

A REVENUE-BASED MODEL FOR MAKING RESOURCE INVESTMENT DECISIONS IN IP NETWORKS

Srinivasan Jagannathan,¹ Jörn Altmann,² and Lee Rhodes³

¹University of California Santa Barbara, ²University of California Berkeley, ³Hewlett-Packard Company

¹jsrini@cs.ucsb.edu, ²jorn.altmann@acm.org, ³lee_rhodes@hp.com

Abstract: Capacity planning is a critical task in network management. It identifies how much capacity is needed to match future traffic demand. It directly affects customer satisfaction and revenues. In this work we present a network usage analysis tool called *Dynamic Netvalue Analyzer (DNA)*, which helps alleviate a big problem that network engineers and marketing executives face—making optimal resource investment decisions. Marketing executives have to project customer growth while network engineers have to project traffic volume based on the entire customer population. DNA helps the prediction process by presenting actual network usage data from a business perspective, in a form that is useful to both network engineers and marketing executives. Using these projections, decisions on how to upgrade resources can be made. We show that information from DNA can be used to: (1) quantify revenue earned on each link, (2) quantify return-on-investment on performing a link upgrade, and (3) quantify the loss due to customer dissatisfaction when a link is not upgraded. We also illustrate how these formulations based on business information can be used to improve capacity planning decisions.

Keywords: Network and Systems Monitoring, Investment Cycle, Business Process, Network and Service Management

1. Introduction

The explosive increase in the number of Internet users as well as in volume of usage poses significant challenges to the network infrastructure and, by extension, to the network service providers. Network service providers are faced with two challenges today. On the one hand, they want larger number of customers in order to increase revenues. On the other hand, they want to manage the data volume efficiently. Capacity planning plays a crucial role in helping network providers tackle these challenges. Capacity planning is the process of predicting tomorrow's needs and preparing for them today. Network capacity planning involves combining marketing information and traffic analysis results to predict future resource requirements. Intelligent capacity planning can result in enormous cost savings and increased customer satisfaction. At the other extreme, poor capacity planning can result in enormous expenditures, poor customer service, and loss of revenues. The importance of capacity planning cannot be over-emphasized.

Capacity planning for Internet infrastructure requires good understanding of network traffic growth. The overall traffic growth depends on the number of new subscribers on the network as well as the usage increase per subscriber. Predicting these growth factors determines the decision on the investment size for upgrading network capacity. Currently, there is no standardized process to combine both growth factors.

Prediction of traffic growth and intelligent decision-making can be greatly facilitated by correlating network data with business-relevant information. Such information can be classified into:

- 1 Subscriber usage information: Considering usage of individual subscribers (or of subscriber segments) reveals more information than considering the aggregated usage of all subscribers.
- 2 Value of the subscriber to the business: The revenue and costs of an individual subscriber are important in evaluating the value of an investment. Clearly, an investment is worthwhile only if the revenue per subscriber outweighs the costs incurred per subscriber. Therefore analyzing the value of customers can help make educated investment decisions.
- 3 Competition in the market: Since competition in different geographical regions can impact subscriber loyalty, it is important to consider its impact before making costly investment decisions.

By associating such business information with raw traffic data, the network service provider can make better decisions about investments. However, today's capacity planning for the Internet does not incorporate this kind of information. The dominant reason is that the data collection and analysis process is not in place. The current process suffers from the problem that the data volume overwhelms conventional database management systems. Moreover, marketing and engineering disciplines of the service provider business lack a common vision. Marketing managers concentrate only on customer numbers and ignore the traffic-volume aspects of the business, while network engineers concentrate only on traffic volumes and ignore the customer aspect of the business. Each discipline views data in isolation, which results in myopic decision making. A holistic view of the data is necessary for informed investment decisions. In this paper, we present an innovative tool called *Dynamic Netvalue Analyzer (DNA)* that overcomes these lacunae. DNA aggregates, analyzes and models network data streams on the fly [1]. We show how this tool can be used to combine marketing, revenue, and engineering aspects of a service provider's business in order to make efficient capacity planning decisions.

In this paper, we only focus on capacity planning for regional networks. Our solutions are not directly applicable to capacity planning in backbone networks because of the following differences. First, the cost structures in these two categories of networks are dissimilar. For instance, laying optical fiber from coast to coast imposes very different costs than laying cable in an urban neighborhood. Second, the ratio between the traffic of an individual subscriber and the total traffic is significantly higher in case of a regional network. On the Internet backbone however, a huge number of individual flows is aggregated, thereby decreasing the impact of individual flows on overall traffic. Conversely, traffic of one individual subscriber in a regional network has much more impact on the overall traffic. Third, the size of the network and the volume of traffic in regional networks is much smaller. Therefore, capacity planning solutions for regional networks are not limited by scalability issues.

We now briefly discuss related work. Some of the previous work on capacity planning has focused on IT issues. Diao et al. propose an approach to maximize profits in service level agreements by designing feedback loops at application level [3]. Menasce and Almeida discuss performance issues and capacity planning for client-server systems [8]. On the network side, Robertazzi presents practical aspects of plan-

ning telecommunication and telephone networks [12]. Keshav discusses network capacity planning from the perspective of traffic management [7]. In addition, there are many software products that perform network capacity planning [11, 9, 6, 10].

The rest of the paper is organized as follows. In Section 2, we describe the capacity planning problem in greater detail. We describe current approaches to capacity planning and illustrate how a combination of marketing, revenue, and engineering analysis can greatly improve the planning process. In Section 3, we describe the DNA tool in more detail, and its versatility in collecting varied kinds of network data. In Section 4, we describe new algorithms to improve capacity planning decisions. We conclude the paper with a discussion on architecture and implementation issues in Section 5.

2. Capacity Planning

Network Capacity planning has three phases– (1) predicting future growth in customers, (2) predicting future volume of traffic, and (3) planning resource upgrades for the future. In the first phase, the marketing team estimates how many new customers will join the service and how many old customers will leave the service. The marketing team can use historical growth patterns, advertising budget, channel strength and customer satisfaction reviews, etc. to determine future growth and churn. This allows prediction of total number of users in the network. In the second phase, network engineers translate the number of customers into possible network traffic. This helps identify hot-spots in the network. Once the hot-spots are identified, in the third phase, the service provider must decide where investments are necessary in order to provide a good network service to customers.

For example, one simple approach to make investment decisions could be the following. The service provider sets a policy that all links should have overall utilization less than a threshold, $\tau_{suggested}$. Consider link l with capacity β_l that has a projected volume of traffic $b_{l,future}$ bits over the future time period T_{future} . Therefore, overall utilization of l , denoted by $\tau_{l,future}$ can be computed as $\frac{b_{l,future}}{T_{future} \beta_l}$. If this exceeds $\tau_{suggested}$, then the service provider marks the link to be upgraded. The quantum of upgrade, χ_l , should be greater than $\frac{1}{\tau_{suggested}} \times \frac{b_{l,future}}{T_{future}} - \beta_l$. This is obtained by solving the following inequality representing the service provider’s policy decision:

$$\tau_{suggested} \geq \frac{b_{l,future}}{T_{future} (\beta_l + \chi_l)} \quad (1)$$

Using some such mechanism, the service provider can construct a set of resources S , that need to be upgraded/purchased, and also determine the capacities of these resources. In this paper, we assume that the service provider has done some analysis to determine this initial set S of resources that need to be upgraded/purchased.

In the rest of this section, we examine the state-of-the-art in capacity planning and ways to improve it using an illustrative example.

2.1 State-of-the-Art

The example we present to describe the state-of-the-art is contrived and simplified, but it serves to illustrate the inefficiencies of capacity planning as practised today. Consider the example network shown in Figure 1

Figure 1 shows a small network with two edge routers, A , and B , and a border router C . There are two main links in the network– AC and BC . A back-up link AB




Figure 1. A Capacity Planning Example

is used for fault-tolerance. Customers are connected either to A or to B . Further, let us assume that there are two kinds of customers—residential and business. Residential customers pay a flat fee of \$20 every month, while business customers pay a flat fee of \$1000 every month. Currently, there are 20 business customers, 6 at A and 14 at B , and 200 residential customers, 150 at A and 50 at B . The problem of capacity planning involves analyzing existing traffic, predicting growth and making intelligent decisions on: (1) scheduling maintenance of AC and BC , (2) upgrading links AC , BC , AB , and routers A , B , and C , and (3) changing physical topology with new nodes and links.

Currently, service providers monitor the traffic using protocols like Simple Network Management Protocol (SNMP) [2]. A huge amount of data is collected, usually at 5 minute intervals. Let us assume the example network also uses SNMP and that a network engineer has analyzed SNMP data collected over the past 6 months (sampled in 5 minute intervals). The analysis result is that link AC has overall utilization of 60% and a peak utilization of 70% (over a 5 minute interval). Similarly BC has overall utilization of 51% and a peak utilization of 65%. The engineer also observes that link AB has less than 5% utilization.

The marketing team believes that it will acquire 5 residential customers each at A and B every month over the next 6 months. In addition, it estimates that it will acquire 1 business customer each at A and B every month over the next 6 months. Further, the marketing team believes that there will be no loss of customers in the next 6 months. This information is presented to the network engineer.

Currently, there is no industry-wide standard for interpreting marketing projections. Different service providers use different metrics. To the best of our knowledge (based on direct inquiries to network service providers), network engineers use some rule-of-thumb to convert marketing data to traffic volume predictions. For instance the rule-of-thumb may translate every new residential customer into 64kbps network capacity and every business customer into 512kbps network capacity. In addition they may assume that current customers' usage will increase 50% every 6 months. Some service providers use the traffic volume predictions obtained from such rules-of-thumb in sophisticated network simulation tools [11, 9, 6, 10] to analyze points of failure and then decide on investments and expenditures. But most service providers make decisions using simpler analysis. For instance, in this example, the network engineer can compute that after 6 months the requirement would be as follows. Current customers in A would require $0.6 \cdot 10240 \cdot 1.50 = 9216$ kbps. New residential customers will require $30 \cdot 64 = 192$ kbps, and new business customers will require $512 \cdot 6 = 3072$ kbps. This means that traffic on link AC will be 12480kbps which exceeds the link capacity. Similar analysis reveals that traffic on link BC will be 11098kbps which also exceeds the link capacity. Based on these results, the engineer may conclude that the capacity

of AC and BC needs to be doubled. Furthermore, this capacity increase may require router C also to be upgraded. Suppose that upgrading AC and BC costs \$20000 and \$10000 respectively. And let a high-capacity router cost \$20000. Thus, using this analysis, capacity planning expenditures total \$50000.

In summary, today network capacity planning is an art and not a science. We now present our vision for capacity planning.

2.2 Improving Capacity Planning

Our vision of capacity planning uses information that is available, yet unused. For example, the service provider can ascertain how much data on each link belongs to each customer segment. This can be done by observing the source or destination IP address of network flows and then correlating this address to the customer segment assigned that address. The network engineer generates a histogram to study how usage by residential customers has grown over the past 6 months. Using this information, the network engineer can use standard mathematical techniques of extrapolation to estimate how much volume of traffic new customers will generate. By performing a similar analysis, the engineer can also estimate how much data the business customers will generate. Thus the network engineer can estimate overall traffic growth in a more objective manner.

For the example presented in the previous subsection, let us assume that the engineer performs subscriber-specific usage analysis and diagnoses that usage per subscriber has not grown over the past 6 months for both residential as well as business customers. Based on this result, he predicts that the same trend will continue. Since the marketing team predicts negligible customer churn, and that there will be many new customers over the next 6 months, the service provider needs to analyze if current capacity can sustain future traffic. Based on past statistics, the overall utilization on both the links AC , and BC will exceed 50%. Therefore if more customers are added, the service provider may decide on upgrading one or both the links.

The monetary value generated by the customers can also provide critical information for capacity planning. In this example, customers at A generate $6*1000 + 150*20 = \$9000$ each month, and almost all the data they generate traverses AC . Similarly, customers at B generate $14*1000 + 50*20 = \$15000$ each month, and almost all the data they generate traverses link BC . Clearly, link BC has been more lucrative over the past 6 months. Furthermore, based on the marketing projections, adding new customers to region A will increase revenues by $(1000 + 5*20)*(1 + 2 + 3 + 4 + 5 + 6) = \23100 over the next 6 months. Similarly, adding new customers to B will increase revenues by $(1000 + 5*20)*(1 + 2 + 3 + 4 + 5 + 6) = \23100 over the next 6 months. For simplicity, let us assume that link capacities can only be doubled. Since doubling capacity of link BC will cost only \$10000, while gain in revenue is \$23100, it is worthwhile to double the capacity of link BC . On the other hand, in case of link AC , the gain in revenue almost matches the cost of doubling the capacity. Furthermore, if capacity of AC is also doubled, then the router will also need to be replaced. There will therefore be a loss if one invests in link AC and BC together. Based on this analysis, the service provider can decide to: (1) double capacity of only link BC , (2) ask marketing team to not campaign aggressively in region A , and (3) ask marketing team to intensify campaign in region B .

The example illustrates the power of combining marketing, revenue and network usage data in devising better capacity planning solutions. A tool that collects, ag-

gregates, and visualizes network usage and revenue data is a prerequisite for such a capacity planning scenario. In the next section, we describe one such tool.

3. Dynamic Netvalue Analyzer

HP OpenView Dynamic Netvalue Analyzer (DNA) is a business intelligence and decision support tool targeted for network service providers [1]. It transforms raw customer usage data into business information, supporting business managers and network engineers to model revenues and profitability for new and existing services as well as network capacity upgrades. It helps network service providers to understand the usage behavior of their subscribers in real-time. In general, it enables new revenue and return on investment paradigms.

DNA uses statistical models to analyze customer usage data. A statistical model can be thought of as a histogram-based distribution of observed data. By converting raw usage data into statistical models (streaming data analysis), instead of storing all the raw data as in a warehouse approach, DNA frees storage space. It thus enables the analysis of more historical data in a shorter time period. Also, a benefit of this approach is that statistical models are small in size and do not grow with increased traffic volume - instead they get more accurate. The statistical models are based on business dimensions such as pricing plans, services, or geography. The only drawback of using such statistical models is that it is static; the business dimensions have to be decided well in advance. If new business dimensions are selected then new data has to be collected to populate the statistical models along these dimensions.

DNA is comprised of a backend server, which aggregates usage data, and clients to view and model business decisions. The backend server is built on HP's Internet Usage Manager (IUM) mediation platform. At the data collection stage, three modules of IUM are used to aggregate customer information, session information, and the actual usage data of all end-users. The second level correlates the output of the previous stage and transforms the data into an internal format, which allows rapid access of individual usage data. The third stage is comprised of statistical models, which are specified by the user of the tool. The models are populated by the second stage's output. The third stage also interacts with clients, which request model data. DNA clients allow network usage data to be viewed in numerous formats with varying levels of detail. Data is presented in statistical histograms, tables, and summaries. For an in-depth analysis, it allows the business manager to view individual subscriber usage.

How DNA can improve communication between network engineers and marketing managers, especially with respect to the capacity planning process, can be demonstrated using the analysis example shown in Figure 2. Instead of aggregating all usage data across all subscribers to one number, DNA splits up the data into those business dimensions (e.g. pricing plan, link, subscriber, etc) that are important for the decision making process. Figure 2 illustrates the usage of individual subscribers over a time period of 30 days for one customer segment under a flat-rated pricing plan. By looking at this distribution and the historical trends (such as mean and variance), the network engineers can predict the usage of new subscribers in this customer segment. The marketing manager, who analyzes this figure, notices the imbalance between light users and heavy users. As a consequence, he can work on introducing pricing plans that resolve this problem.

In detail, DNA allows an on-the-fly analysis of raw data coming from a NetFlow-enabled router. Considering the example of Figure 1, DNA would analyze the NetFlow data coming from router A and B. It would aggregate usage data in units of bytes with

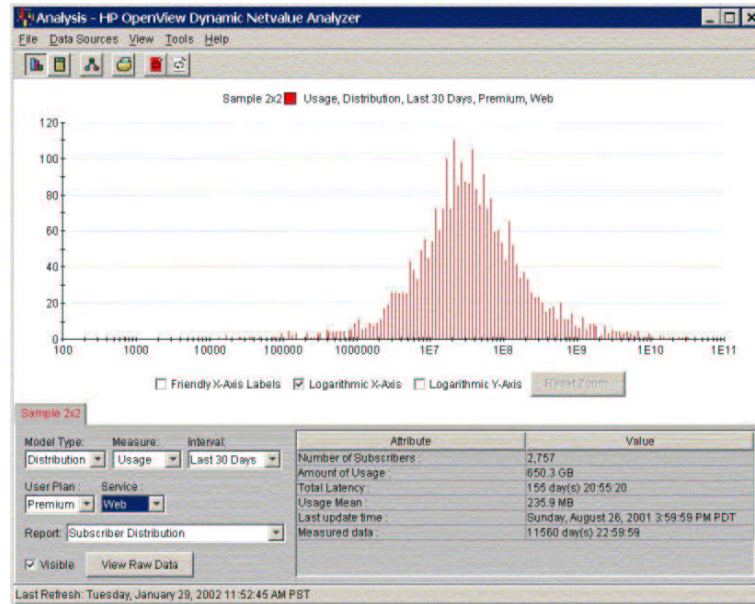


Figure 2. DNA analysis result of subscriber usage behavior

respect to the customer segment (based on the pricing plan), the destination IP address, and link congestion. The result of the aggregation would be used to make the decision for resource investments for all links as described in the following section.

4. DNA Decision Making Process for Resource Investments

We now present an approach to make smarter resource investment decisions. The approach assumes that there is an initial set of resources S that need to be upgraded and that the quantum of upgrade for each resource is known. Our main contribution is to associate a number called *investment gain* with each resource upgrade in S . The investment gain associated with a resource upgrade indicates how much benefit is accrued from the upgrade per unit money spent on that upgrade. The problem of making the right upgrade decisions then reduces to solving an optimization problem where the total investment gain is maximized subject to budgetary constraints on the investment costs. The investment gain $\Phi(l)$ for each link l in the network is based on two business criteria:

- 1 The size of the return-on-investment (RoI). The RoI depends on the quantum of upgrade, growth in customer base as well as on projected volume of traffic per customer segment.
- 2 The loss of customer satisfaction. The loss is directly correlated with the loss caused by not upgrading. Customers who are unhappy with the service may choose a different service provider thereby decreasing revenue.

The investment gain $\Phi(l)$ is calculated according to the following formula:

$$\Phi(l) = \frac{\gamma_{l,T_{future}} + \zeta_{l,T_{past}}}{FC_l + M_{l,T_{future}}} \quad (2)$$

where $\gamma_{l,T_{future}}$ is defined as the revenue earned on link l in future time period T_{future} , $\zeta_{l,T_{past}}$ represents the monetary value of customer unhappiness with link l during the past time T_{past} , FC_l is the fixed cost for upgrading link l by a quantum χ_l , and $M_{l,T_{future}}$ is the additional maintenance cost of l in time period T_{future} .

The numerator represents the sum of revenue earned and the monetary value of gain in customer satisfaction as a result of the upgrade. We assume that the quantum of upgrade is such that all the customers who were discontent with service before the upgrade are no longer discontent. The denominator represents the cost of the investment. The ratio reflects what is gained per unit of investment. The higher the ratio, the greater the priority for upgrade. Suppose that a network provider has to choose the right set of resource upgrades from set S , given a budget constraint B . For each resource r_i in S , we ascertain $\langle FC_i, \Phi(r_i) \rangle$ where FC_i is the cost of the upgrade and $\Phi(i)$ the investment gain. Let δ_i denote the decision to upgrade r_i , i.e., $\delta_i = 1$ represents a decision to upgrade and $\delta_i = 0$ otherwise. Then the network provider can choose the right set by solving the following optimization problem: choose $\delta_i \in \{0, 1\}$ to maximize $\sum \Phi(r_i) \delta_i$ subject to: $\sum FC_i \delta_i \leq B$. This optimization problem is a version of the 0-1 Knapsack problem which is known to be NP-Complete. If the cardinality of the set S is large, then any of the well known polynomial time approximation algorithms for 0-1 Knapsack problem can be used to solve the resource investment problem.

In our model, even though we focus on link upgrades, the approach is easily extended to include other resources like routers. Router investment costs can be subsumed¹ into the costs of the first link that necessitates a router investment.

In order to quantify return-on-investment and loss of customer satisfaction, we introduce two algorithms to compute $\gamma_{l,T_{future}}$ and $\zeta_{l,T_{past}}$ in the following two subsections.

4.1 Quantifying Loss of Customer Satisfaction

Rational customers are discontent with the service if the service provider does not meet QoS requirements. This occurs if one of the links of the data route is congested. By observing the degree of congestion at each link, we can compute the degree of customer dissatisfaction at a particular link.

For each link, we assume that the business manager sets a threshold. If the data transmitted per unit time on that link exceeds the threshold, then we shall assume that the link is congested. For example, the network provider can follow a policy that if the data transmitted on a link in a 5 minute interval exceeds 70% of the total link capacity then that link is considered to have been congested in that interval.

Since all customers who are using a link when it is congested are equally affected, we shall quantify customer dissatisfaction as follows. Compute the monetary value of bytes transmitted during congested intervals on a link. This represents revenue earned when customers are discontent with the service. Now, let us assume that the business manager weights the importance of customer satisfaction to his business on a scale

¹In this case FC_i is not known initially, and will vary with the candidate solutions of the knapsack problem.

between 0 and 1. Suppose this weight is θ . Then, the monetary value of customer dissatisfaction ζ can be computed as described in Figure 3.

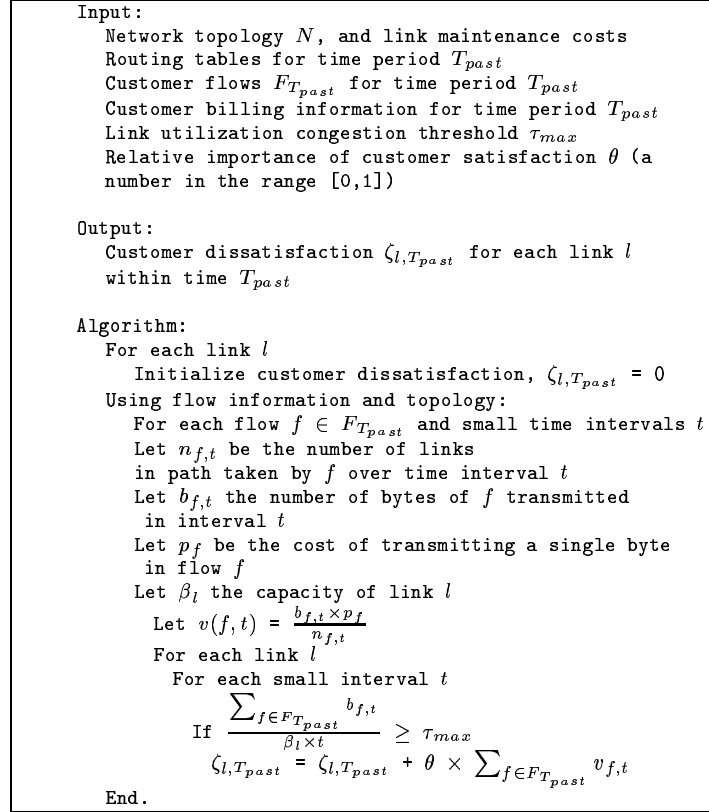


Figure 3. Computing customer dissatisfaction

4.2 Quantifying Return-on-Investment

Return-on-Investment depends on future revenues that the service provider hopes to earn. Therefore, to quantify RoI, we also need a mechanism to estimate future revenues. Suppose that there exists some mechanism to project growth of data volume as well as growth in the number of customers. Let us call this algorithm *Pred- Algo*. Then the return-on-investment can be computed as shown in Figure 4. The algorithm takes as input, the network topology, future routing information, current maintenance costs, quantum of upgrade for different links, the cost of these upgrades, and future maintenance costs. It produces as output, the projected revenue on all links for a future time period T_{future} .

The algorithm for quantifying return-on-investment assumes that we know how to quantify revenue earned on a link. To this end, we now develop an algorithm to distribute revenues among the links in the network. Our algorithm considers how much each link has been used by each customer. Then, using pricing plan information, the algorithm partitions the customer segment's payment to each of the links based on vol-

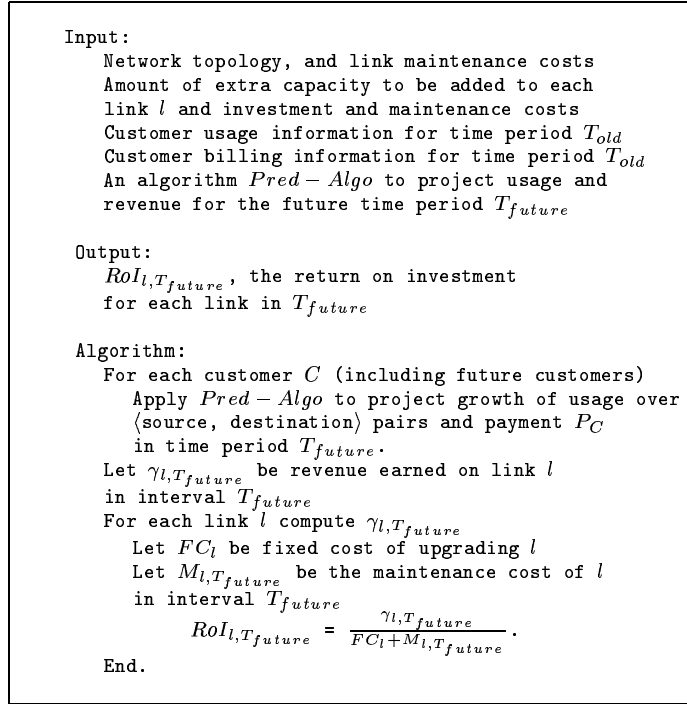


Figure 4. Estimating Return-on-Investment on upgraded links

ume of usage. The value of each link then, is the sum total of the money earned by that link from all customers. By comparing the value of each link with the investment and maintenance costs, we get a reasonable estimate of the importance of each link to the NSP's business over a given time period. The algorithm is presented in Figure 5

The algorithm outlined in Figure 5 estimates the price of each flow and equally distributes the payment among all the links in the path taken by that flow. Notice that the way we choose to distribute revenues and estimate value of a link is just for "accounting simplicity" and by no means an exact measure. It however gives us adequate perspective about which links are on popular revenue-generating routes. The algorithm in Figure 5 makes an implicit assumption that there exist alternative paths to every destination, and that should a link in one of the paths fail, the alternate path can sustain the traffic. Otherwise, a bridge link will be worth as much as the revenue earned from traffic traversing that link. The path redundancy assumption is justified because of two reasons: fault tolerance of network and computational complexity. Path redundancy is highly desirable from the perspective of fault tolerance. Hence almost every real-world network can be expected to be a strongly connected graph with redundant paths to every destination. The computational complexity of estimating whether an alternate path can sustain the traffic load is prohibitive². Assuming that alternate paths can sustain

²For each edge in a given path in an undirected graph, one can estimate the impact of removing an edge in that path, using a modified shortest path algorithm [4, 5]. But this approach cannot be used for directed graphs, and the problem is provably harder.

Revenue-Based Resource Investment Decisions in IP Networks

the traffic however allows us to perform a fast computation to estimate link value at the cost of some loss in accuracy.

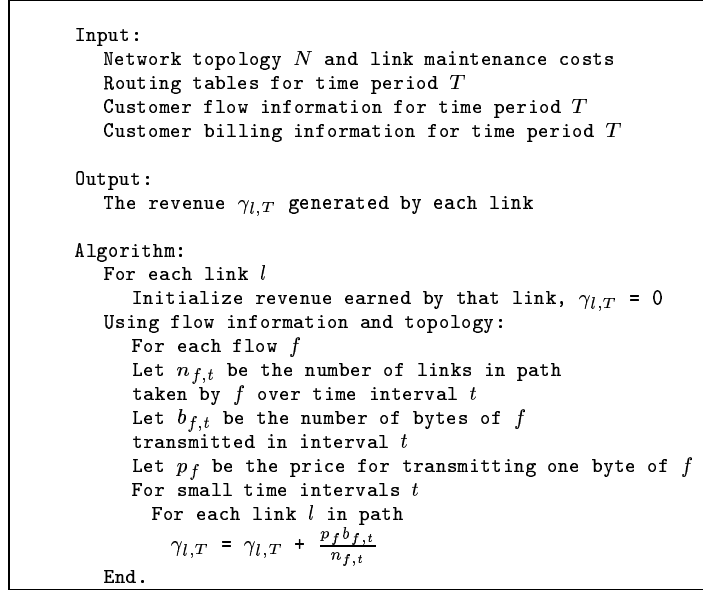


Figure 5. Distributing revenues among links

We now illustrate our algorithm using an example. Figure 6 shows a network service provider (NSP) with two customers—XYZ and ABC. XYZ pays a monthly rate of \$20000 for time critical data that it transmits to other locations. ABC pays a monthly rate of \$5000 as data transmission costs. The NSP pays \$5000 a month to the backbone provider for data transmission. The maintenance costs for each link are indicated in the figure. The figure also shows how much data has been transmitted by each customer (say in the past month) and which links the data traverses on its way to destinations. XYZ sends/receives 50MB data on link A-B and A-C and 200 MB data on link B-D. ABC sends/receives 300 MB data on link A-C. For simplicity, we shall ignore the links connecting ABC and XYZ to nodes A and B respectively.




Figure 6. Estimating the Value of a Link

Let us use algorithm to estimate the revenue from each link in this network. We assume that the entire volume of data over a link was generated in a single flow. Since payment of customer ABC is \$5000, the price of each Megabyte is $p_{ABC} = \frac{5000}{300} = \16.67 . For customer XYZ, the payment is \$20000. Hence $p_{XYZ} = \frac{20000}{250} = \80 . ABC has

only one flow of volume $300MB$ over path A-C which contributes \$5000 to $\gamma_{A-C,T}$. Customer XYZ has two flows, one over B-A-C and the other over B-D. The former has a data volume of $50MB$ and contributes $\frac{1}{2} \times \frac{50}{250} \times 20000 = \2000 each to $\gamma_{A-B,T}$ and $\gamma_{A-C,T}$. Similarly, the second flow contributes $\frac{200}{250} \times 20000 = \16000 to $\gamma_{B-D,T}$. Thus we have, $\gamma_{A-C,T} = \$7000$, $\gamma_{A-B,T} = \$2000$, and $\gamma_{B-D,T} = \$16000$.

5. Discussion

In this paper, we have described a new approach to capacity planning. We illustrated the power of combining marketing, revenue, and customer usage information. We then showed that the Dynamic Netvalue Analyzer (DNA), a tool that can aggregate and analyze raw network data on-the-fly, can be used for this new approach to capacity planning. In more detail, we outlined algorithms that use DNA-aggregated data for making investment decisions.

We now discuss some architectural issues and tradeoffs associated with some possible implementations. An important issue that arises is the location of the data collecting agents within the network. Network managers are very reluctant to make changes to the network or install monitoring equipment on important or busy routers. Our architecture does not require DNA at any of the internal or border routers. DNA runs at the network edge. Because all traffic has to enter or exit through one of these edge-routers, information collected from these routers is sufficient to capture the entire network scenario. Our algorithms do require routing information from the network. One possible approach is to get a daily update of routing tables from all routers in the network. A daily snapshot will suffice if internal routes are relatively stable over time. Even in case of route instability over short intervals of time, our analytical results will not be significantly affected because capacity planning is a long-term process and uses data over long intervals of time (often using data gathered over months).

The algorithms we have described are implementation independent. They can be implemented over a sophisticated network simulation tool, or they can be implemented inside the monitoring tool itself, thus performing all the analysis on-the-fly. A simulation based approach has significant data storage and transportation overheads. To perform a realistic simulation of observed traffic, we need detailed information about network flows to be collected, transported and stored at a centralized location. Furthermore, the simulation will be computation intensive and could be time consuming. However, since capacity planning is a long-term process, a dedicated machine can accomplish the task. A possible optimization is to use “aggregated network data” or traffic models instead of detailed flow information. This will reduce the accuracy of the analysis, but will also significantly reduce the data transportation and storage costs. DNA can be used to perform such traffic modelling [1].

If the second approach is adopted, our algorithms that compute monetary value of links and customer dissatisfaction are implemented within DNA. This may require approximations of the price of transmitted data because the actual price charged for transmission may not be known in real-time (e.g. tiered-pricing plans). Moreover, some coordination may be required with a network management protocol like SNMP. Specifically, the network management protocol should trigger alarms whenever link utilization over observation intervals exceed pre-configured thresholds. This is required to compute the customer dissatisfaction with links. But at the cost of administrative overhead, this approach saves significant data transportation and storage costs.

An important issue concerning our capacity planning vision is in evaluating the trade-off between information utility and data collection overhead. Data collection and

analysis, if done efficiently, can deliver business information, whose value outweighs the cost of the data collection. We believe that DNA and our algorithms for investment decisions deliver this efficiency.

References

- [1] J. Altmann and L. Rhodes, "Dynamic Netvalue Analyzer - A Pricing Plan Modeling Tool for ISPs Using Actual Network Usage Data", IEEE WECWIS2002, International Workshop on Advance Issues of E-Commerce and Web-Based Information Systems, 2002.
- [2] J. Case, K. McCloghrie, M. Rose and S. Waldbusser, "Protocol Operations for Version 2 of the Simple Network Management Protocol (SNMPv2)", RFC 1905, 1996.
- [3] Y. Diao, J. L. Hellerstein, and S. Parekh, "A Business-Oriented Approach to the Design of Feedback Loops for Performance Management", Distributed Operations and Management, 2001.
- [4] J. Hershberger and S. Suri, "Vickrey Prices and Shortest Paths: What is an Edge Worth?", IEEE Symposium on Foundations of Computer Science, 2001.
- [5] J. Hershberger, S. Suri and A. Bhosle., "On the Difficulty of some Shortest Path Problems," STACS, Berlin, 2003.
- [6] IPAT, http://www.wandl.com/html/ipat/IPAT_new.cfm
- [7] S. Keshav, "An Engineering Approach to Computer Networking: ATM Networks, the Internet, and the Telephone Network", Addison-Wesley, 1997.
- [8] D. A. Menasce and V. A. F. Almeida, "Capacity Planning for Web Services: Metrics, Models, and Methods", Prentice Hall PTR, 2002.
- [9] NetQuad, <http://www.anitetelecoms.com/products/NetQuad.html>
- [10] NetRule, <http://www.analyticalengines.com/>
- [11] Opnet, <http://www.opnet.com> .
- [12] T. Robertazzi, "Planning Telecommunication Networks", Wiley-IEEE Press, February 1999.