Customer-centric Network Upgrade Strategy: Maximizing Investment Benefits for Enhanced Service Quality

J. Xiao, R. Boutaba
School of Computer Science, University of Waterloo
200 University Ave. W.
Waterloo, ON, Canada
{j2xiao|rboutaba}@bbcr.uwaterloo.ca

Abstract

With the ever increasing demand for network resources, network operators and Internet Service Providers are under constant pressure to accommodate more network bandwidth and offer better service quality via periodic network upgrades. Given a budget constraint, a sound network upgrade decision should maximize investment benefit which is contingent on the degree of customer satisfaction. This paper presents a customer-centric approach in making network upgrade decisions, where customer satisfaction is the key evaluation criterion. Network performance is related to customer's perceived service quality and component upgrades are assessed based on their profitability. As demonstrated using a case scenario, our approach results in effective upgrade decisions that enhance service quality, improve customer satisfaction, and maximize revenue.

Keywords

service management, network upgrade, customer satisfaction, investment decision

1. Introduction

As networked applications become increasingly prevalent in daily business operations and home activities, the demand for network services is growing in leaps and bounds. While this trend affords great marketing potential for Internet Service Providers (ISPs) and network operators, it also exerts immense pressure on the suppliers to accommodate more network bandwidth and offer better service quality via periodic network upgrades. A sound network upgrade decision should maximize revenue given a limited budget.

In practice, the network upgrade process is a long-term planning strategy consisting of three crucial steps: identification, assessment, and decision. Network components impacting service quality should be identified, the effect of their upgrades should be quantified as monetary benefits, and an upgrade decision that maximizes revenue should be determined. The entire process requires in-depth analysis of the underlying network infrastructure, the network performance, the characteristics of its supported service operations, and the customer behaviors. In this paper, we focus on regional networks, where the network size and service subscriber population are relatively small. Nevertheless, our approach is generally applicable to large scale networks. Current practice often produces ineffective investment decisions that do not achieve the desired level of service quality. We believe the inadequacy lies fundamentally at its network-centric view. When a component reaches some network performance threshold, it is tagged for potential upgrade, and the network infrastructure is scrutinized in terms of its QoS performance (e.g. delay, throughput, availability, etc.). While such an approach reflects the network status, it does not link the performance of a component to its severity or scope of impact on the customer population. Our work departs from this network-centric paradigm by establishing customer satisfaction as the key evaluation criterion. Our approach is motivated by two observations: 1) an ISP's revenue is solely based on its customers' willingness to use its services; 2) an ISP maintains competitiveness in the market by meeting the customer expectations.

In our customer-centric approach, we first establish a linkage between customer satisfaction, service quality and network performance using an analytical framework: the metrics tree model. This affords us the possibility of identifying candidate replacement components that have the greatest impact on the customer population. We then evaluate the benefits of component upgrades as changes to customer satisfaction, and consequently changes in revenue. Finally, we formalize the upgrade decisions as a profitability-based optimization problem.

The remainder of this paper is organized as follows. Section 2 presents background and related works. Section 3 gives an overview of our approach. In Section 4, we establish the relation between network performance and customer satisfaction for the identification of candidate upgrade components. Section 5 assesses the monetary benefit of component upgrade. Section 6 presents the optimization model for making investment decision. The effectiveness of our approach is demonstrated in Section 7. Section 8 summarizes our approach and presents some future works.

2. Background and Related Works

Investment decision making has been an established practice in business planning. The soundness of a decision is largely dependent on the ability to correctly analyze the status of the operations and the trend of market growth. In the context of the network service market, the basis for such in-depth analysis relies heavily on the information gathering capability of key service aspects: network performance, service operations, and customers. Existing works in Internet measurement offer rich reservoir of network statistics, ranging from statistical information collection from Management Information Base [3] (via SNMP [2]), monitored network QoS performance (e.g. RMON [4]), to actively measured end-to-end path information (e.g. via pinging). Today's routers are even capable of keeping track of individual traffic flow information (e.g. Cisco's NetFlow [1]). At the service and customer level, customer access can be tracked by customer-side monitors, and service difficulties are recorded in the form of trouble-ticketing logs. Due to the lack of mapping between network performance and its impact on customers, trouble spots are currently identified on the basis of simple network QoS metrics (e.g. component utilization). Such identification does not reflect the customer's perception of a service, and adversely affects the outcome of the investment decisions. In our customer-centric approach, we leverage the existing information gathering capabilities to establish a sensible mapping between the network QoS and customer satisfaction. This serves as our basis for assessing component upgrade benefits.

Some previous works, such as [10] and [12], emphasize the importance of analyzing both the customer and the network profiles in a business decision process. However, [10] does not offer any means of correlating the two. [12] highlights the importance of customer utility and devises a utility model based on customer's service preferences. The work assumes the existence of a mapping between the network performance and a customer's perceived service quality. Motivated by these previous attempts, our work provides a well-structured linkage between the network, the service, and the customer. The resulting mapping is specific to the underlying network infrastructure, the requirements of service operations, and the particular characteristics of each customer.

Using real option pricing, [14] tries to determine the best investment time for link capacity upgrades. Their work evaluates the profitability of an investment in terms of revenue generated from the network usage. To account for customer dissatisfaction due to congestion, a simple discount factor is associated with each congested link. In our work, revenues are estimated based on the customer's perception of service quality rather than the traffic volume. This approach elevates our analysis from the underlying networks to the services and customers, and leads to investment assessments emphasizing better service quality.

[11] proposes a revenue-based approach to component upgrade optimization given a fixed budget constraint. The profitability of each network component is estimated based on the amount of customer traffic it bears, with the assumption that previously unsatisfied customers are satisfied after the upgrade. Similarly, we establish the profitability of a network upgrade by analyzing customer traffic flows. In contrast, our approach attempts to estimate the benefit of an upgrade to the customers, taking into account key factors influencing a customer's perception of service quality: interdependencies among network components, the increase in traffic demand, and the access behavior of customers.

To a customer, his/her perception of service quality is only related to the network components bearing the customer traffic. Therefore, it is important to focus on part of the network that serves the customer. Considering the amount of customer flows generated over time (e.g. months), the issue of scalability should be considered. [6] demonstrates the effectiveness of tracking only frequently used flows to achieve better efficiency. In the same spirit, we utilize pruning to achieve better scalability in our customer flow analysis.

3. Overview of Approach

An overview of our component upgrade approach is depicted in Figure 1. The metrics tree model is central to the identification and assessment stages of our analysis, as it provides the essential relationship between network QoS performance and customer satisfaction. In the identification stage, component dissatisfactions are computed as the service dissatisfaction of customers caused by "troubled" network components. We describe a "troubled" component as a component impacting network performance (e.g. high delay). The output of the identification process is a set of candidate upgrade components selected according to their component dissatisfaction ratings.

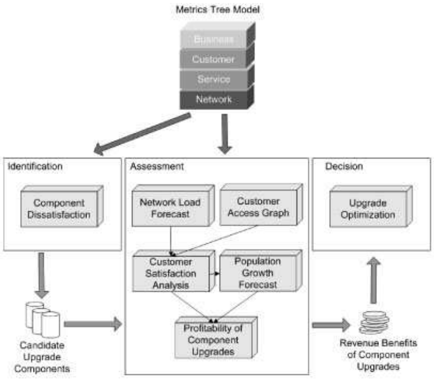

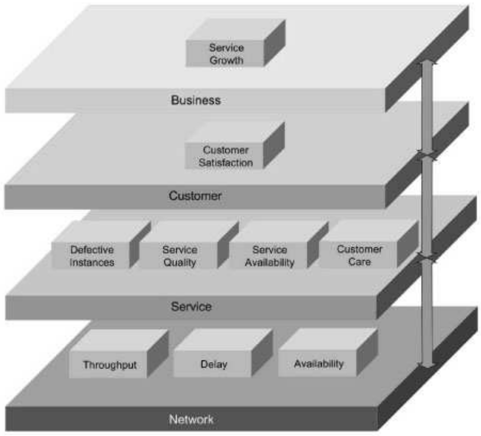

Figure 1: Process overview for component upgrade decision

Figure 2: Overview of the metrics tree model

During the assessment stage, network load forecasts for some future time T are obtained. Combined with the set of customer access graphs, which abstract the customers' network access behaviors as frequently used paths, the change in customer satisfactions due to upgrade is computed via applying the metrics tree model. The changes in customer satisfaction is then mapped to changes in business revenue, and the population growth at time T is estimated based on an economic model. The monetary profitability of any component upgrade is then evaluated as the combination of changes in revenue due to customer satisfaction changes, and the increase in revenue due to subscriber population growth. The output of the assessment stage is a set of revenue benefits for each candidate upgrade component.

The decision process involves the optimization of revenue benefits given a set of candidate upgrade components and a fixed budget constraint. By applying existing optimization techniques, we obtain the set of component upgrades that maximizes investment benefits. The process can be repeated multiple times to improve the quality of a upgrade decision.

4. Relating Network Performance to Customer Satisfaction

In this section, we present the linkage between QoS performance at the network level and customer satisfaction at the customer level. Using this relation, we then identify network components that cause the greatest customer dissatisfactions.

We hereby distinguish network quality of service (QoS), service quality, and customer satisfaction as follows:

- *Network quality of service*: a set of network QoS metrics that measures the performance status of a network. QoS metrics include throughput, delay, jitter, utilization, etc.
- *Service quality*: a set of service level metrics that measures all quality aspects of a specific service, such as network performance, service availability, customer care, etc.

- *Customer satisfaction*: a measure of customer's satisfaction towards his/her subscribed service.

4.1 The Metrics Tree Model

The metrics tree model is an analytical framework relating network level performance metrics to service, customer, and business level metrics. The model is constructed as a strictly layered tree, where metrics at upper layers of the tree relate to other metrics in the same or lower layer via mapping functions. Figure 2 illustrates this layering and some metrics appearing in this paper.

At the network level, QoS metrics are obtained for network flow paths. This information is used to evaluate defective service instances at the service level. Customer satisfaction is then computed at the customer level by analyzing each customer's defective service instances with regard to customer's service preferences, application characteristics, and access behaviors. The business layer abstracts the service layer and customer layer information into aggregate metrics that are meaningful to the business operations (e.g. growth rate).

4.2 Relating Network QoS to Customer Satisfaction

We now focus our attention on how network QoS is related to service quality in the metrics tree model. The demonstration below is concerned with ADSL service quality, nevertheless our approach is generalized to all common ISP service offerings. Consider an ADSL service subscriber, whose frequent activities (e.g. connecting to work from home) go through a set of paths.

$$cust = \{P_1, P_2, \ldots, P_n\}, P = \{x_1, x_2, \ldots, x_m\}, \tag{1}$$

where x are network components along a path P.

Since ISPs rarely have control over networks outside their administrative domain, for end-to-end paths involving components in foreign administrative domains, P is modified to only span across components residing inside the ISP's domain.

We quantify the impact of low QoS performance at the service level as defective instances. A defective instance is a series of consecutive network QoS measurements whose values are below a specified performance tolerance limit. The limit is deduced from service level agreement and based on the particular customer flow. During the course of a network trouble, the QoS measurements may fluctuate wildly above and below the performance tolerances. We introduce an network flux parameter to account for this fluctuation. We say that after the first recorded QoS measurement that is below tolerance limit, the onset of such defective instance ends only if up to network flux number of consecutive QoS measurements are recorded as above the performance tolerance limit. The ADSL service performance is concerned with defective throughput instances and defective delay instances. The time length of such an instance on a path P is denoted as $Len(TTI_P)$ and $Len(DTI_P)$ for throughput and delay respectively.

From the set of recorded throughput and delay defective instances over some evaluation time interval I_{ADSL}, we can define throughput satisfaction $TSat_P$ and delay satisfaction $DSat_P$, using the concept of cumulative defective tolerance levels. Cumulative defective tolerance levels ($AggTDL_{ADSL}$ for throughput and $AggDDL_{ADSL}$ for delay) are the total amount of defective service time a service can tolerate. Their values are derived from service level specifications, and are specialized to each customer as $AggDD_P$ (delay tolerance of path P) and $AggTD_P$ (throughput tolerance of path P).

$$AggTD_P = \frac{AccessT_P \times AggTDL_{ADSL}}{I_{ADSL}}, AggDD_P = \frac{AccessT_P \times AggDDL_{ADSL}}{I_{ADSL}} \tag{2}$$

Then throughput and delay satisfaction of path P ($TSat_P$ and $DSat_P$) are computed as the extent to which the total amount of defective service times ($\sum Len(TTI_P)$ and $\sum Len(DTI_P)$) exceed cumulative defective tolerance levels ($AggTD_P$ and $AggDD_P$). For simplicity, the satisfaction deduction for excess defective service times is a linear relationship.

$$TSat_P = \begin{cases} 1 & \text{if} \sum Len(TTI_P) \leq AggTD_P \\ h(\frac{\sum Len(TTI_P) - AggTD_P}{AggTD_P}) & \text{otherwise} \end{cases} \tag{3}$$

$$DSat_P = \begin{cases} 1 & \text{if} \sum Len(DTI_P) \leq AggDD_P \\ h(\frac{\sum Len(DTI_P) - AggDD_P}{AggDD_P}) & \text{otherwise} \end{cases} \tag{4}$$

$$\textit{where } h(x) = \begin{cases} 0 & \text{if } x > 1 \\ 1 - x & \text{otherwise} \end{cases}$$

The customer's perceived ADSL service quality satisfaction $QSat_{ADSL\_P}$ of path P is then a weighted average of throughput and delay satisfactions. The weighing factor α_1 and α_2 are preference parameters concerning the type of applications used over path P.

$$QSat_{ADSL_P} = \alpha_1 \times TSat_p + \alpha_2 \times DSat_p, \textit{where } \alpha_1 + \alpha_2 = 1 \tag{5}$$

Up to now, we have considered the service quality satisfaction for a particular path. The customer's overall satisfaction of ADSL service quality can be obtained as the aggregation of service qualities along the ADSL service paths a customer uses, weighed by their respective access frequency.

$$QSat_{ADSL} = \frac{\sum_{P \in cust}(QSat_{ADSL_P} \times AccessT_P)}{\sum_{P \in cust} AccessT_P} \tag{6}$$

We compute customer satisfaction by considering service quality $QSat_{ADSL}$, service availability $ASat_{ADSL}$, and customer care $CCSat_{ADSL}$.

$$CSat_{ADSL} = \beta_1 \times QSat_{ADSL} + \beta_2 \times ASat_{ADSL} + \beta_3 \times CCSat_{ADSL}, \quad \textit{where } \beta_1 + \beta_2 + \beta_3 = 1 \tag{7}$$

β_1, β_2, and β_3 are the customer's service preference parameters. Due to length limitation, we will not detail service availability and customer care, and simply assume these results are given by network operators.

4.3 Identify Candidate replacement components based on customer dissatisfaction

Instead of using component utilization as the selection criterion for identifying candidate replacement components, we use the notion of component dissatisfaction. Given a number of components identified as the "trouble spots" along a path experiencing service difficulties, we can assert the degree of dissatisfaction such a component x_d causes to the customer population. Let Q be a subset of $cust$ where x_d is a member of all paths in Q, $Sat\_X_d\_cust$ be the satisfaction rating of x_d for a customer, then the component dissatisfaction rating $Dissat\_X_d$ can be expressed as:

$$Sat\_X_d\_cust = \frac{\sum_Q (CSat_{ADSL_Q} \times AccessT_Q)}{\sum_{P \in cust} AccessT_P} \tag{8}$$

$$Dissat\_X_d = \sum_{customers} \begin{cases} 0, & \textit{if } x_d \notin \Phi_{cust} \\ 1 - Sat\_X_d\_cust & \textit{otherwise} \end{cases} \tag{9}$$

where Φ_{cust} is the set of trouble components for the customer

A higher $Dissat\_X$ value not only indicates a severer impact to customer satisfaction, but also a broader scope of impact to customer population.

5. Assessment of component upgrade benefits

To make informed upgrade decisions, it is essential to estimate the benefit of each component upgrade. Due to the inherent intricacies among network components and the potential growth in network bandwidth usage, the assessment process should consider the effect of component upgrades not only on the upgraded components themselves, but also on the other impacted components in the network. We evaluate the benefits of upgrade at a future time T, the benefit horizon, which is the time in near future when network upgrades are to be re-evaluated. The introduction of benefit horizon serves two purpose: 1) since the process of upgrade only occurs periodically, to avoid short-sightedness, the upgrade benefits should be significant in the present and future; 2) the benefit of upgrades should be evaluated against the consequence of no upgrades at future time.

5.1 Overview of assessment process

As presented in the previous section, we can determine the customer satisfaction based on a number of key network QoS metrics: throughput, delay, and availability. Using load forecast technique, we estimate the network load conditions at future time T, before and after component upgrades. Then, we compute the change in customer satisfaction by analyzing customer traffic flows on the load graphs. Due to network interdependencies, it is essential to analyze all customer traffic flows. This requirement raises a scalability issue: per flow analysis is expensive. In response, we introduce the concept of customer access graph, which only considers frequently used flow paths. Based on the changes in customer satisfaction, we quantify the profitability of enhancing the performance of particular customer paths, and then compute the profitability of component upgrades. Customer population growth is also considered using a modified Bass model.

5.2 Network Load Analysis

To estimate the network load condition at some future time, we abstract our view of the network load status as a directed graph, where the nodes and edges of the graph are weighed by the load condition at time T+t. Time T is the benefit horizon and time t is a small time interval (e.g. time of day in a week). Studies on typical ISP operation data suggest that aggregate network load exhibits both daily and weekly traffic patterns. We establish a series of network load graphs over a week to capture the load characteristic of the network around future time T. A number of works done in the area of network demand

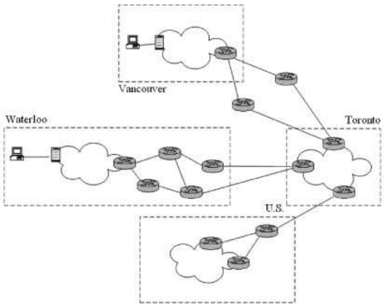

Figure 3: Example of a complete customer access graph

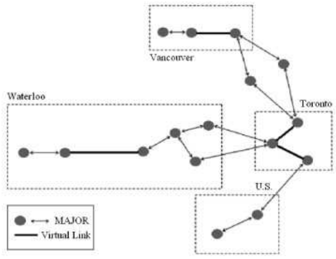

Figure 4: Illustration of pruning process

forecast, such as [13] and [9], can serve as guidelines for establishing the load condition on each network component at time T+t.

Based on the network load graphs, we can estimate the changes in network load at T+t due to component upgrades. The process involves rippling load changes from a upgraded component to its downstream neighbors in an iterative fashion.

5.3 Generation of Customer Access Graph

The customer access graph consists of nodes and links weighed by its total access time of a customer. We prune the graph by discarding nodes/links infrequently used by the customer. The following procedure outlines the pruning process (see Figure 3, 4). The resulting graph is the customer access graph.

1. Mark the nodes and links whose weight is no less than Minimum Access Frequency * Total Access Time as MAJOR. The total access time is the total amount of time a customer uses the network. The minimum access frequency is a predetermined value between 0 and 1, which moderates the degree of pruning. We observe that the marked nodes and links do not necessarily form a connected graph. In such cases, each disjoint sub-graph is denoted as a zone. For each zone, we denote the set of nodes that takes incoming traffic from other zones as entry-nodes of the zone, and the set of nodes that sends outgoing traffic to other zones as exit-nodes. A node can be classified as both an entry and exit node.
2. For each zone, examine its exit-nodes interconnecting entry-nodes of other zones. For every pair of exit-node to entry-node interconnection, establish a “virtual link” between them.

The nodes marked MAJOR are the focus of our analysis as they are the frequently accessed portion of the network. However, as we are conducting analysis on end-to-end paths, the access graph must be a completely connected graph spanning these zones. Virtual links are used for this purpose. It is interesting to note that the components marked MAJOR are parts of a network where the customer must access (e.g. access network) or where most customer flows would congregate (e.g. gateways to national transit network). The virtual links are abstraction of the “cloud-like” portions of the network where the routes could change frequently (e.g. parts of a regional network interconnecting local access networks to IP backbone gateways). By pruning away the "cloud-like" sub-paths, we are able to merge together large number of similar flow paths that only differ in their sub-paths traversing the clouds. The QoS performance of a virtual link is estimated as the average QoS performance of paths in the "cloud". Although simplistic, we only aim at having an estimate of the "cloud" performance and in doing so, significantly reduces the number of flow paths to be analyzed.

5.4 Quantify Component Upgrade Benefit

We first evaluate the changes in customer satisfaction due to component upgrades. This process involves the quantification of customer satisfaction using network load forecast and customer access graphs. Suppose QoS performance can be estimated based on our network load graph for any end-to-end path on the graph, then the customer satisfaction of a path P' ($CSat_{P'}$) can be computed using relations described in section 4.

The following process details the evaluation of changes in customer satisfaction:

1. Construct network load graphs by performing load readjustment.
2. For each customer, for each path P' on the access graph, compute $CSat_{P'\_old}$ based on the network load graphs generated in Step 1, add virtual links on the load graph where necessary. $CSat_{P'\_old}$ quantifies the effect of delaying all upgrades to future time T.
3. Conduct all upgrades on the network load graph of Step 1, and perform load readjustment to obtain the set of upgraded network load graphs.
4. For each customer access graph, mark the paths containing at least one candidate upgrade component as UPGRADE path. For any component downstream from an upgrade, with utilization level exceeds threshold κ, mark it as IMPACT component. The threshold κ is set to a value where noticeable performance deterioration is expected. Mark all paths on customer access graphs that intercepts an UPGRADE path at an IMPACT node as IMPACT path. A path identified as UPGRADE path cannot be marked as IMPACT path.
5. For each customer, for each path P' on the modified path list that is identified as either UPGRADE or IMPACT path, compute $CSat_{P'\_new}$ based on the upgraded network load graphs generated in Step 3, add virtual links where necessary. $CSat_{P'\_new}$ quantifies the effect of upgrades around future time T.

The effect of upgrades on customer flows can be evaluated as the changes to overall customer satisfaction associated with any modified path P' of a customer:

$$\Delta CSat\_P' = \left(\frac{(CSat_{P'\_new} - CSat_{P'\_old}) \times AccessT_{P'}}{\sum_{P' \in cust} AccessT_{P'}}\right) \times \beta_1 \tag{10}$$

For any UPGRADE path G, the revenue change ΔR_G is as considering the monetary benefit $Upgrade_G$ and penalty $Impact_G$ of upgrading G. Since $\Delta CSat\_G$ is normalized between 0 and 1, the values of $Upgrade_G$ and penalty $Impact_G$ are derived from $\Delta CSat\_G$ using linear revenue mapping. The increases in traffic volume downstream from an upgrade node on G may exert negative impacts on other customer flow paths intercepting G. $Impact_G$ considers these monetary penalties associated with all IMPACT paths affected by UPGRADE path G. For each customer, every such IMPACT paths V adds a fraction of its penalty to G, depending on the access time of all UPGRADE paths

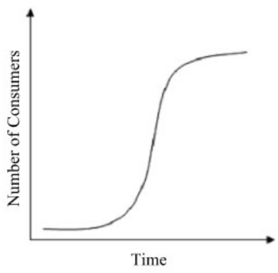

Figure 5: Bass model

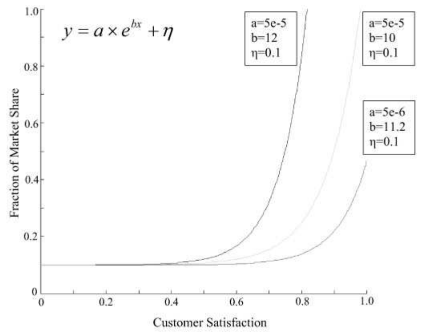

Figure 6: Customer satisfaction

U intercepting V at IMPACT nodes. SC_{cust} denotes the service charge of a customer in the relations below.

$$\Delta R_G = Upgrade_G + Impact_G, Upgrade_G = \Delta CSat\_G \times SC_{cust} \tag{11}$$

$$Impact_G = \sum_{cust} (\sum_{V \in cust} \frac{\Delta CSat\_V \times AccessT_G}{\sum_U AccessT_U} \times SC_{cust}) \tag{12}$$

Let μ be a upgrade component on a UPGRADE path G, $Cost\_U_\mu$ and $Cost\_M_\mu$ be the upgrade and additional maintenance costs of μ respectively, then the profitability of an UPGRADE path G $\Delta Profit_G$ is:

$$\Delta Profit_G = \frac{\Delta R_G}{\sum_{\mu \in G} (Cost\_U_\mu + Cost\_M_\mu)} \tag{13}$$

Two interesting observations can be made about $\Delta Profit_G$. First, a path upgrade is profitable only if its benefit is significant compared to its adverse effects on other flows. Second, a path is "expensive" to upgrade when it involves many candidate upgrade components. The profitability of upgrading a component x ($\Delta Profit_x$) can then be computed as the profitability aggregation of each path Q spanning over x. Let w denotes the number of upgrade components on a path.

$$\Delta Profit_x = \sum_Q \frac{1}{w} \Delta Profit_Q \tag{14}$$

5.5 Estimation of Population Growth

We estimate population growth based on a modified Bass growth model. The classic Bass model [5] divides the consumer market into innovators and imitators, where the innovators purchase the service regardless of service maturity and the imitators are attracted to the service due to positive responses from others (Figure 5).

Assuming all competitors in the market have equivalent pricing and technology attractiveness, we augment the Bass model by considering the presence of competitors and the customer satisfaction level of the service provider. Let N be the size of current subscriber population, S be the size of the consumer market, f_{inn} be the percentage of consumers that are innovators, p be the probability of innovator purchasing the service, q be the probability of imitator purchasing the service, t be the evaluation interval, M be the number of competitors in the market, and $CSat_{avg}$ be the average customer satisfaction, then the modified Bass model can be expressed as:

$$\frac{dN}{dt} = \frac{p}{M} f_{inn}(S-N) + q(ae^{b \times CSat_{avg}} + \eta)(1-f_{inn})\frac{N}{S}(S-N) \quad (15)$$

η is a natural growth factor, a fraction of potential growth regardless of customer satisfaction level; a and b are parameters that influence the competitiveness of the ISP. By varying the value of a and b, the "sensitive" regions of service satisfaction can be adjusted to best reflect the competition environment (Figure 6).

The modified Bass function takes mean customer satisfaction as input, and outputs the estimated number of new customers ΔN. The derivation of mean customer satisfaction before ($CSat_{avg\_old}$) and after ($CSat_{avg\_x}$) component upgrade is similar to component profitability computation. Then a modified profitability formula for upgrading component x ($\Delta Profit^*_x$) is:

$$\Delta N_x = Bass(CSat_{avg\_x}) - Bass(CSat_{avg\_old}) \quad (16)$$

$$\Delta Profit^*_x = \Delta Profit_x + \frac{\Delta N_x \times SC_{cust}}{Cost\_U_x + Cost\_M_x} \quad (17)$$

We only considered single-class service in the above computation. The analysis can be easily extended to multi-class service by computing $CSat_{avg\_x}$ and $CSat_{avg\_old}$ for each service class.

6. Maximizing Investment Benefit of Component Upgrades

In the previous section, we have computed the profitability of each candidate replacement component. As component profitabilities are independent and additive, we can maximize the investment benefit of the upgrade decision by reducing our investment decision problem to the classic "capital budgeting" problem. Our objective function is:

Given a set of candidate components $X = \{x_1, \ldots, x_n\}$*, its corresponding profit set* $P = \{\Delta Profit^*_{x_1}, \ldots, \Delta Profit^*_{x_n}\}$*, and upgrade cost set* $C = \{Cost_1, \ldots, Cost_n\}$

Determine a subset $\{y_1, \ldots, y_m\}$ *of X and subset* $\{Cost_1, \ldots, Cost_m\}$ *of C, such that:*

$$MAX(\sum_{i=1}^{m} \Delta Profit^*_{y_i}) \text{ and } \sum_{i=1}^{m} Cost_i \leq B, \text{where B is the fixed monetary budget.} \quad (18)$$

The optimal solution to capital budgeting is known to be NP-hard. However, efficient approximation algorithms exist by applying the 0-1 integer programming technique. Indeed, the formulation of our optimization model is similar to [11], but our derivation of investment benefit differs. More advanced solution techniques can be found in the works of [7] and [8]. Our decision process can benefit from such multi-objective optimization models by introducing additional constraints, expressed in terms of linear arithmetic expressions, such as:

$$\begin{aligned} &\sum(C_1, \ldots, C_{10}) \leq 5 \quad \text{At most five routers should be replaced} \\ &C_2 + C_5 \neq 1 \quad \text{Switch 2 and Switch 5 must be replaced together} \end{aligned} \quad (19)$$

As the optimization is performed on the profitability of upgrades derived from customer satisfaction, the solution not only maximizes return of investment, but also ensures visible service quality improvements to the customers, with regard to the specific network infrastructure and customer requirements.

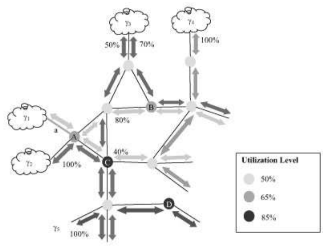

Figure 7: A simple scenario

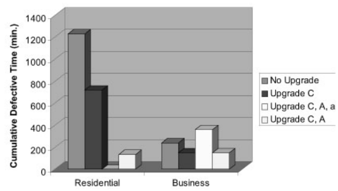

Figure 8: Cumulative service defective time

7. An Illustrated Case Scenario

In this section, we illustrate the effectiveness of our approach in comparison to traditional investment decision process. Figure 7 depicts a simple regional network consisting of four groups of customers.

γ_1 and γ_2 are residential areas with 10,000 users each, charged at $30 per user per month. γ_3 and γ_4 are business areas with 2,000 users each, charged at $300 per user per month. γ_5 is not part of the customer population. The utilization level of each component is shown in the figure, as well as the customer traffic flows and their access frequencies. We note that additional traffics exist in the network, whose effects on subscriber population are captured in terms of the load conditions on each component.

The ISP decides to perform network upgrades with a budget of $600,000. The link a and nodes A, B, C, and D are chosen as the candidate upgrade components. The cost of replacing B, C, D, and E is $300,000 each, and the cost of replacing A and a is $100,000 and $150,000 respectively. For simplicity, we assume no additional maintenance cost.

We set the benefit horizon T at 6 months after. The estimated network performance before and after upgrade is presented in Figure 9. The observation period is one week and the cumulative tolerance levels are computed from their normalized form. Residential users mostly use best effort traffic (α_1=0.5, α_2=0.5), while business users require support for highly interactive applications (α_1=0.4, α_2=0.6). The analysis is simplified to only consider the throughput and delay performance (i.e. β_1=1.0, β_2=0.0, and β_3=0.0).

Without customer flow information, a pure utilization based investment strategy would pick $\{$C,D$\}$, since they have the highest utilization level. We observe that replacing component D does not improve the service quality for the customers. After examine the customer flows, the set $\{$a,A,C$\}$ seems to be the best choice, as it would significantly improve the service quality for residential areas. We now assess component upgrade benefits and see if our approach confirms with this choice.

All paths except γ_4 are marked as UPGRADE paths, and γ_4 is marked as IMPACT path due to visible congestion at node downstream from B. We computed the profitability of upgrading each path, and then the profitability of each component upgrade (Figure 10). Note that $\gamma_4 - 100$ shows no profitability since it is not an UPGRADE path. Both $\gamma_1 - 80$ and $\gamma_3 - 70$ have an impact on $\gamma_4 - 100$, and hence suffer a fraction of $\gamma_4 - 100$'s penalty.

We skip the estimate on customer population growth, and directly apply the result obtained from Figure 10. It seems surprising that only component A and C should be

6 MONTHS AFTER (One Week)

Flow	Cumulative Throughput Defects (days)		Cumulative Delay Defects (days)		Throughput Tolerance (days)	Delay Tolerance (days)
	No Upgrade	Upgrade	No Upgrade	Upgrade		
γ_1-80	1.5	0.5	1.5	0.5	1.2	1.2
γ_1-40	1.5	0.25	1.5	0.25	0.6	0.6
γ_2-100	3.0	0.75	3.0	0.75	1.5	1.5
γ_3-50	0.725	0.325	0.4	0.225	0.375	0.188
γ_3-70	0.5	0.4	0.25	0.175	0.525	0.263
γ_4-100	0.5	0.725	0.25	0.45	0.75	0.375

Figure 9: Performance six months after

PROFITABILITY OF PATH UPGRADE

Flow	ΔCSat	Upgrade	Impact	ΔR	ΔProfit
γ_1-80	0.17	\$51,000	-\$72,000	-\$21,000	-0.038
γ_1-40	0.33	\$99,000	\$0	\$99,000	0.18
γ_2-100	1.0	\$300,000	\$0	\$300,000	0.75
γ_3-50	0.36	\$216,000	\$0	\$216,000	0.72
γ_3-70	0.0	\$0	-\$72,000	-\$72,000	-0.24
γ_4-100	-0.12	N/A	N/A	N/A	N/A

PROFITABILITY OF COMPONENT UPGRADE

Component	ΔProfit
a	-0.013
A	0.422
B	-0.253
C	1.155
D	0.0

Figure 10: Profitability computation

upgraded, but not link a. By tracing the computation of our analysis, the reason becomes apparent. While upgrading link a, we congest node B which will severely impact service quality for γ_3. The upgrade of node B will only marginally improve the performance for γ_3, at the cost of making γ_4 suffer a performance hit. Hence it is much more profitable to maintain γ_1's status quo, and delay the replacement of node B, since it has enough capacity to support the growth of γ_3.

The above case scenario is simulated using a packet level simulator. The links are all OC-1 and nodes are gigabit routers. Background traffics are introduced in the network to satisfy the load condition (before upgrades) illustrated in Figure 7. The load condition varies according to time of day. One day is simulated in the network (1440 min.). Delay statistics on the six flows are collected every 15 minutes. Figure 8 shows the cumulative service defective time for residential and business users under different upgrade options.

8. Conclusion

In this paper, we have discussed the deficiencies in current network upgrade strategies and presented a new customer-centric approach to making network upgrade decisions under budget constraint. By linking the performance of underlying network infrastructure to customer's perceived service quality, we established customer satisfaction as the primary criterion in the identification of candidate components and the assessment of upgrade benefits. Our mapping of network performance to customer satisfaction is unique to each customer, with regard to the particular service characteristics, customer preferences, and application requirements. Then, using customer flow analysis and network load readjustment, we established an efficient scheme to assess the profitabilities of upgrading network components, taking into account the interdependencies inherent to networks, the customer's access behavior, the changing network demands, as well as the potential customer population growth. Utilizing the results of our analysis, we can obtain sound investment decision that maximizes upgrade benefits. Through a case scenario, we have demonstrated the soundness of our approach compared to conventional practices.

Using our customer-centric approach, the resulting investment decisions enhance service quality, promote better customer satisfaction, and maximize investment return.

Some areas of this work can be further investigated. The linkage between customer satisfaction and revenue could be refined and modified to account for customer loss due to dissatisfaction. The assessment of component upgrade profitability could be extended to include the determination of best upgrade time for each candidate component. Such analysis would be extremely useful for incremental update scheduling. Although the simulation results show much promise, additional validations are planned on real regional networks to better evaluate the efficiency and effectiveness of our approach.

References

[1] Cisco netflow. http://www.cisco.com/warp/public/732/Tech/nmp/ netflow/index.shtml.

[2] A simple network management protocol (snmp). Technical Report RFC 1157, May 1990.

[3] Concise mib definitions. Technical Report RFC 1212, March 1991.

[4] Remote network monitoring management information base. Technical Report RFC 1757, Feburary 1995.

[5] F. M. Bass. New product growth model for consumer durables. *Management Science*, Vol.15:p215, 1969.

[6] C. Estan and G. Varghese. New directions in traffic measurement and accounting. In *ACM SIGCOMM'02*, 2002.

[7] R. F. Deckro J. L. Corner and R. W. Spahr. Multiple-objective linear programming in capital budgeting. *Advances in Mathematical Programming and Financial Planning*, Vol.3:p241–264, 1993.

[8] T. W. Lin. Multiple-criteria capital budgeting under risk. *Advances in Mathematical Programming and Financial Planning*, Vol.3:p231–239, 1993.

[9] H. Bailey R. Yelton A. Ajibulu M. Hopkins, G. Louth and M. Niva. A multifaceted approach to forecasting broadband demand and traffic. *IEEE Communications Magazine*, Vol.33:p36–42, Feb. 1995.

[10] K. Curran P. Flynn and T. Lunney. A decision support system for telecommunications. *International Journal of Network Management*, (No.12), 2002.

[11] J. Altmann S. Jagannathan and L. Rhodes. A revenue-based model for making resource investment decisions in ip networks. In *IEEE/IFIP Symposium on Integrated Network and System Management (IM)*, 2003.

[12] K. F. Li S. Khan and E. G. Manning. The utility model for adaptive multimedia systems. In *International Conference on Multimedia Modeling*, Nov. 1997.

[13] K. Stordahl and E. Murphy. Forecasting long-term demand for services in the residential market. *IEEE Communications Magazine*, Vol.33:p49, Feb. 1995.

[14] P. Forsyth Y. d'Halluin and K. R. Vetzal. Managing capacity for telecommunications networks under uncertainty. *IEEE/ACM Transactions on Networking*, Vol.10(No.4), August 2002.