CyberPlanner: A Comprehensive Toolkit For Network Service Providers

Raouf Boutaba, Jin Xiao, and Issam Aib

David R. Cheriton School of Computer Science
University of Waterloo
200 University Ave. W., Waterloo, ON, Canada
Email:{rboutaba,j2xiao,iaib}@cs.uwaterloo.ca

***Abstract*—With the increased network size and diversity, and the proliferation in applications and services, the network service providers are faced with a flood of information from many levels of their network and service operations, often in uncorrelated forms. At the same time, the focus of network services is shifting from managing networks to managing services and customers. CyberPlanner[1] is a comprehensive toolkit for the collection, correlation and filtering of metrics at all levels of a network service provider's business. In this paper, we present its concept, modeling methodologies and implementation. Experience gained from CyberPlanner testing and operations in both academic and industrial settings demonstrates its effectiveness and potential in assisting many key provider operations such as customer-oriented trouble diagnostics, network dimensioning and upgrading, service planning and business revenue forecasting.**

***Index Terms*—Service management, metrics, service and customer modeling**

I. Introduction

The ability to collect, correlate and present information in an effective and timely manner is paramount to the success of today's network service management. With the increased network size, the diversity in network technologies, and the proliferation in user applications and services, the network service providers are becoming overwhelmed with a flood of information coming from many levels of their network and service operations. Often in uncorrelated and unfiltered forms, these information come from a wide array of performance measures ranging from low level network device MIBs, service level data probing and flow tracking logs, to customer trouble ticketing and accounting, market reports, etc. There is the apparent need for a comprehensive tool that can organize, correlate and filter/aggregate these information in flexible and meaningful ways, and present them through graphical and human-friendly interfaces. Furthermore, with the focus of network services shifting from managing networks to managing services and customers, there lacks a good understanding of the relations among network performance, customer satisfaction and business profitability. Thus high level decision questions such as: "how can we re-dimension our networks so that our DSL service revenue could be raised by 10%?"[2] can not be resolved or even readily analyzed today. These factors significantly impact the effectiveness and efficiency of many management and business processes in today's network service provider operations.

CyberPlanner is a six-year project jointly funded by Bell Canada, NSERC[3] and CITO[4]. Participants include researchers from University of Waterloo, Bell Canada, and POSTECH. The aim of the project is two fold. One, to create a pragmatic toolkit that can collect, correlate and filter key metrics at all levels of the network service provider business, and to provide customizable yet unified presentations of the network and service "health" information to a wide range of personnel (e.g. network administrators, business planners, technicians, customer support staffs, etc.). Two, to establish theoretical and analytical correlations among metrics, especially among key metrics at the network, service and customer levels. Because of its comprehensive nature, the CyberPlanner toolkit can assist in many crucial network adminstration and planning operations such as customer-oriented trouble diagnostics, network dimensioning and upgrading, service planning and business revenue forecasting, etc. At the core of CyberPlanner is the concept of Metrics Tree Models (MTMs) that structure and correlate metrics at the technology, network, service, customer and business levels of a network service provider's business. Metrics Tree Instances (MTIs) are constructed based on a large library of MTMs that model the specifics of the provider's network topology, technology composition, service types, and customer profiles. Experience gained from preliminary testing in industrial and academic setting shows the effectiveness and potential of CyberPlanner. In this paper, we explain the metrics tree concept central to CyberPlanner and detail some of the metrics correlations. Furthermore, we present the architecture and implementation of the CyberPlanner prototype and its key operations. Finally, we discuss our experience with the testing and operations of CyberPlanner. Much of our previous work [1][2] have focused on the specific metrics correlations, especially at the

[1]The project jointly is sponsored by Bell Canada, NSERC CRDPJ grant 311309-04 and CITO 7.2 RPP grant 140381-01.

[2]A similar question was posted to the project by network planners in Bell Canada

[3]Natural Science and Engineering Research Council of Canada

[4]Communications and Information Technology Ontario

service, customer and business level, thus we will not delve into metrics correlation specifics in this paper.

The rest of the paper is organized as follows. Section II presents the metrics tree concept and related works, Section III provides details on example MTMs, while Section IV shows the CyberPlanner architecture and implementation. Section V discusses our experience with CyberPlanner testing and operations. Section VI concludes the paper.

II. The CyberPlanner Concept

A. The metrics layering

At the heart of CyberPlanner are the Metrics Tree Models (MTMs) that describe the correlation among metrics. Metrics are concrete, quantifiable and specific performance measures whose definitions and interpretations are context specific. For example, at the network device level, metrics are used to record the various performance attributes of a DSL box or a IP router; at the network level, metrics are used to measure network-wide performance attributes such as end-to-end delay or throughput; at the customer and service level, they are defined to track customer satisfaction and customer care efficiency, etc. Therefore, we depict the metrics as a layered concept (Figure 1), where each layer is concerned with a specific facet of the network service provider's operations:

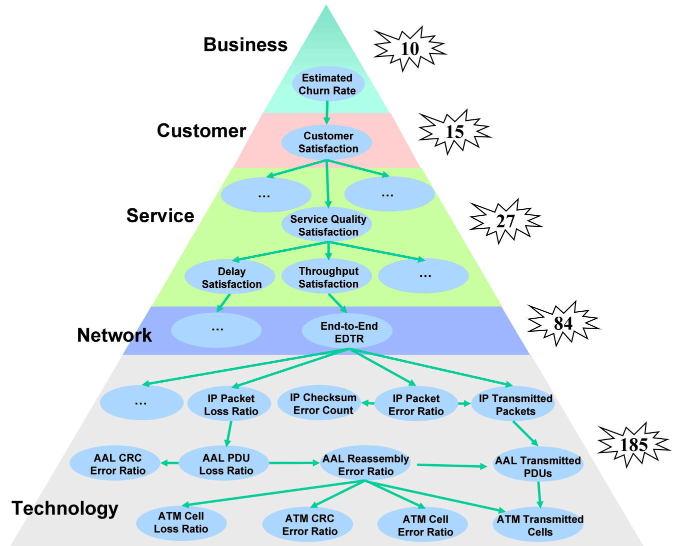

Fig. 1. Example Metrics Tree

- **Technology layer**: is concerned with the performance characterization of specific network technologies including IP, ATM, DSL, Ethernet, MPLS, etc. An individual network element such as a router, a switch or a link is modeled based on its technology compositions (e.g. the protocol stack). Interval unidirectional/bidirectional connectivity, packet transfer delay, and packet error ratio are some of the example metrics in this layer.

- **Network layer**: deals with the measurement of end-to-end quality of service in the networks. A network metric, such as one way end-to-end delay, end-to-end packet loss ratio, relates to a collection of network elements that together compose the end-to-end path spanning two points in the network.

- **Service layer**: contains metrics describing the service performance of a particular network service bundle (e.g. xDSL service). They measure aspects including service quality, service availability, customer care, etc. Metrics at the service layer are technology neutral and service specific. Some example metrics are mean time between service failures, billing accuracy, etc.

- **Customer layer**: relates customer specific metrics including service fulfillment metrics (e.g. perceived service quality), service assurance metrics (e.g. helpdesk response time for the customer), and billing metrics (e.g. billing accuracy). A key metric in this layer is customer satisfaction, which quantifies the likelihood of the customer to continue his service in near future. Some details on the computation of this metric is presented in Section III.

- **Business layer**: has metrics organized into three categories: customers, profit, and CAB (Charging, Accounting and Billing). Customer metrics provide an overview of the customer related information such as average customer satisfaction and potential churn rate. Profit metrics present general evaluation on the profitability of the services and the CAB metrics report on the efficiency of the bills collection and accuracy.

In Figure 1, we show the total number of metrics currently present in each layer. The number of metrics at the technology layer is significantly larger due to the diverse technology combinations and their many performance attributes. In contrast, the higher layer metrics tend to be some form of aggregation or abstraction of the lower layer metrics (e.g. end-to-end throughput is jointly influenced by metrics in the lower layers such as packet loss rate, packet error rate, and effective data throughput). Thus we depict the overall structure as a pyramid. In accordance to the design of a layered structure, when defining the MTMs, a metric has correlation with metrics in the same layer and/or the layers directly above and below. The ordering of the business, customer and service layers in CyberPlanner reflects a customer-oriented view of the network service operations. The customers are the foundations upon which the business objectives can be realized (e.g. customer satisfaction and potential churn directly impact business revenue), while the services are the vehicles with which the operators fulfill their customers' needs.

A Metrics Tree Instance (MTI) is formed in a top-down manner, by first taking a metric of interest (at any layer of the pyramid) as the root and then attaching to it the MTMs of the correlated metrics at each layers. Figure 1 shows an example where the root metric is the estimated churn rate at the business layer. It is correlated to customer satisfaction at the customer layer, which is determined based on the type of services the customer uses (the customer is using xDSL service in this example and hence the MTM for xDSL service is attached). One aspect of the xDSL service quality is throughput satisfaction which is then correlated with the end-to-end effective data throughput (EDTR) at the network layer. And this metric is correlated to the network elements that may service this customer. In this example, we show one such correlated MTM: IP over ATM. It is apparent that a MTI could be very detailed down to the technology level of the provider network. Thus, the MTMs yield a powerful modeling construct with sufficient level of details for modeling and understanding the dynamics of the network service provider's operations, and assessing the impact of network performance changes on the customers and the business objectives. In practice, the MTI is fully customizable such that much of the modeling details can be filtered away and only enabled on-demand. This aspect is further elaborated in Section IV.

B. background and related works

A considerable amount of standardization efforts have been carried out by various organizations to characterize network performance. IETF defines a set of criteria and terminologies to facilitate a common understanding of the IP network performance [3]. Metrics such as one way delay and one way packet loss are given concise treatment [4][5][6]. There also exists technology specific bench-marking techniques [7][8][9]. These collective efforts are network technology specific, and the measurement data from these facilities are uncorrelated. Thus much human efforts are involved in understanding the network construction and then accessing various parts of the network to assemble the raw data into a uniform presentation. Not only are such activities human-intensive, they are also extremely time-consuming and error-prone.

Some works (e.g. [10][11][12]) have emphasized on the importance of analyzing both the customer and network profiles in business decision processes of a network service provider. However, there are no methodologies on metrics correlation between the network and service profiles, and the derivations of customer utility models are very rudimentary. IST projects, such as TEQUILA, have investigated the monitoring and analysis of specific end-to-end network metrics for traffic engineering and quality assurance. TMF's efforts in Operational Support Systems (OSS) cover some aspects of network service and business issues such as customer care and CAB. However, since the focus is on defining methodologies to structure network service business processes, metrics modeling and correlations are not specifically investigated. Through the course of the CyberPlanner project, we have investigated in great length the relations among service performance, customer behavior and business revenue. Some methodologies for estimating customer satisfaction from network performance are presented in [1] and a set of comprehensive models of network performance and customer behavior are discussed in [2].

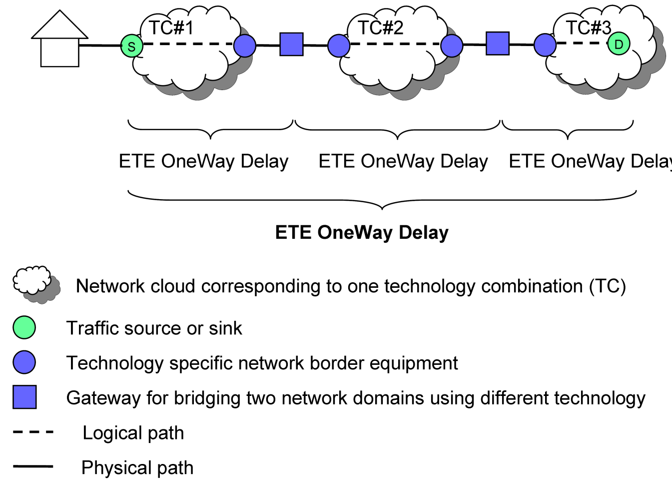

Fig. 2. Horizontal Division for ETE One Way Delay Computation

III. THE METRICS TREE MODELS

In this section, we present details of the MTMs and their relations. Rather than attempt to present all of the metrics, which is vast, the focus will be on the metric modeling methodologies and some detailed examples are provided. We first discuss the MTMs of the network and technology layer which are network and device dependent, then we present the modeling of metrics in the service, customer and business layers.

A. The network and the technology layer

A MTM at the network layer typically involves some form of end-to-end performance metric (e.g. end-to-end one way delay). This metric is related to a particular network path between two end points in the network and can consist of multiple technologies. Hence the general approach we take is to break down the end-to-end network path into consecutive domains of common technologies (TCs) and then address each domain separately. Figure 2 depicts our methodology. This is the *horizontal division*. The computation of the end-to-end metric is then the aggregation of the individual metrics in the domains. Taking the end-to-end one way delay example, its computation is then:

$$OneWayDelay = \sum_{TC} IPOneWayDelay$$

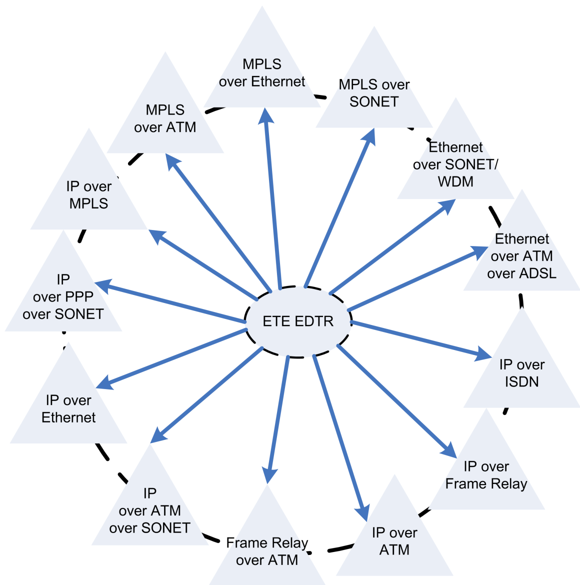

Fig. 3. MTMs of Different Technology Combinations

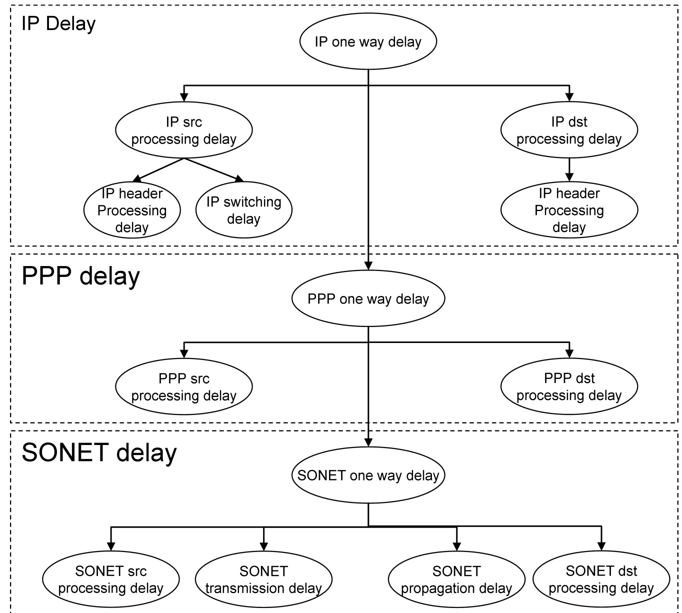

Fig. 4. IP over PPP over SONET Metrics Tree Model

The specific technology combination in a TC domain depends on the particular type of network elements being deployed. Currently, CyberPlanner includes a sizable library of MTMs for various technology combinations (Figure 3). The IP over PPP over SONET MTM is shown in Figure 4 as an example. The correlations of the metrics are defined as follows:

$$IPOneWayDelay(L) = IPSrcProcDelay(L) + IPDstProcDelay(L) + PPPOneWayDelay(L+11)$$

$$PPPOneWayDelay(L+11) = PPPSrcProcDelay(L+10) + PPPDstProcDelay(L+10) + SONETOneWayDelay$$

$$SONETOneWayDelay = SONETSrcProcDelay + SONETDstProcDelay + SONETTransmissionDelay + SONETPropagationDelay$$

$$SONETTransmissionDelay = n \times 125\mu 2$$

$$SONETPropagationDelay = 5ms \times \frac{SONETPathLength(km)}{1000}$$

L: IP Packet Length
P: SONET Payload Length
n: Number of SONET frames required

$$n = \left\lceil \frac{L}{P} \right\rceil + 1$$

Thus a MTM for end-to-end one way delay depends on the MTMs of many technology combinations. An instantiated MTI for the metric would then involve a list of specific MTMs customized to the particular technology combinations required. Due to limitation of space, we have shown a very simple example, most of the technology MTMs are much more complex (one such example can be seen in Figure 1).

B. The service, customer and business layers

Network performance metrics are associated with the service layer metrics through a set of MTMs for specific network services. Customers of the service derive their customer satisfaction based on the performance of the services they subscribe to. Figure 5 shows an example xDSL service MTM for xDSL subscribers. Service quality, service availability and customer care are the three xDSL service aspects that contribute to the customer's satisfaction. The service quality is related to throughput and delay performance metrics at the network layer, and the evaluation follows a threshold based mechanism under which *defective service intervals* [1] are computed. The determination of defective service intervals is based on whether the network performance metrics meet the SLA or SLS specification of the minimum performance guarantees. Our prior work [1] details the constructions of service quality from network performance.

Customer satisfaction is then modeled as the interaction between perceived utility and customer expectation. The perceived utility is related to service performance as:

$$xDSLPerceivedUtil = \beta_1 \times PerfQual + \beta_2 \times AvailQual + \beta_3 \times CCareQual$$

The parameter β is used to specify the customer's preference to each of the service quality aspects, with the condition $\beta_1 + \beta_2 + \beta_3 = 1$. In [2], we have detailed the modeling of service performance, customer satisfaction and business revenue. The overall process is depicted in Figure 6. the computation of

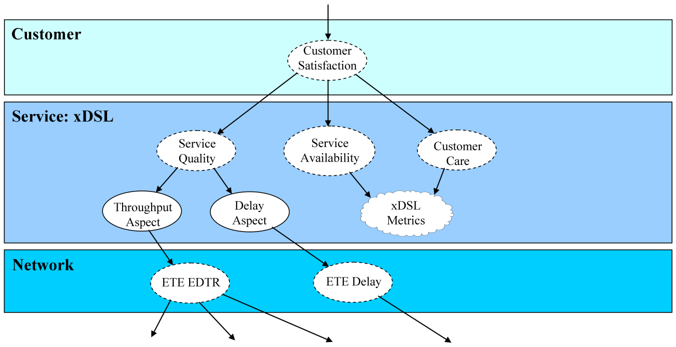

Fig. 5. Example xDSL Service Metric Tree Model

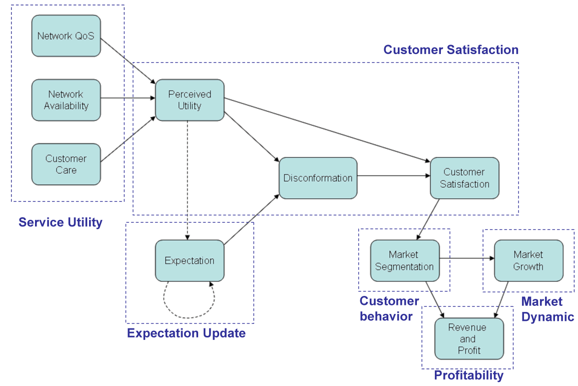

Fig. 6. Customer Satisfaction and Business Revenue Model

customer satisfaction is a function of perceived utility (f_1) and disconfirmation (f_2):

$$CustomerSat = f_1(PerceivedUtil) + f_2(PerceivedUtil - ExpectedUtil)$$

$$f_1(x) = \frac{\mu_1}{6}x^3 - \frac{\mu_1}{2}x^2 + \left(\frac{\mu_1}{3} + \omega_p\right)x$$

where $\mu_1 \leq 6\omega_p$, $\mu_1 \geq 0$, *and* $\omega_p > 0$

$$f_2(x) = \begin{cases} \omega_{dp}x & x \geq 0 \\ \omega_{dn}(x+1)^{\mu_3} - \omega_{dn} & x \leq 0 \end{cases}$$

where $\mu_3 \geq 0$, *and* ω_{dp},$\omega_{dn} > 0$

The disconfirmation function models the dissatisfaction a customer feels when his perceived service utility does not match his expectation. Through an expectation update process, the expectation of a customer is formed and adjusted over time based on the customer's past service experience. The effect of the customer satisfaction is that the service provider gains popularity based on customer responses, which in turn influences the customer's repurchase intentions and the attractiveness of the provider's service to non-subscribers. This effect is further moderated by the market condition to form a forecast of the subscriber market segmentation in the near future. Business level metrics such as potential churn rate and estimated revenue is then derived from the result of the market segmentation.

IV. The CyberPlanner Architecture and Implementation

The CyberPlanner architecture consists of a service core, a topology editor, and a database (Figure 7). The topology editor (Figure 8) is capable of creating network elements of various technology combinations and generating complete network topologies. The database hosts internal data to the CyberPlanner including network topologies, MTMs, saved MTIs, etc. The service core consists of a set of service modules whose invocations are facilitated through a common request dispatcher. The request dispatcher has both a Web interface and a Web Service interface with which management applications and human users could create and manipulate MTIs. Some of the key service modules in the service core are:

- **Topology management**: deals with the storage and retrieval of network topology data. It is primarily invoked by the topology editor through the request dispatcher, but is also requested by the metrics tree management module to supply topology information for MTI creation, and by the metrics tree operations module for MTI computation.

- **Data management**: facilitates database access for the storage, retrieval and manipulation of network topology information, MTMs, stored MTIs and various associated metrics data. The data management module also takes care of data collection from data sources or network simulators.

- **Metrics tree management**: is responsible for the generation of new MTIs based on user/application requests, and the storage and retrieval of stored MTIs. A component called the Instance Factory resides here that facilitates MTI generations based on corresponding network topologies.

- **Metrics tree operations**: takes care of MTI manipulations such as metrics computation, metrics data update, and metrics parameter configuration.

- **Session management**: is responsible for maintaining application/user sessions involving one or more MTIs and to ensure multiple sessions operating on the same MTIs are synchronized and do not cause database inconsistencies. this module is currently a stub.

Each metric of a MTI in CyberPlanner has a type. Of particular interests are the types *Normal* and *Default* which have implications on MTI creation and computation. When a metric is set to *Normal*, it indicates that this metric is to be computed from its correlated metrics, and hence when the MTI is constructed, this metric and its correlated metrics are

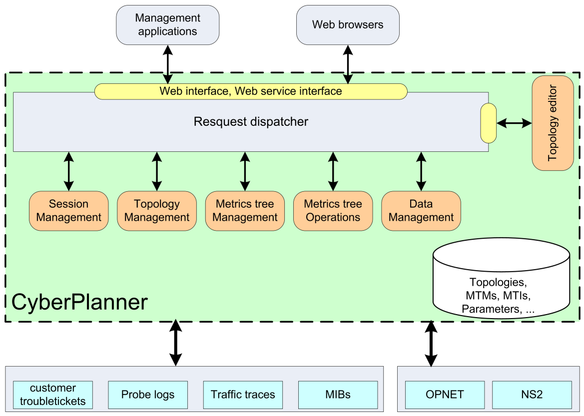

Fig. 7. The CyberPlanner Architecture

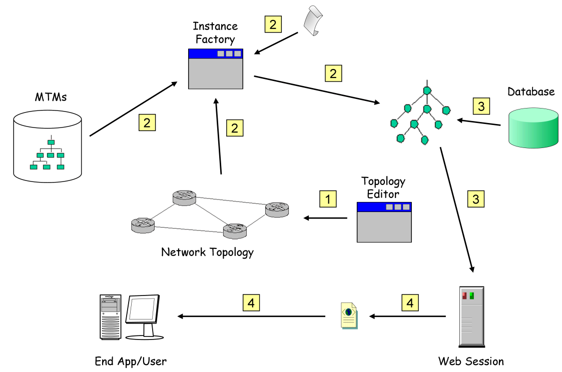

Fig. 9. The MTI Creation in CyberPlanner

all instantiated. When a metric is set to type *Default*, it is treated as a leaf metric during MTI instantiation and hence no correlated metrics at the lower layers are instantiated. During computation, the value of this metric is read from an external data source. In a nutshell, a metric with type *Normal* is computed based on MTMs and a metric with type *Default* is measured. Thus, the size and extensiveness of any MTI is customizable by the application/user depending on the level of details required. This is particularly useful for troubleshooting and network diagnostics where detailed information are required only in specific areas. The metric types *Minimum* and *Maximum* are similar to the type *Default*, except that their values are taken as theoretical values rather than real measurements.

Two key operations in CyberPlanner are the *Create_MTI* operation and the *Compute_MTI* operation. The *Create_MTI* operation is depicted in Figure 9. The topology editor is used to create a particular network topology at some prior time (Step 1). When a request for MTI creation is received by the metrics tree management module, the instance factory component takes as input the root metrics for the new MTI, the corresponding network topology ID in the CyberPlanner database, the customer traffic flows, the service information (e.g. types of services), and the metrics configuration parameters. A MTI is then generated (Step 2). The newly created MTI is established as a service session and the *Compute_MTI* operation is invoked by the metrics tree operations module (Step 3). The session and its initial metrics information is then presented to the end user/application (Step 4). The *Compute_MTI* operation involves the re-computation of the MTI, the metrics values are refreshed depending on their types (i.e. computed, obtained from measurement, or use a theoretical value). A newly created MTI could be stored in the CyberPlanner database for future use. Figure 10 shows the graphical web interface a user is presented with. The left panel shows a directory style display of the metrics and their correlations. The value and type of the currently selected

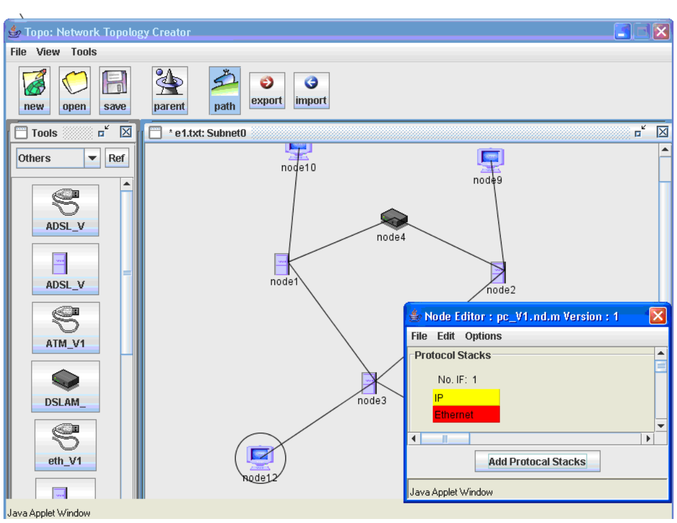

Fig. 8. The Topology Editor

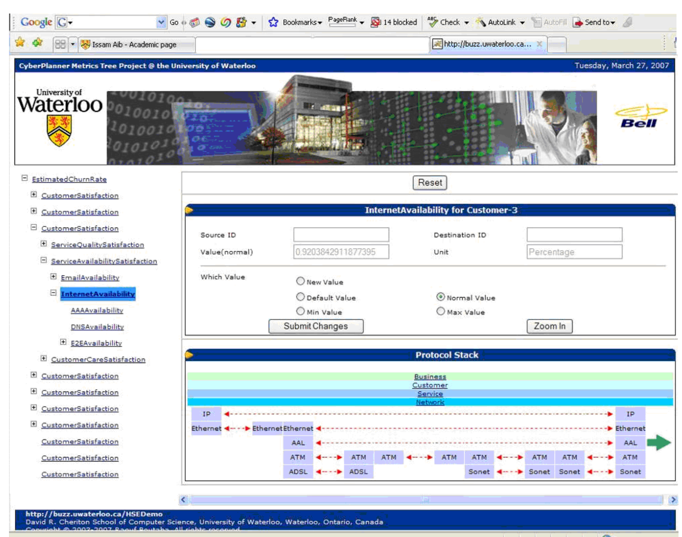

Fig. 10. The Web-based Graphical Presentation in CyberPlanner

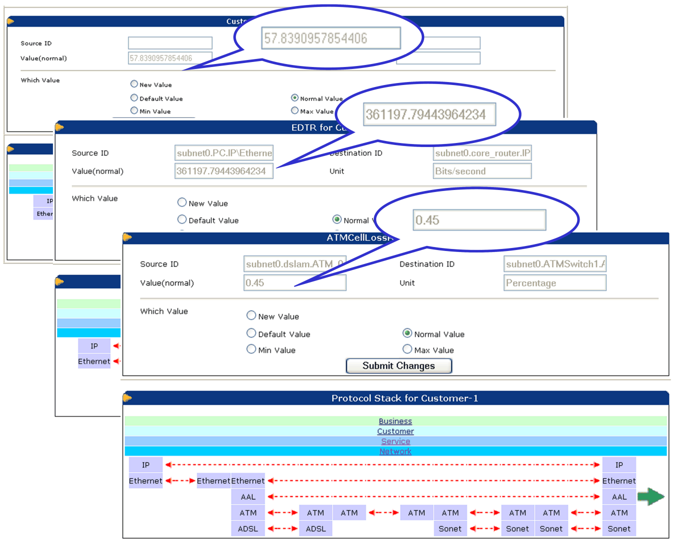

Fig. 11. CyberPlanner Example Customer Troubleshoot

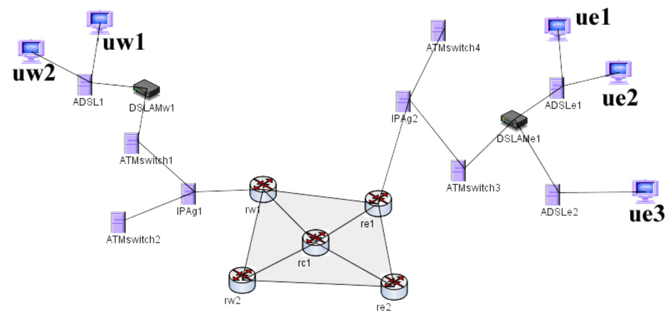

Fig. 12. Example Network Topology

metric are displayed in the right side window, and both of these attributes could be modified at runtime. The metric layers and protocol stacks (technology layer) are displayed at the lower right hand side for easy browsing.

Currently, the prototype implementation is done in Java and sessions are created as Java servlets. Majority of the testing and experimentations are conducted through applet interactions via web browsers. Access to the same MTI among concurrent sessions are not supported right now.

V. Experience with CyberPlanner

Over the years, CyberPlanner has been used in example test cases and industrial experiments, and served as the platform for a number of research ventures. This section presents some of the experiences gained and demonstrates the potentials of CyberPlanner.

One of our use cases dealt with trouble diagnostics using logs. A set of real network performance data was obtained from Bell Canada with example traffic traces from Toronto to Montreal. We have generated a partial regional topology using the topology editor. The original data was presented in plain logs and raw spreadsheet forms which make diagnostics very difficult and there was no correlations of these network data to the service and customer profiles. Using CyberPlanner, we generated a number of MTIs for the customer satisfaction metric (Figure 11), and have found that some of the computed customer satisfaction values are quite low (100 being the maximum rating). Exploring the correlated metrics in the service layer revealed that some of the customer's service sessions experienced low effective data throughput which is the major contributor to the customer's dissatisfaction. Further exploring the network segments at the technology level quickly isolated the problem to a defective ATM switch in the network. Thus, CyberPlanner offers a quick and structured metrics view to the network service providers, while the same diagnostic result would otherwise have to be obtained manually through cross departmental collaborations and is extremely time consuming.

Figure 12 shows an example regional network created based on a real Bell Canada network. It has a general access-transit-core topology, with network technologies similar to those found in Bell Canada networks. We examine the traffic flows of two customers that share much of the same path. Customer 1 uses path uw1-ue1 and customer 2 uses path uw2-ue1. An individual MTI for customer satisfaction is generated for customers 1 and 2 respectively. Figure 13 shows the computation results. Although both MTIs have the same end-to-end effective data throughput (quite low for xDSL service), the computed customer satisfactions differ significantly. This is due to the difference in customer profiles. Customer 1 has less emphasis on the delay aspect of the service, with coefficient 0.1 (i.e. the delay aspect only contributes to 10% of the customer satisfaction over service quality), while customer 2 has a coefficient of 0.9 for the delay aspect and thus is dissatisfied with the performance. Through customer profiling and service performance correlations, CyberPlanner is able to capture the customer aspects in meaningful ways.

Other experiments are conducted to test the metrics correlation models. In one such experiment, the network topology and device technologies are given by Bell Canada for network points between Toronto and London (Ontario, Canada). A full MTI is constructed for end-to-end delay and the metric is then computed based on theoretical values at the device level. The CyberPlanner computed result is then compared with real network measurements obtained off peak hours between these points. Figure 14 shows the absolute error between the measurement data and the theoretical computation. We observe that 80% of the measurement are at most 0.34ms off from the CyberPlanner's computed result. This error also include natural computation error and latency of measuring devices. Interestingly, initial measurements obtained from the

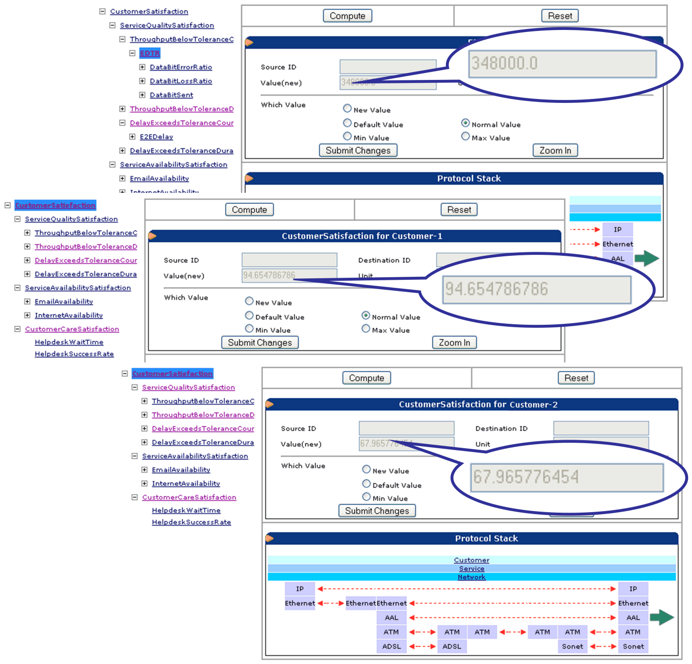

Fig. 13. Example Customer Satisfaction Computation

network showed significant discrepancies with the computed value. After much tracking and testing, it was found that one of the measurement device was not calibrated properly and gave erroneous readings.

A number of research projects leverage or build on the CyberPlanner concept and further demonstrate the potentials of the toolkit. Works on customer-centric network upgrade and planning [1] and business revenue forecast [2] analyze the effect of network upgrades using CyberPlanners and propose new modeling methodologies for metrics at the service, customer and business layers. With the aid of CyberPlanner, extensive simulation and use case studies are performed to demonstrate the effectiveness of such an approach. Zero-budget traffic dimensioning [13] focus on satisfying all the customer demand while maximizing business profit. The work evaluates various traffic engineering plans based on CyberPlanner's forecasting capability. Extensive simulation studies demonstrated that the approach fares significantly better over traditional network dimensioning approaches.

VI. Conclusion

In this paper, we have presented the CyberPlanner concept, architecture and implementation. Central to CyberPlanner are the metrics tree models that correlate and aggregate metrics at all levels of the network service operations. Flexible metrics tree instances that are generated based on a library of MTMs can model the specifics of a provider networks, service attributes, customer profiles, and business objectives. The CyberPlanner toolkit is a powerful tool capable of assisting in many critical network diagnostic, planning, and business

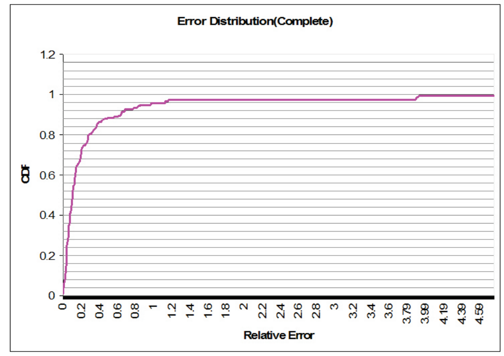

Fig. 14. Comparison of Computed vs. Measured ETE Delay

decision processes. Our experience with CyberPlanner testing and operations have demonstrated its effectiveness and potential.

Acknowledgement

We would like to thank the project participants: Martin Barth, Youssef Iraqi, Tianshu Li, Wenli Liu, Loubna Mekouar, Wayne Szeto, and Dan Xie from University of Waterloo; Michael Desroches and Anwar Haque from Bell Canada; Mijung Choi, James W-K. Hong, and Young J. Won from POSTECH.

References

[1] J. Xiao and R. Boutaba, "Customer-centric network upgrade strategy: Maximizing investment benefits for enhanced service quality," in *IEEE/IFIP Network Operations and Management Symposium (NOMS'04)*, 2004.

[2] ——, "Assessing network service profitability: Modeling from market science perspective," *IEEE/ACM Transactions on Networking*, to appear last quarter 2007.

[3] V. Paxson, G. Almes, J. Mahdavi, and M. Mathis, "Framework for IP performance metrics," May 1998.

[4] J. Mahdavi and V. Paxson, "IPPM metrics for measuring connectivity," September 1999.

[5] G. Almes, S. Kalidindi, and M. Zekauskas, "A one-way delay metric for IPPM," September 1999.

[6] ——, "A one-way packet loss metric for IPPM," September 1999.

[7] S. Bradner, "Benchmarking terminology for network interconnection devices (RFC 1242)," July 1991.

[8] S. Bradner and J. McQuaid, "Benchmarking methodology for network interconnection devices (RFC 2544)," July 1999.

[9] R. Mandeville, "Benchmarking terminology for lan switching devices (RFC 2285)," February 1998.

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

[11] S. Jagannathan, J. Altmann, 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] Y. d'Halluin, P. Forsyth, and K. R. Vetzal, "Managing capacity for telecommunications networks under uncertainty," *IEEE/ACM Transactions on Networking*, vol. Vol.10, no. No.4, August 2002.

[13] W. Liu, Y. Iraqi, and R. Boutaba, "Zero-budget network dimensioning," in *IEEE/IFIP International Symposium on Integrated Network Management (IM'05)*, May 2004.