. The carried load at 2% blocking more than doubles when a new carrier is added, even though the maximum load only doubles. Table 1 gives a synopsis of an Erlang table.

3 Mathematical Model

As noted above, traffic data on a typical switch appears to contain a random component (see Figure 2). A simple model for network traffic is geometric Brownian motion. Letting Q represent the bouncing busy hour network traffic in minutes per hour, we have

$$dQ = \mu Q dt + \sigma Q dz \tag{3.1}$$

[1] In telecommunications, a switch is a network device that selects a path or circuit for sending a unit of data to its next destination.

Users	Erlang	Minutes of traffic in the hour
20	13.2	792
40	31.0	1860
60	49.6	2976

TABLE 1: *Correspondence table between users, Erlang and minutes of traffic in the hour that can be approximately handled by a single sector of a carrier at 2% blocking. Carriers usually cover three sectors.*

where μ is the drift or growth rate, σ is the volatility, and dz is the increment of a Wiener process. Equation (3.1) implicitly assumes that traffic is lognormally distributed. Note that if $\sigma = 0$, then equation (3.1) implies exponential growth at rate μ, i.e. $Q = Q_0 \exp(\mu t)$ where Q_0 is a constant. Because of their scaling with time, the volatility and drift terms have different effects. For short time periods, the volatility (uncertainty) term will dominate, while for longer time periods the drift term becomes important. Appendix A provides details on our estimation of the parameters μ and σ of equation (3.1) based on traffic data time series.

In order to determine the optimal time for an equipment upgrade, we must determine the value $\mathcal{V}$ of an investment in equipment for various level of capacity. $\mathcal{V}$ depends on the following factors:

1. Actual demand for service is uncertain and given by the stochastic differential equation (SDE) (3.1).
2. Upgrade decisions can be made at various times. There is a time lag between when the equipment is ordered and when it comes on-line.
3. Revenues, capital costs, and maintenance costs must be taken into account correctly.

Our goal is to determine the optimal course of action so as to maximize the value of the investment, in light of the fact that actual usage of the network in the future is uncertain. The procedure is similar to that used in pricing financial options. We will employ a dynamic programming approach, and solve for the optimal policy by proceeding backwards in time.

We define a set of observation times $t_{obs} = \{0, \Delta t_{obs}, 2\Delta t_{obs}, \ldots\}$. At these times, we assume that any of the following events take place:

- Maintenance costs are paid.
- Partial or complete payments are made for capital expenditures.
- Decisions about possible upgrades are made.
- Upgrades come on-line.

Typically, we will take $\Delta t_{obs} = 1$ month. We assume that the set of times when possible upgrade decisions are made is discrete and denote this set by t_{up}. If $t_{up}^{\alpha} \in t_{up}$, then $t_{up}^{\alpha}/\Delta t_{obs}$ is an integer.

If a decision is made to upgrade at t_{up}^{α}, then the actual upgrade is completed at $t_{up}^{\alpha} + \gamma$. We assume that $\gamma/\Delta t_{obs}$ is an integer. Note that γ corresponds to the time necessary to order and set up the equipment.

Since we use a dynamic programming approach to find the optimal policy, we work backwards in time, so that at any time $t_{obs}^{\beta} \in t_{obs}$ we cannot know when a decision was made to upgrade to a higher level of capacity. Hence, we have to solve for all possible times at which an upgrade could occur. Consequently, we need an additional discrete state variable $\mathcal{F}_l$ for $l = 0, \ldots, l_{\max} - 1$, which we define as

$$\begin{aligned} \mathcal{F}_l = \ & \left(\text{elapsed time since last decision to upgrade}\right)/\Delta t_{obs} + 1 \\ & \text{as observed the instant after the previous observation date.} \end{aligned}$$

For example, assume that $\gamma = 3$ months and $\Delta t_{obs} = 1$ month. Suppose that $t \in [12 \text{ months} + \varepsilon, 13 \text{ months} - \varepsilon]$. Then:

- $\mathcal{F}_1 = 1 \rightarrow$ a decision to upgrade was made at $t = 12$ months. The new equipment will come on-line at $t = 15$ months.
- $\mathcal{F}_2 = 2 \rightarrow$ a decision to upgrade was made at $t = 11$ months. The new equipment will come on-line at $t = 14$ months.
- $\mathcal{F}_3 = 3 \rightarrow$ a decision to upgrade was made at $t = 10$ months. The new equipment will come on-line at $t = 13$ months.

We need to keep track of the maximum capacity of each cluster, which we assume is a discrete variable denoted by $\overline{Q}_j$, where $j = 0, \ldots, j_{\max} - 1$. In addition, we will allow different levels of upgrades to occur. In other words, we can upgrade from $\overline{Q}_j$ to $\overline{Q}_u$, where $u \in \{j+1, \ldots, j_{\max} - 1\}$. Note that $u = j$ corresponds to an existing cluster of capacity $\overline{Q}_j$ where no decision to upgrade has been made. By convention $\mathcal{F}_{l=0} = 0$ when $u = j$. Hence, $\mathcal{V}(Q, \overline{Q}_j, j, 0, \tau)$ represents a cluster of capacity $\overline{Q}_j$ where no decision to upgrade has been made.

Consequently, the value of an investment is given by $\mathcal{V} = \mathcal{V}(Q, \overline{Q}_j, u, \mathcal{F}_l, \tau)$. This represents the value of an investment with

- Maximum capacity $\overline{Q}_j$,
- A decision made to upgrade to capacity $\overline{Q}_u$,
- The elapsed time since the decision to upgrade was made (as observed at the instant after previous observation date) is $\mathcal{F}_l$.

The above holds for all times $\tau = T - t \neq \tau_{obs}$, where $\tau_{obs} = T - t_{obs}$. Furthermore, we must carefully distinguish between the instant before and after observation times, since $\mathcal{F}_l$ is incremented at these times. We assume that once a decision to upgrade has been made, it cannot be reversed. Note that this assumption can be easily changed in our model, but irreversibility appears to be consistent with practice (i.e. equipment is rarely if ever removed or downgraded once it has been deployed).

At each observation time t^{α}_{obs}, the value of $\mathcal{F}_l$ will be changed. Let $t^+ = t^{\alpha}_{obs} + \varepsilon$, $t^- = t^{\alpha}_{obs} - \varepsilon$. Then

$$\mathcal{F}_l^+ = \mathcal{F}_{l+1}^-, \tag{3.2}$$

since the elapsed time will be incremented by one. Absence of arbitrage implies

$$\mathcal{V}(Q, \overline{Q}_j, u, \mathcal{F}_l^+, t^+) = \mathcal{V}(Q, \overline{Q}_j, u, \mathcal{F}_l^-, t^-), \tag{3.3}$$

or

$$\mathcal{V}(Q, \overline{Q}_j, u, \mathcal{F}_{l+1}^-, t^+) = \mathcal{V}(Q, \overline{Q}_j, u, \mathcal{F}_l^-, t^-). \tag{3.4}$$

Since we work backwards in time, let $\tau^+ = t^-, \tau^- = t^+$. Thus

$$\mathcal{V}(Q, \overline{Q}_j, u, \mathcal{F}_{l+1}^-, \tau^-) = \mathcal{V}(Q, \overline{Q}_j, u, \mathcal{F}_l^-, \tau^+). \tag{3.5}$$

But we need to account for the cost of upgrading in equation (3.5). Let $\mathcal{C}_{j \rightarrow u}(t)_l$ denote some designated fraction of the cost of upgrading from a cluster of maximum capacity $\overline{Q}_j$ to one with maximum capacity $\overline{Q}_u$[2]. Then we subtract the cost of the upgrade from (3.5)

$$\mathcal{V}(Q, \overline{Q}_j, u, \mathcal{F}_l^-, \tau^+) = \mathcal{V}(Q, \overline{Q}_j, u, \mathcal{F}_{l+1}^-, \tau^-) - \mathcal{C}_{j \rightarrow u}(\tau^-)_l, \tag{3.6}$$

[2] The designated fraction can be specified in a variety of ways. For instance, with four months lead time, one quarter of the cost could be paid in each of the four months. Alternatively, all of the cost could be paid up front.

where $l = 1, \ldots, l_{\max} - 2$.

At $l = l_{\max} - 1$, we set

$$\mathcal{V}(Q, \overline{Q}_j, u, \mathcal{F}^-_{l_{\max}-1}, \tau^+) = \mathcal{V}(Q, \overline{Q}_u, 0, 0, \tau^-) - \mathcal{C}_{j\to u}(\tau^-)_{l_{\max}-1}. \tag{3.7}$$

Equation (3.7) indicates that it is possible to upgrade from $\overline{Q}_j$ to $\overline{Q}_u$ but the new equipment will not be ready for some time.

At each upgrade date, t_{up}, we maximize the value of the investment $\mathcal{V}$. For a given cluster j, we have for $l = 0$

$$\mathcal{V}(Q, \overline{Q}_j, j, 0, \tau^+_{up}) = \max\left(\mathcal{V}(Q, \overline{Q}_j, j, 0, \tau^-_{up}), \mathcal{V}(Q, \overline{Q}_j, u, \mathcal{F}_1, \tau^-_{up}) - \mathcal{C}_{j\to u}(\tau^-_{up})_0\right), \tag{3.8}$$

for $u = j+1, \ldots, j_{\max} - 1$. Equation (3.8) indicates that a decision to add capacity will only be justified if the value of the investment exceeds the value of not investing.

Based on standard hedging arguments, a partial differential equation for the value of an investment $\mathcal{V}(Q, \overline{Q}, u, \mathcal{F}, \tau)$ where cash flows are a function of Q is found to be

$$\frac{\partial \mathcal{V}}{\partial t} + \frac{1}{2}\sigma^2 Q^2 \frac{\partial^2 \mathcal{V}}{\partial Q^2} + (\mu - \kappa\sigma)\, Q \frac{\partial \mathcal{V}}{\partial Q} - r\mathcal{V} + \mathcal{R}(Q, \overline{Q}, t) = 0, \tag{3.9}$$

where $\mathcal{V}$ is the value of the investment in dollars, $\mathcal{R}(Q, \overline{Q}, t)$ is the revenue term in dollars per year, r is the risk free interest rate and κ is the market price of risk. Informally, the tradeoff between the risk of an investment that depends on Q and its anticipated return is captured by κ. In Appendix B, we describe how we estimate the market price of risk κ from market data.

In the valuation of financial options, the value of the option at the expiry date is a known function of the underlying stock price. However, the value of the option prior to expiry is not known, but may be found by solving a partial differential equation similar to (3.9). In our case, we consider an investment horizon T (analogous to the expiry date of a financial option). Mathematically, we then have

$$\mathcal{V}(Q, \overline{Q}, u, \mathcal{F}, \tau) = f(Q).$$

Although the methods discussed in this paper can be used with any suitable choice of $f(Q)$, for simplicity we will restrict our attention to the case where the value of all capital investment at $t = T$ is equal to the salvage value of the upgraded equipment. We will assume that $f(Q) = 0$, i.e. the salvage value at $t = T$ is zero. We choose the investment horizon to be $T = 5$ years. Implicitly, we assume that new technology will render all existing equipment obsolete at $t = T$.

Since the value of the investment is known at $t = T$, the forward equation (3.9) is transformed into a backward equation by substituting $\tau = T - t$ to give

$$\frac{\partial \mathcal{V}}{\partial \tau} = \frac{1}{2}\sigma^2 Q^2 \frac{\partial^2 \mathcal{V}}{\partial Q^2} + (\mu - \kappa\sigma)\, Q \frac{\partial \mathcal{V}}{\partial Q} - r\mathcal{V} + \mathcal{R}(Q, \overline{Q}, \tau). \tag{3.10}$$

Define $\mathcal{V}(Q, \overline{Q}_j, u, \mathcal{F}_l, \tau)$ to be $\mathcal{V}(Q, \tau)_{j,u,l}$ where the indices $j = 0, \ldots, j_{\max} - 1$, $u = j+1, \ldots, j_{\max} - 1$, and $l = 0, \ldots, l_{\max} - 1$. $\mathcal{F}_l$ represents the elapsed time since the decision to upgrade was made. A set of equations must be solved for each possible cluster capacity $\overline{Q}_j$, $j = 0, \ldots, j_{\max} - 1$, where $j_{\max} - 1$ is the maximum number of cluster capacities. This can be accomplished using a general numerical partial differential equation (PDE) solver [13, 14]. The numerical PDE approach involves a finite volume discretization of equation (3.10) along the axis representing the demand Q [6]. The finite volume method has been extensively studied in [13, 14]. As it is beyond the scope of this paper, we will not present the details of the discretization scheme. Interested readers should see [4, 6].

Consider a set of clusters with a maximum capacity $\overline{Q}_j$. For example, $\overline{Q}_1$ could be a one-carrier cluster (i.e. a cluster where there is only one carrier per cell site), $\overline{Q}_2$ a two-carrier cluster, and so on (see Figure 4). Let $\mathcal{V}(Q,\tau)_{j,u,l}$ be the value of an investment in a cluster with maximum capacity $\overline{Q}_j$. We must solve a set of PDEs (3.10) for each upgrade possibility, i.e.

$$\frac{\partial \mathcal{V}_{j,u,l}}{\partial \tau} = \frac{1}{2}\sigma^2 Q^2 \frac{\partial^2 \mathcal{V}_{j,u,l}}{\partial Q^2} + (\mu - \kappa\sigma)\, Q \frac{\partial \mathcal{V}_{j,u,l}}{\partial Q} - r\mathcal{V}_{j,u,l} + \mathcal{R}(Q, \overline{Q}_j, t). \tag{3.11}$$

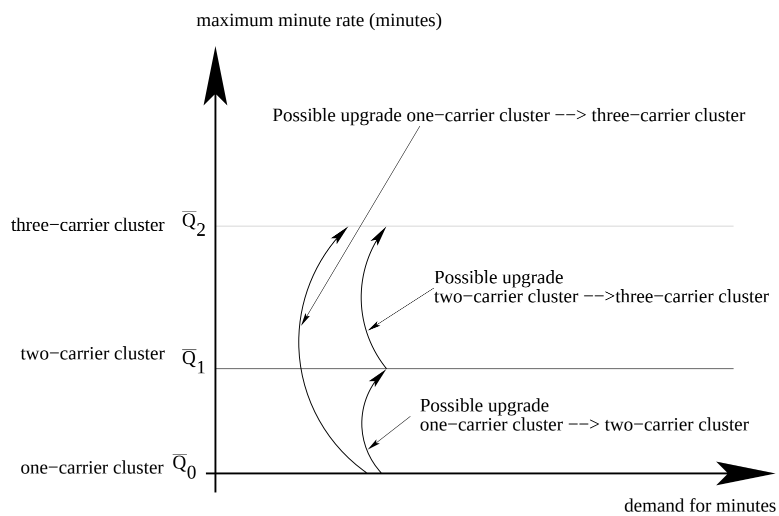

FIGURE 4: *For a particular plane with fixed l, we consider a set of clusters Q_j with maximum capacity $\overline{Q}_j$. We solve a set of PDEs* (3.11) *for each upgrade possibility.*

Let N be the number of nodes along the axis representing the discrete values of demand Q. Each time step of the solution requires solving approximately $(l_{\max}-1)(j_{\max}-1)(j_{\max}-2)/2$ one dimensional problems of size N.

3.1 Revenue

We assume that the owner of the cluster receives continuous payments. For a maximum capacity $\overline{Q_j}$ (in minutes per bouncing busy hour), we have

$$\mathcal{R}(Q, \overline{Q}_j, \tau) = \min(Q, \overline{Q}_j)\mathcal{P}(\tau). \tag{3.12}$$

The function $\mathcal{P}(\tau)$ is given by

$$\mathcal{P}(\tau) = \mathcal{P}_0 \exp(-\alpha(T-\tau)), \tag{3.13}$$

where $\mathcal{P}_0 Q$ has the units of dollars per year and α is a decay parameter. The payment received can be no larger than the maximum capacity of the cluster multiplied by the price. We assume that the price is a known decreasing function of time. Note that this does not create an arbitrage opportunity because unused minutes cannot be stored for later use.

The data underlying Figure 1 shows that today the average revenue per user is approximatively .229$/min. However, this value corresponds to non-marginal revenue received based on daily traffi c, and not bouncing busy hour traffi c. Consequently, to estimate $\mathcal{P}_0$ we need to adjust the average revenue per user (ARPU) appropriately. Based on estimates and discussions with industry personnel, marginal revenue represents seventy percent of the average revenue per user $.229 \times .7 \approx .1603$ [9]. Our studies indicate that bouncing busy hour represents approximatively 10% of total daily traffic. We assume that zero revenue is obtained on weekends. Consequently, $\mathcal{P}_0$ is given by

$$\begin{aligned}\mathcal{P}_0 = & \text{ Max minutes at 2\% blocking per hour} \\ & \times 10 \text{ (busy hr represents 10\% of the total daily traffic)} \\ & \times .7 \text{ (marginal revenue)} \times \text{ARPU} \times 250 \text{(excludes weekends)},\end{aligned}$$

where $\mathcal{P}_0$ has units $\left[\frac{\text{dollars}}{\frac{\text{minutes}}{\text{bbhr}}\text{year}}\right]$. The total revenue per year for various number of carriers per cell site for the cluster is given in Table 2.

Furthermore, we require that the adjusted revenue per user $\mathcal{P}_0$ be the same at all points in a cluster so as to avoid arbitrage. For example, if $\mathcal{P}_0$ of a three-carrier cluster was less than that of a two-carrier cluster, we could buy minutes on the three-carrier cluster and then immediately sell it at the two-carrier cluster spot price making a free profit. Consequently, equation (3.11) becomes

$$\frac{\partial \mathcal{V}_{j,u,l}}{\partial \tau} = \frac{1}{2}\sigma^2 Q^2 \frac{\partial^2 \mathcal{V}_{j,u,l}}{\partial Q^2} + (r - \kappa\sigma)\, Q \frac{\partial \mathcal{V}_{j,u,l}}{\partial Q} - r\mathcal{V}_{j,u,l} + \min(Q, \overline{Q_j})\mathcal{P}_0 \exp(-\alpha(T-\tau)). \tag{3.14}$$

Number of carriers per cell site	Revenue for the cluster in \$ per year
1	19,043,640
2	44,723,700
3	71,557,920

TABLE 2: *Maximum total revenue per year for the cluster based on bouncing busy hour traffic.*

3.2 Maintenance Costs

A cluster will have some unavoidable maintenance costs. These maintenance costs are assumed to be constant over time and paid at discrete time intervals Δt_{maint} (e.g. monthly). Given a cluster j with maximum capacity $\overline{Q_j}$, we have

$$\mathcal{V}_{j,u,l}(Q, \tau^+_{maint}) = \mathcal{V}_{j,u,l}(Q, \tau^-_{maint}) - \mathcal{M}_j \Delta\tau_{maint}, \tag{3.15}$$

where $\mathcal{M}_j$ is the maintenance cost in \$/year. Since maintenance costs are paid at discrete time intervals, it is important to appropriately determine if the maintenance costs are paid before or after the partial upgrades are made (3.7). We assume that $\Delta t_{maint} = \Delta t_{obs}$. However, we assume that, going forward in time, maintenance costs are paid before upgrade decisions are made. If $t^+ = t^{\alpha}_{obs} + \varepsilon$, and $t^- = t^{\alpha}_{obs} - \varepsilon$ then

$$\begin{aligned} t^+_{maint} &= (t^{\alpha}_{obs} - \varepsilon) - \varepsilon = t^- - \varepsilon, \\ t^-_{maint} &= t^+_{maint} - \varepsilon. \end{aligned}$$

Consequently, going backwards in time, upgrade decisions are made before maintenance fees are paid.

3.3 Upgrade Decision

We assume that while new equipment is ordered and tested, the current stream of revenue is not interrupted. In other words, there is no down time. Using our dynamic programming approach, we solve the PDEs

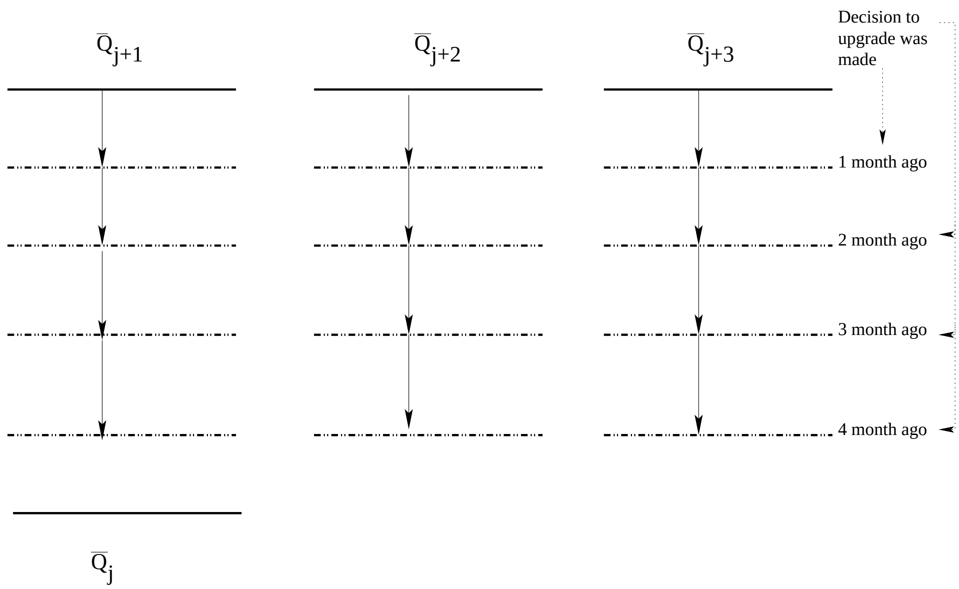

FIGURE 5: *At each observation date* τ_{obs}, *update the solution of the intermediate problems between the cluster of maximum capacity* $\overline{Q}_j$ *and* $\overline{Q}_u$ *where* $\overline{Q}_j < \overline{Q}_u$. *This figure shows the case where we are considering upgrading to 1, 2, or 3 higher levels. The solid lines represent the capacity, while the dotted lines represent the times at which a decision to upgrade to a higher capacity was made (e.g. one month ago, two months ago, etc). In this example, we assume there is a four month period between when the upgrade decision is made and when the new equipment is available for use.*

(3.11) backwards in time (τ increasing) and determine the optimal decision at each upgrade decision date τ_{up}. Investments are assumed to be decided upon at the beginning of each month.

Consider the clusters ordered as $j = 0, \ldots, j_{\max} - 1$ where $\overline{Q}_{j+1} > \overline{Q}_j$ (see Figure 4):

- **Stage I**: At each observation date τ_{obs}, update the solution of the intermediate problems between the cluster of maximum capacity $\overline{Q}_j$ and $\overline{Q}_u$ where $\overline{Q}_j < \overline{Q}_u$. Figure 5 provides a graphical representation of the execution of Algorithm 1 at each observation date. The pseudo-code for this stage is given in Algorithm 1.

- **Stage II**: At each upgrade decision date τ_{up}, compare the solution $\mathcal{V}(Q, \tau)_{j,0,0}$ with the solution of the partial investment into a cluster with higher capacity. Figure 6 provides a graphical representation of the execution of algorithm 2 at each upgrade decision date. The pseudo-code for this stage is given in Algorithm 2.

- **Stage III** At each maintenance date τ_{maint}, the maintenance costs are paid according to (3.15). The pseudo-code for this stage is given in Algorithm 3.

- **Stage IV** If $\tau = T$, terminate. Otherwise, solve the PDE for each problem to the next observation date and repeat the above process.

Algorithm 1 At each observation date τ_{obs}, update the solution of the intermediate problems between the cluster of maximum capacity $\overline{Q}_j$ and $\overline{Q}_u$.

for $j = 0, \ldots, j_{\max} - 2$ **do**
  // loop over the different clusters
  for $u = j+1, \ldots, j_{\max} - 1$ **do**
    // loop over the upgrade clusters possibilities for cluster j
    for $l = 1, \ldots, l_{\max} - 2$ **do**
      // copy the solution of the intermediate plane above to below
      $\mathcal{V}_{j,u,l}(Q, \tau^+_{obs}) = \mathcal{V}_{j,u,l+1}(Q, \tau^-_{obs}) - \mathcal{C}_{j\to u}(\tau^-_{obs})_l$
    end for
    at $l = l_{\max} - 1$
    $\mathcal{V}_{j,u,l_{\max}-1}(Q, \tau^+_{obs}) = \mathcal{V}_{u,u,0}(Q, \tau^+_{obs}) - \mathcal{C}_{j\to u}(\tau^-_{obs})_{l_{\max}-1}$
  end for
end for

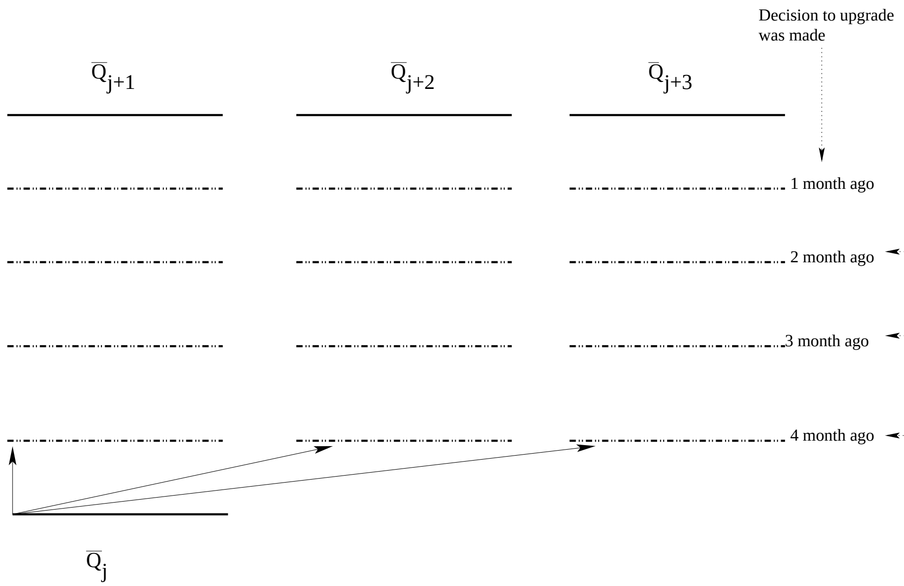

FIGURE 6: *At each upgrade decision date τ_{up}, compare the solution $\mathcal{V}(Q,\tau)_{j,0,0}$ with the solution of the complete investment into a cluster with higher capacity $\mathcal{V}(Q,\tau)_{j,u,1}$. This figure shows the case when we are considering upgrading to 1, 2, or 3 higher levels. The solid lines represent the capacity, while the dotted lines represent the time at which a decision to upgrade to a higher capacity was made (e.g. one month ago, two months ago, etc.). In this example, we assume there is a four month period between when the upgrade decision is made and when the new equipment is available for use.*

3.4 Upgrade Costs

In Algorithm 1, equation (3.6) and equation (3.7), $C_{j\to u}(\tau^-_{obs})_l$ is a fraction of the upgrade cost from cluster of maximum capacity $\overline{Q}_j$ to a cluster of maximum capacity $\overline{Q}_u$. The average revenue per user (ARPU) charged for usage of the wireless network is assumed to be declining over time. Upgrade costs are assumed to follow the same decreasing pattern as ARPU. Thus we will use the same decay factor α as for ARPU (see

Algorithm 2 At each upgrade decision date τ_{up}, compare the solution $\mathcal{V}(Q,\tau)_{j,0,0}$ with the solution of the partial investment into a cluster with higher capacity.

```
for j = 0,...,j_max − 2 do
  // loop over the cluster with capacity j
  for u = j+1,...,j_max − 1 do
    // loop over the upgrade clusters possibilities for cluster j
    V j,j,0(Q,τ_up^+) = max( V_{j,j,0}(Q,τ_up^−), V_{j,u,1}(Q,τ_up^−) − C_{j→u}(τ_up^−)_1 )
  end for
end for
```

$$\mathcal{V}j,j,0(Q,\tau_{up}^{+}) = \max\left(\mathcal{V}_{j,j,0}(Q,\tau_{up}^{-}),\mathcal{V}_{j,u,1}(Q,\tau_{up}^{-}) - \mathcal{C}_{j\to u}(\tau_{up}^{-})_1\right)$$

Algorithm 3 At each maintenance date τ_{maint}, the maintenance costs are paid.

```
for j = 0,...,j_max − 1 do
  // loop over the different clusters
  for u = j+1,...,j_max − 1 do
    // update the solution with the appropriate maintenance cost
    for l = 0,...,l_max − 1 do
      V_{j,u,l}(Q,τ_maint^+) = V_{j,u,l}(Q,τ_maint^−) − M_j Δτ_maint
    end for
  end for
end for
```

$$\mathcal{V}_{j,u,l}(Q,\tau_{maint}^{+}) = \mathcal{V}_{j,u,l}(Q,\tau_{maint}^{-}) - \mathcal{M}_j\Delta\tau_{maint}$$

equation (3.13)), i.e.

$$\mathcal{C}_{j\to u}(t)_l = \hat{\mathcal{C}}_{j\to u}\exp(-\alpha(t - \text{elapsed time since start of upgrade})), \tag{3.16}$$

where $\hat{\mathcal{C}}_{j\to u}$ is the initial fraction of the upgrade cost from a cluster of maximum capacity $\overline{Q}_j$ to a cluster of maximum capacity $\overline{Q}_u$.

4 Parameter Values

This section briefly describes the estimation of parameter values. Details regarding estimates of the growth rate μ and the volatility σ can be found in Appendix A. The volatility σ and the growth rate μ for the minutes per busy hour have been estimated by averaging σ_i and μ_i for three different time series i (corresponding to data obtained from [9] for three separate representative switches). It is found that $\mu \approx .30$ per year and $\sigma \approx .65$ per year$^{\frac{1}{2}}$.

We next consider the market price of risk κ. In Appendices B and C, two approaches are presented to estimate the market price of risk. Using these different methods, we obtain values of $\kappa \approx .08$ and $\kappa \approx .03$. The difference between these two values is due to the fact that in the second approach (see Appendix C), we estimate the correlation between the demand and a stock market index (i.e. TSE300), while in the first approach, we estimate κ from the β of the stock of companies whose revenue is primarily from wireless networks. However, since we have only one year of bouncing busy hour traffic data, it is difficult to have an accurate estimate of the correlation in the second approach. Fortunately, our results do not appear to be very sensitive to $\kappa \in [.03, .08]$.

The hardware cost of adding one carrier to a cell site is approximately \$100,000. This does not include engineering/commissioning costs and costs required for hardware upgrades at the switch. \$150,000 per carrier is our total approximate cost with everything included. Table 3 contains a summary of different upgrade costs.

Number of carriers per cell site	Cluster upgrade cost
1	$3,000,000
2	$6,000,000
3	$9,000,000

TABLE 3: *Given a cluster with* 20 *cell sites, this table presents the upgrade costs of adding one or more carriers to each of the cell sites of the cluster. Including all additional costs,* $150,000 *is the approximate cost per carrier.*

Once an order has been placed to upgrade a cell site with an additional carrier, it takes approximately two months before the hardware is delivered, one month to install the hardware, and one month to set up and optimize the new carrier. Hence, it takes about four months from order placement for the equipment to be online.

The monthly maintenance fee of a cluster is the cost of a cell site technician and a T1 cable connection. A cell site technician maintains approximately 20 cell sites. The cost of a cell site technician is assumed to $150,000 per year. Hence the monthly salary cost per cell site is assumed to be $\$150{,}000/20/12 = \$625$ per month or $7,500 per year. Cell site leasing costs are approximately $1,500 per month and T1 backhaul costs are approximately $500 per month per T1. Since, there is one T1 per carrier, a 3 carrier site would have 3 T1s provisioned. Electricity and warranty costs are respectively $250 and $200 per cell site. Thus the total maintenance cost for the cluster cost per month is about $\$150{,}000/12 + \$1{,}500 \times 20 + \$500 \times 20 + (\$250 + \$200) \times 20 = \$61{,}500$. Table 4 contains a summary of the maintenance costs.

Number of carriers per cell site	Cluster monthly maintenance cost
1	$61,500
2	$71,500
3	$81,500

TABLE 4: *Maintenance costs for a cluster in dollars per month. The maintenance costs vary depending on the number of carriers installed on the cluster.*

For voice traffic, the price per carrier is decreasing every year, but not by a significant amount. Vendors offer features to increase the traffic handling capability of each carrier every couple of years. By offering enhancements, the vendor feels justified in keeping the dollars per carrier rate relatively stable. Five percent per year is probably a good assumption for the decay factor of the average revenue per user.

Note that this decay factor is not the same as in Figure 1 (i.e. 8%). Figure 1 was constructed using Bell Canada quarterly financial reports. The decay factor and volatility were estimated using the methods described in Appendix A. From our discussions with network operators [9], we found that a decay factor of 5% was more representative for the expected future decrease of the average revenue per user. Furthermore, the initial installation cost for a cluster is approximatively $20 million and we will assume that the upgrade cost (Table 3) decreases at a rate of 5% per year.

As for cluster characteristics, we assume:

- A cluster is composed of 20 cell sites.
- Each cell site starts with one carrier and can be upgraded to contain up to three carriers.
- Traffic is homogeneous throughout the cluster.

- It takes 4 months after order placement until the upgraded cluster is completely operational.
- Each cell site of the cluster has three sectors.
- The maximum number of minutes per hour that can be handled at 2% blocking for a particular sector of a carrier of a cell site is given in Table 5.
- The maximum number of minutes that can be handled at 2% blocking for the cluster (20 cell sites) in the hour per sector and in the year per sector is given in Table 6. It is computed by multiplying the per hour capacity (see Table 5) times the number of hours per year.

Users	Erlang	Minutes of traffic in the hour
20	13.2	792
40	31.0	1860
60	49.6	2976

TABLE 5: *Maximum number of minutes that can be handled at* 2% *blocking for a particular sector of a carrier of a cell site.*

Minutes of traffic in the hour	
One cell site	Cluster (20 cell sites)
792	15,840
1860	37,200
2976	59,520

TABLE 6: *Maximum number of minutes that can be handled at* 2% *blocking for the cluster (20 cell sites) in the hour per sector.*

The remaining parameters are summarized in Table 7.

Parameter	Value
Investment horizon (T)	5 years
Decay in price (α)	.05/year
Growth rate (μ)	.3/year
Volatility (σ)	.65/year$^{\frac{1}{2}}$
Risk free rate (r)	.04/year
Market price of risk (κ)	[.03 .08]

TABLE 7: *Model parameters that are used to solved equation* (3.14).

5 Results

The volatility σ and growth rate μ have been estimated from reliable bouncing busy hour traffic data time series (see Appendix A). The tests conducted on these time series indicated that there was a large amount of volatility σ in the traffic data despite the general belief in the industry that traffic is very seasonal and predictable. Consequently, in the following we assume that the volatility σ and growth rate μ are given, and we focus on the upgrade time interval and the safety level factor. For a more detailed analysis of the effects of the volatility and growth rate, readers are referred to [3].

Table 8 contains different simulation results for the same safety factor level. The safety level is the percentage of the maximum capacity, at 2% blocking, of the cluster that is allowed to be used. In Table 8, the safety level is set to 100%, meaning that all of the cluster capacity, at 2% blocking, is available and the stream of revenues becomes capped as the maximum capacity is reached. In Table 8, we present our results for both values of the estimated market price of risk κ. We notice that for the higher value of $\kappa = .08$, the percentage in terms of the maximal cluster capacity at which it is optimal to upgrade is higher than when $\kappa = .03$. This result is in accordance with our modeling framework, since as κ increases, the drift term $(r - \kappa\sigma)$ of (3.10) decreases, and the upgrade should occur later. However, in Table 8 we observe that the difference between the two market prices of risk, in terms of the optimal upgrade percentage, is negligible ($\approx$ 4%). This indicates that our results are not very sensitive to our estimate of κ.

In Table 8, we find that at present it would be optimal to add a new carrier to each cell site of the cluster if 80-82% of its maximum capacity is reached when considering monthly upgrade decision dates. Table 8 presents the results for other upgrade decision intervals. As the upgrade decision interval is increased from monthly upgrade decisions to annual upgrade decisions, the upgrade percentage decreases from 80% to 50%. Intuitively, this simply reflects the fact that with less frequent decisions it is better to upgrade earlier, since there are fewer opportunities to make decisions. This behavior is consistent with the results found in [2].

Upgrade decision interval	Add one carrier $\kappa = .03$	Add one carrier $\kappa = .08$	Add two carriers $\kappa = .03$	Add two carriers $\kappa = .08$
monthly	80%	82%	94%	97%
quarterly	69%	71%	78%	83%
semi-annually	62%	64%	71%	77%
annually	49%	51%	60%	62%

TABLE 8: *Today's upgrade decision in terms of upgrade percentage with respect to the maximum capacity of the cluster at* 2% *blocking. We allow* 100% *usage of the total cluster capacity. Above* 100%*, the revenue stream is capped by the cluster maximum capacity. Parameters are* $r = .04$, $\sigma = .65$, $\mu = .3$, *and* $T = 5$ *years. It takes four months between the time the equipment is ordered and the time it is online.*

Finally, we observe from Table 8 that in all simulations we conducted, it may be optimal to add two carriers for each cell site of the cluster instead of just one, if the traffic is high enough.

5.1 Quality of Service Modeling

Our modeling framework enables us to take into account criteria such as quality of service. For instance, it is conceivable that engineers prefer a safety buffer between the maximum capacity (at 2% blocking) and the capacity available to customers. Once this threshold is reached, the quality of service deteriorates. To compensate, customers may receive rebates or free calls. Of course, customers may also seek alternative vendors.

We adopt the following simple model to investigate how quality of service can affect the upgrade decision. Let ϕ be the safety factor, representing a percentage of the maximum capacity of the cluster. Then

Safety factor	Add one carrier
No safety	64%
100%	57%
90%	55%
80%	53%

TABLE 9: *Today's upgrade decision with different safety factor values in terms of upgrade percentage with respect to the maximum capacity of the cluster at* 2% *blocking. Parameters are* $r = .04$, $\sigma = .65$, $\mu = .3$, $\kappa = .08$, $T = 5$ *years. Upgrade time is four months. Cluster upgrade decisions are taken every six months.*

specify revenue as

$$\mathcal{R}(Q, \overline{Q}_j, \tau) = \begin{cases} \mathcal{R}(Q, \overline{Q}_j, \tau)\,(= Q \times \mathcal{P}(\tau)) & \text{if } Q \leq \phi \overline{Q_j} \\ \mathcal{P}(\tau)\overline{Q_j}\max(1 - \frac{Q - \phi\overline{Q_j}}{\phi\overline{Q_j}}, 0.0) & \text{otherwise.} \end{cases}$$

For example, $\phi = .9$ implies that only 90% of the maximum capacity of any given cluster (at 2% blocking) can be used. Note that the above expression implies that revenues eventually drop to zero as demand keeps increasing. Effectively, we are adding a financial penalty as the quality of service deteriorates. It might also be possible to develop penalty functions based on the effect of quality of service on customer "churn rates" (i.e. the loss of customers to other vendors as a result of poor service), or other criteria. This differs from the cases considered above in Table 8 where revenue was simply capped once capacity was reached.

In Table 9, we present the results for several simulations with the safety level ranging from 100% to 80% when the upgrade decision is considered every six months for $\kappa = .08$. We observe that as the safety level decreases, the upgrade occurs sooner in terms of the percentage of the total capacity of the cluster. Similar behavior is reported when upgrade decisions are made for shorter time intervals (e.g. quarterly, or monthly). This phenomenon is intuitively correct: if there is a financial penalty for poor quality of service, upgrades will occur sooner.

6 Conclusion

In this paper, we considered the issue of management of wireless network capacity. While the method presented here is similar to that described in [3], the limitations embedded in the algorithm in that work are alleviated by developing a four dimensional model. This enables us to consider different upgrade decision intervals independent of the time period before the new equipment becomes operational.

As previously noted, in practice current upgrade decisions are often based purely on quality of service criteria. We believe that it is important to also consider financial criteria in terms of maximizing net revenues. By developing appropriate penalty functions which assign a cost to poor quality of service, we can combine both financial and quality of service criteria. This approach will require managers to assign a cost to quality of service issues. Penalty functions could be real financial incentives provided to users (e.g. during high blocking periods, all calls are free), or they could be based on customer churn rates.

Acknowledgment

This work was funded by Bell University Laboratories at the University of Waterloo. We would like to thank Brian O'Shaugnessy and Darlene Farrow for many useful discussions. In particular, we would also like to thank Derek MacAvoy for providing us with invaluable technical information on numerous occasions.

References

[1] C. Chatfield. *The Analysis of Time Series: An Introduction*. Chapman and Hall, London, Fourth edition, 1989.

[2] T. Dangl. Investment and capacity choice under uncertain demand. *European Journal of Operational Research*, 117(3):415–428, September 1999.

[3] Y. d'Halluin, P. A. Forsyth, and K. R. Vetzal. Managing capacity for telecommunications networks under uncertainty. *IEEE/ACM Transactions on Networking*, 10(4):579–588, August 2002.

[4] Y. d'Halluin, P. A. Forsyth, K. R. Vetzal, and G. Labahn. A numerical PDE approach for pricing callable bonds. *Applied Mathematical Finance*, 8:49–77, 2001.

[5] J. Hull. *Options, Futures, and Other Derivatives*. Prentice Hall, Inc., Upper Saddle River, NJ, fifth edition, 2002.

[6] D. Kröner. *Numerical Schemes for Conservation Laws*. Chichester, New York, 1997.

[7] J. S. Lee and L. E. Miller. *CDMA Systems Engineering Handbook*. Artech House, 1998.

[8] WestBay Engineers Limited. What is an Erlang? `http://www.erlang.com/whatis.html`.

[9] D. MacAvoy. Private communication, 2002. Bell Mobility, Associate Director Access Technology Planning.

[10] S. Romero. Success of cellphone industry hurts service. *The New York Times*, November 18, 2002.

[11] E. S. Schwartz and M. Moon. Rational pricing of internet companies revisited. Working Paper, The Anderson School, UCLA, 2001.

[12] P. Wilmott. *Derivatives*. John Wiley and Sons Ltd, Chichester, 1998.

[13] R. Zvan, P. A. Forsyth, and K. R. Vetzal. Robust numerical methods for PDE models of Asian options. *Journal of Computational Finance*, 1:39–78, Winter 1998.

[14] R. Zvan, P. A. Forsyth, and K. R. Vetzal. A finite volume approach for contingent claims valuation. *IMA Journal of Numerical Analysis*, 21:703–731, 2001.

Appendices

A Estimation of Growth Rate and Volatility Parameters

We obtain daily bouncing busy hour traffic data from [9] for three representative switches for a time period of a little over one year (from mid-2000 to mid-2001). We will refer to these different switches as A, B, and C. An initial analysis of the traffic data showed a strong autocorrelation of the time series within each week (see Figure 7(a)). This is not surprising, since we expect that there will be repetitive patterns within each week. To filter out this effect, we average the yearly network traffic for each day of the week separately, and choose the day with the highest average bouncing busy hour traffic (see Figure 7(b)). We then use this same day each week to estimate week to week effects.

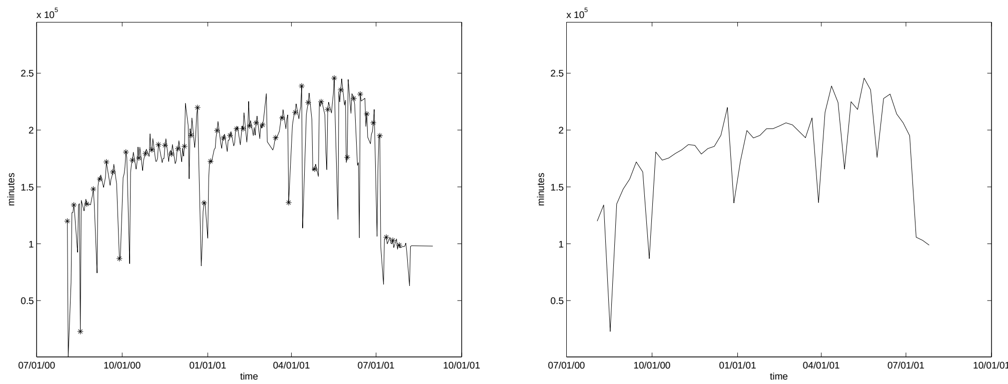

(a) *One year total daily bouncing busy hour traffic, switch A. The stars indicate the weekly highest traffic days.*

(b) *One year total daily bouncing busy hour traffic, switch A. The data corresponds to the highest average day traffic per week. This day is chosen based on the average network traffic that day encounters throughout the year. Once that day is found, it remains the same for the rest of the year.*

FIGURE 7: *Total daily bouncing busy hour traffic time series for switch A.*

Prior to estimating the drift term μ and the volatility σ, an inspection of Figure 7(b) reveals several suspicious large drops in traffic. Basing our volatility estimate on this data would produce a very high value. Some of the short term declines are simply due to statutory holidays or scheduled maintenance. As these are known events for low network traffic, we should not take them into account when estimating volatility. Consequently, the dates corresponding to statutory holidays and suspicious changes are smoothed out using interpolation. The month of December is also ignored since it is a known low traffic period. Figure 8 presents the time series once the holidays and the large drops have been removed and smoothed out.

Equation (3.1) implicitly assumes that bouncing busy hour traffic is lognormally distributed. Consequently, based on the bouncing busy hour traffic logarithmic weekly changes (i.e. $u_i = \log(\frac{Q_{i+1}}{Q_i})$ for $i = 1, \ldots, M-1$, where M is the number of points in the time series once the weekly cycles have been removed) the drift term is estimated using a least squares method. Next the estimated trend is removed. The final results are presented in Figure 9.

As a diagnostic test, we use the Ljung-Box Q-statistic (see, e.g. [1]) on the detrended logarithmic relative traffic change. We calculate this for 4 lags. For each switch, this test statistic indicates that there is no remaining serial correlation in the data.

The volatility is finally estimated under the assumption that relative changes in network traffic are lognormally distributed, as implied by equation (3.1). Table 10 presents our results for different time series.

From Table 10, we notice that network traffic is highly volatile. The volatility values can be compared to volatilities in the range of 15% to 30% for stock market indices. For our simluations we will average the

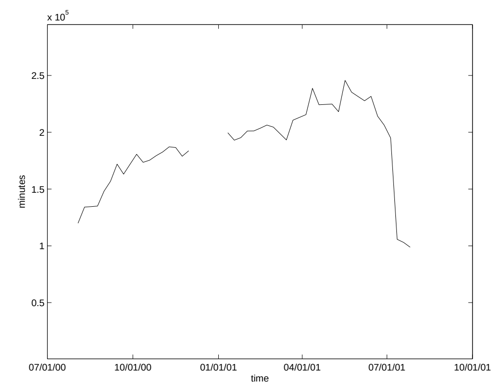

FIGURE 8: *One year total daily bouncing busy hour traffic on switch A without major holidays and suspicious events. The month of December is ignored. The data corresponds to the highest average day traffic per week. This day is chosen based on the average network traffic that day encounters throughout the year. Once that day is found, it remains the same for the rest of the year.*

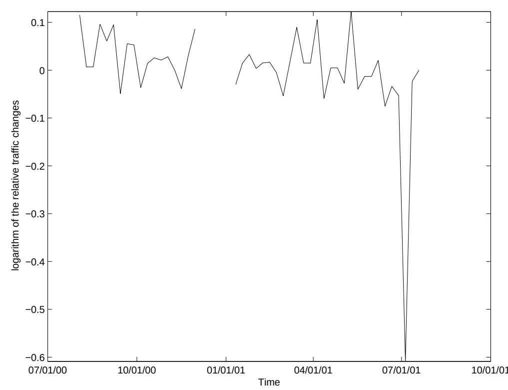

FIGURE 9: *One year logarithmic daily bouncing busy hour relative traffic changes. Known holidays and suspicious changes have been replaced using interpolation. The month of December is ignored and the trend of the time series has been removed.*

results for both the drift term and volatility, using $\mu \approx .30/\text{year}$ and $\sigma \approx .65/(\text{year})^{\frac{1}{2}}$.

Bouncing busy hour daily traffi c time series		
	Drift term $\mu\,/(year)$	Volatility (fi ltered data) $\sigma\,/(year)^{\frac{1}{2}}$
Switch A	-.24	.90
Switch B	.41	.74
Switch C	.73	.32

TABLE 10: *Summary table of results for the time series for three representative switches provided by [9].*

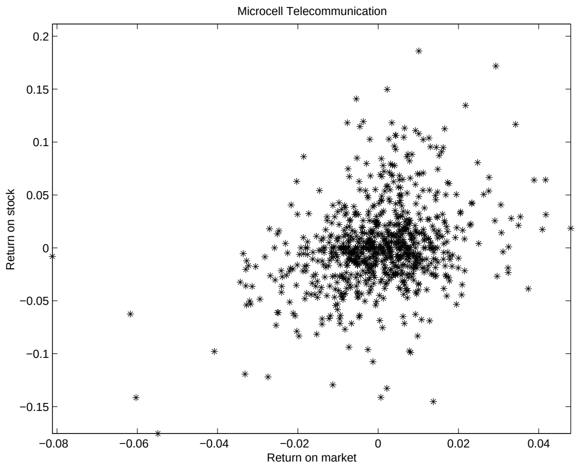

FIGURE 10: *Return on TSE300 index versus return on Microcell Communications. Each point represents pairs of daily returns. The vertical axis measures the daily return on the stock and the horizontal axis that of the TSE300.*

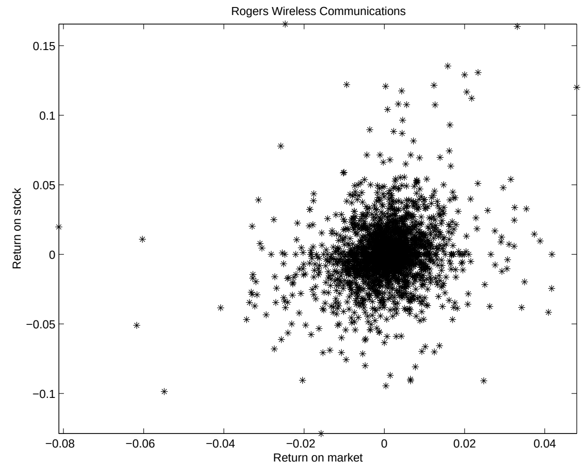

FIGURE 11: *Return on TSE300 index versus return on Rogers Wireless Communications. Each point represents pairs of daily returns. The vertical axis measures the daily return on the stock and the horizontal axis that of the TSE300.*

B Estimation of the Market Price of Risk

We follow the general approach described in [11] which uses stock market data. However, Bell Mobility is not a publicly traded company, so we need to find public Canadian companies whose revenue streams are similar to that of Bell Mobility. Two such companies are Rogers Wireless Communications Inc. and Microcell Telecommunications Inc.

For each company, we must first compute the systematic risk exposure using the standard capital asset pricing model. Because each firm has significant amounts of debt outstanding, we will initially use the levered equity's beta β as the measure of the systematic risk. To estimate βs, we run linear regressions of returns for each stock versus the return on the market. We use the TSE300 index as a proxy for the return on the market. Figures 10 and 11 present the return for the market versus the return for Microcell and Rogers.

Figures 12 and 13 show the line of best fit superimposed on each point representing pairs of daily return data. For Microcell Telecommunications we find $\beta = 1.0455$, while $\beta = .5603$ for Rogers Wireless.

Once the levered betas for both firms have been estimated, we calculate the beta for the hypothetical unlevered firm. To compute the unlevered firm's beta from the levered equity beta, the firm's debt, market capitalization and corporate tax rate must be estimated. The unlevered firm's beta is then given by

$$\beta^{\text{unlevered}} = \frac{E}{E+(1-T_c)D}\beta, \tag{B.1}$$

where

$$\begin{aligned} E &= \text{market capitalization (total number of shares times share price)},\\ D &= \text{long term debt},\\ T_c &= \text{corporate tax rate}. \end{aligned}$$

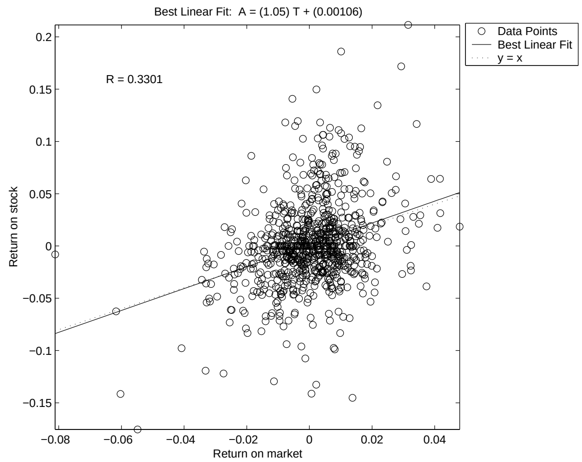

FIGURE 12: *The figure shows the line of best fit superimposed on the points in Figure 10. The slope of the regression line* $\beta = 1.0455$*. For this regression,* $R^2 = .3301$.

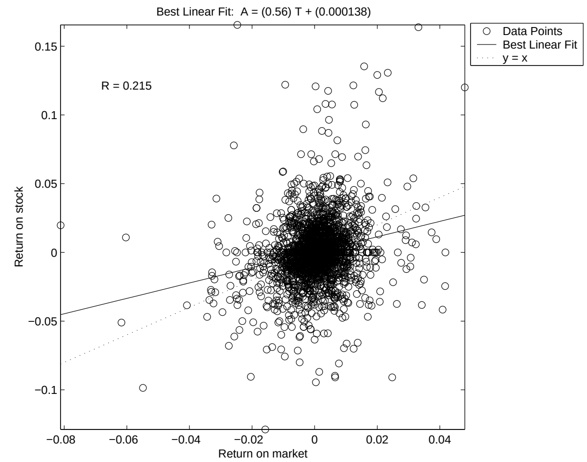

FIGURE 13: *The figure shows the line of best fit superimposed on the points in Figure 11. The slope of the regression line* $\beta = .5603$*. For this regression,* $R^2 = .2150$.

Long term debt	$ 1,887,048,000
Corporate tax rate	40%
Stock price (on February 13, 2002)	$2.68
Number of shares	240,000,000

TABLE 11: *Microcell Telecommunications Inc. corporate data. Dollar figures are in Canadian funds. This data is based on a financial analysis published by the web site* h*ttp://www.globeinvestor.com/ on February 13, 2002 and Microcell Telecommunications Inc.'s quarterly financial reports.*

Note that for both companies, we will use the statutory rate of 40% for the corporate tax rate. Tables 11 and 12 contain the information we used to calculate the unlevered beta for each company.

Using equation (B.1), we find that

$$\beta^{\text{unlevered}}_{\text{Microcell}} = \frac{2.68 \times 240}{2.68 \times 240 + (1 - .4) \times 1887.048} 1.0455$$
$$= .37876,$$

and

$$\beta^{\text{unlevered}}_{\text{Rogers}} = \frac{17.78 \times 144.4}{17.78 \times 144.4 + (1 - .4) \times 2305.638} .5603$$
$$= .36411$$

Averaging these two values, we estimate that the unlevered Bell Mobility β is given by $\beta^{\text{unlevered}}_{\text{Bell Mobility}} = .3714$. Note that our use of unlevered betas means that our real option valuation is biased low for an investment project financed with debt, as interest tax shields have not been accounted for. The value of these tax shields

Long term debt	$ 2,305,638,000
Corporate tax rate	40%
Stock price (on April 18, 2002)	$17.78
Number of shares	144,400,000

TABLE 12: *Rogers Wireless Telecommunications Inc. corporate data. Dollar figures are in Canadian funds. This data is based on a financial analysis published by the web site* h*ttp://www.globeinvestor.com/ on April 18, 2002 and Rogers Wireless Telecommunications Inc.'s quarterly financial reports.*

Long term debt	$ 14,861,000,000
Corporate tax rate	40%
Stock price (on August 2, 2002)	$39.08
Number of shares	808,600,000

TABLE 13: *Bell Canada Enterprises corporate data. Dollar figures are in Canadian funds. This data is based on a financial analysis published by the web site* h*ttp://www.globeinvestor.com/ on August 2, 2002 and Bell Canada quarterly financial reports.*

could be added later, if desired. Given the current financial situation in the telecommunications sector, it appears unlikely that new debt financing would be available at present.

Now that we have estimated the unlevered beta for Bell Mobility, we can compute the market price of risk κ. As mentioned, we follow the methodology described in [11]. For readers unfamiliar with this approach, we present here a short description.

In [11], it is claimed that using the β of the firm's stock to estimate the risk premium in a real options model constitutes a significant improvement over earlier work which used the traditional approach based on the covariance of changes in the state variable(s) with the market portfolio. Briefly, and more simplistically than in [11], the idea is as follows.

Suppose there is a single stochastic factor X (in our context this is Q), which follows the risk-adjusted process:

$$dX = (\mu - \lambda)X\,dt + \sigma X\,dz, \tag{B.2}$$

where μ is the real world drift, σ is the volatility, λ is the risk premium (market price of risk multiplied by volatility $\lambda = \kappa\sigma$ is our new notation), and dZ is an increment of a Wiener process. Let the firm's stock price be S. From Ito's lemma [11, 12], we have:

$$\frac{dS}{S} = \frac{\left[\frac{1}{2}\sigma^2 X^2 S_{XX} + (\mu - \lambda)XS_X + S_t\right]}{S}dt + \frac{\sigma X S_X}{S}dz, \tag{B.3}$$

where the risk premium is:

$$\frac{\lambda X S_X}{S}. \tag{B.4}$$

The authors of [11] then use the intertemporal capital asset pricing model (ICAPM) in the following way. The firm's stock β, denoted by β_S, is the covariance between returns on the market portfolio M and returns on the stock. This can be written as a function of the "β" of the stochastic factor X:

$$\beta_S = \frac{\sigma_{SM}}{\sigma_M^2} = \frac{XS_X}{S}\frac{\sigma_{XM}}{\sigma_M^2} = \frac{XS_X}{S}\beta_X, \tag{B.5}$$

where σ_{SM} is the covariance between changes in S and M and similarly for σ_{XM}. In the ICAPM, the expected return on the stock is:

$$r_S = r_f + \beta_S(r_M - r_f) = r_f + \frac{XS_X}{S}\beta_X(r_M - r_f) \quad \text{(B.6)}$$

where r_f denotes the risk free rate of interest and r_M is the expected return on the market portfolio. Equating the risk premium from (B.4) with that implied in (B.6) gives:

$$\begin{aligned} \frac{\lambda XS_X}{S} &= \frac{XS_X}{S}\beta_X(r_M - r_f) \\ \Rightarrow \lambda &= \beta_X(r_M - r_f). \end{aligned} \quad \text{(B.7)}$$

Using (B.5), we have:

$$\lambda = \frac{S\beta_S}{XS_X}(r_M - r_f), \quad \text{(B.8)}$$

i.e. the risk premium is a function of the expected excess market return, the firm's current stock price, the β of the firm's stock price, the current level of the stochastic factor X, and S_X.

Returning to our context, we then have:

$$\lambda = \kappa\sigma = \frac{S\beta^{unlevered}_{Bell\,Mobility}}{QS_Q}(r_M - r_f), \quad \text{(B.9)}$$

where S is BCE's current stock price, Q is the current level of traffic, and S_Q is the first derivative of the stock price with respect to the level of traffic.

All the parameters from equation (B.9) are known except S_Q, r_M and r_f. For the risk free rate r_f, we assume a value of $r_f = .04$. We assume that the expected market return r_M is 6% higher than the risk free rate (roughly consistent with the average level for the past 50 years of Canadian data). Thus we have $r_M = .1$. S_Q is a more challenging parameter since there is no direct data from which we can determine S_Q. Consequently, to estimate S_Q, we estimate a linear regression based on BCE stock price and traffic data. Figure 14 presents the results. We find that $S_Q = 9.0057 \times 10^{-05}$ dollars per (minutes/busy hour).

Consequently, replacing S_Q by its corresponding value, using $Q = 1.9197 \times 10^5$ minutes per hour(Q corresponds to the busy bouncing hour traffic data as of August 2, 2002) and the stock price information contained in Table 13 and equation (B.9) gives

$$\begin{aligned} \lambda &= \frac{39.08 \times .3714}{1.9197\,10^5 \times 9.0057\,10^{-5}}(.1 - .04) \\ &= \frac{14.5143}{17.2882} \times .06 \\ &= 0.0503 \end{aligned}$$

Hence, the market price of risk $\kappa = \frac{\lambda}{\sigma} = .077$.

C Estimation of the Market Price of Risk, An Alternative Approach

In this section we present another approach to estimate the market price of risk κ. This follows the discussion in [5]. The assumption is that the investment to upgrade capacity depends solely on the network usage. Consequently, if we can determine the correlation between market returns (i.e. TSE index) and the bouncing busy hour changes, it is possible to estimate the market price of telecom risk as follows

$$\kappa = \rho\kappa_M, \quad \text{(C.1)}$$

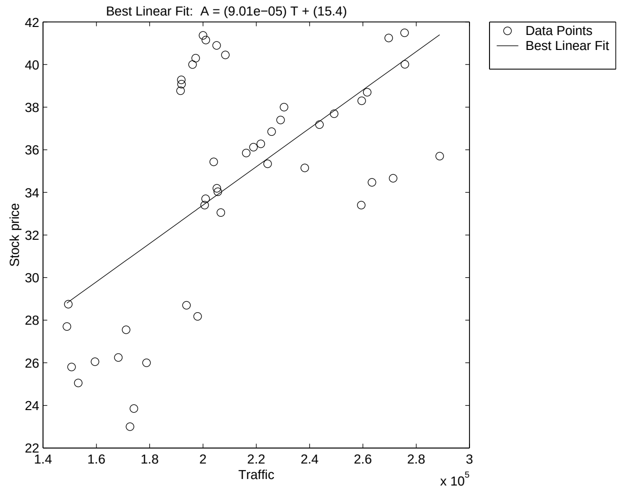

FIGURE 14: *The figure shows the line of best fit superimposed on these points. The slope regression coefficient* $S_Q = 9.0057\,10^{-05}$. *For this regression,* $R^2 = .6198$. *We use the one year time series of bouncing busy hour traffic data obtained from [9]*

where ρ is the instantaneous correlation between the bouncing busy hour changes and returns on a broad index of stock market prices (in our application, the TSE300), and

$$\kappa_M = \frac{r_M - r_f}{\sigma_M},$$

where r_M is the anticipated market return, r_f is the risk free rate, and σ_M is the market volatility. We assume that $r_M - r_f = .06$, and that $\sigma_M = .2383$. The volatility σ_M is determined using the same historical period as the network traffi c time series.

Using the same time series as in Appendix B [9] we fi nd that $\rho \approx 0.1171$. Thus

$$\kappa = 0.1171 \times \frac{.06}{0.2383} \approx .03.$$

Consequently, we fi nd that κ varies between ≈ .08 (Appendix B) and ≈ .03 (Appendix C). The difference between these two values is due to the fact that in the second approach (Appendix C), we estimate the correlation between the demand and a stock market index (i.e. the TSE300). However, since we have only one year of bouncing busy hour traffi c data, it is diffi cult to have an accurate estimate of the correlation.