

Management of Service Level Agreements for Multimedia Internet Service Using a Utility Model

Jong-Tae Park and Jong-Wook Baek, Kyungpook National University

James Won-Ki Hong, POSTECH

ABSTRACT

The efficient management of a quality level of Internet service is becoming increasingly important to both customers and service providers. This article describes how service level agreements for multimedia Internet service can be managed and controlled. We first present a literature survey on the problems of SLA management: SLA parameter definition, SLA measurement, and QoS management. We present a utility model to capture the management and control aspects of SLAs for multimedia Internet service. This utility model has been used in microeconomics theory, but here we have applied it to SLA management. This model provides a computationally feasible solution for admission control and quality adaptation for multimedia Internet service and SLA management. It also allows management policies to be flexibly expressed by service providers. Finally, we apply the utility model to the SLA management of VoIP service and describe how to use it for admission control, dynamic quality adaptation, and resource allocation for SLA assurance.

INTRODUCTION

Unlike a telecommunication network, the Internet is a totally open and distributed environment. It allows users to access all possible services available on the network. Most services are managed by service providers in nonstandard and proprietary ways. The lack of a management framework for Internet service has become the major drawback to the development of further Internet services. Early development of Internet management has focused on network management [1]. However, recently the focus has been on services management and customer management.

The management of Internet service is shifting from a focus on maintaining availability to a focus on maintaining acceptable quality. To meet the quality of service (QoS) requirements,

all components of a service must work correctly. Currently, there are lots of solutions for network management to manage network elements and for systems management to manage hardware and software elements on which a service depends. These solutions are adept at offering views into the network or system, but do not offer a common view of the end-to-end application. Hence, it is very difficult to manage an end-user application as a service. Administrators' needs can no longer be met simply using traditional Simple Network Management Protocol (SNMP)-based troubleshooting tools. Instead, they require service level management and reporting tools for management of Internet service.

There is a trend toward service outsourcing, where providers outsource all or a portion of their service to a third party. In order to guarantee end-to-end QoS to customers, each party must guarantee the availability and performance of the service component it provides. Customers are not concerned about the composition of a service, but only about the QoS. QoS expectations are driving customers to negotiate specific QoS levels with their service providers. This is increasingly done through service level agreements (SLAs) [2]. An SLA is defined as a contract between the service provider and customer that specifies the QoS level that can be expected. It includes the expected behavior of the service and the parameters for QoS. The efficient management of SLAs is a new challenge and very important issue in Internet service management. Research issues on SLA management are briefly summarized:

- SLA parameter definition: There are few common standards for service level parameters: what they are and how their values are computed for SLAs. This concerns the definition of service level parameters such as availability, reliability, latency, and loss for SLA.
- SLA measurement: This issue deals with how to accurately measure the QoS that service providers deliver to their customers.

Valuable network and application performance statistics can serve as the basis for effective SLA management.

- **SLA compliance report:** This deals with mechanisms to satisfy increasingly sophisticated customers who demand real-time reporting to confirm that they are receiving the service levels they were promised.
- **QoS management:** This issue deals with how to manage and control the QoS delivered to customers to ensure compliance with established SLAs. The measurement data plays an important role in this case.

Much research work has been done on SLA measurement [3–5] and SLA compliance reports [6]. However, the problems of SLA parameter definitions and QoS management for efficient management of SLAs have largely been ignored. Extensive research into these problems is needed.

This article describes how SLAs for multimedia Internet service can be managed and controlled. We first present literature surveys about the aforementioned problems of SLA management. We present a utility model to capture the management and control aspects of SLAs for multimedia Internet service. This utility concept has been used in microeconomics theory [7], but here we have applied it to SLA management. This model provides a computationally feasible solution for admission control and quality adaptation for multimedia Internet service for SLA conformance. The utility model also allows management policies to be expressed very flexibly by service providers. The management policies of a service provider can be expressed either indirectly in the service utility function and parameters or as additional system constraints. We apply the utility model to SLA management for voice over IP (VoIP) service and describe how to use it for admission control, dynamic quality adaptation, and resource allocation for SLA assurance. We finally summarize our work on SLA management and discuss future work.

RELATED WORK

SLA PARAMETER DEFINITION

Some research has been done on defining SLA parameters. The QoS team of the TeleManagement Forum has been working on the automation of the interface between service providers and customers for performance reporting with the SLA concept [2]. They have identified common terms and definitions, and have created an industry-wide glossary for performance measurement and reporting. The IP Performance Working Group of the Internet Engineering Task Force (IETF) has been working on the identification of Internet service metrics. These Internet service metrics are:

- Framework for IP Performance Metrics (RFC 2330)
- IPPM Metrics for Measuring Connectivity (RFC 2678)
- A One-Way Delay Metric for IPPM (RFC 2679)
- A One-Way Packet Loss Metric for IPPM (RFC 2680)
- A Round-Trip Delay Metric for IPPM (RFC 2681)

To remain competitive, service providers must offer guarantees not just in terms of availability, but also in terms of performance guarantees such as response time and throughput. There are a number of parameters used in SLAs today. Most are customer support and reliability parameters. Some of the most commonly defined parameters include:

- *Customer support* — These include the typical helpdesk problem of reporting and problem resolution guarantees. Examples include a single point of contact assigned to the customer and problem resolution within 48 hours of reporting.
- *Reliability* — Reliability metrics consist of availability guarantees over a period of time. For example, the Web server will be available 99.999 percent of the time it is accessed over a one-year period.
- *Service provisioning* — This guarantees that the service will be provisioned in a certain manner. For example, a customer will be provided with redundant DS-3 connections to its Web servers.

In addition to reliability and support metrics, service performance metrics are important for business-critical applications. Currently, there are entire new sets of metrics being discussed in the industry, including performance, utilization and security metrics. They include:

- *Performance* — Performance metrics are generally characterized in terms of response time and throughput.
- *Response time* — This metric defines the maximum response time a service is permitted when handling user requests. For example, 95 percent of users will experience a response time of 2 s or less during work hours, where work hours are between 9 a.m. and 5 p.m.
- *Throughput* — This metric defines the rate at which data is delivered to the customer. For example, Intranet users will be able to load a 65 kb GIF file in under 10 s during working hours.
- *Utilization* — This metric defines the maximum service utilization allowed at which a service will perform within guaranteed response times and throughput. An example of this metric is that the system will support 32 simultaneous users during peak hours.

The components of Internet service can be classified into three layers: the application, the system, and the network. A set of SLA parameters at each layer can be included in an SLA. Application level parameters can provide the security, access, configuration, current status, and resource utilization of the Internet service entity and dependent components. System level parameters provide system health information that can affect the performance and reliability of Internet services. Network level parameters are dependent on various transport technologies, such as asynchronous transfer mode (ATM), multiprotocol label switching (MPLS), differentiated services (DiffServ), and so on. In the case of DiffServ networks, an SLA is specified in two parts. One is traffic conditioning specification (TCS), which specifies detailed service parameters for each service level. These parameters

In addition to reliability and support metrics, service performance metrics are important for business-critical applications. Currently, there are entire new sets of metrics being discussed in the industry, including performance, utilization, and security metrics.

Parameter	VoIP	IP VPN
Throughput	Should be based on a five minute collection period at the destination	Should be based on a 5 min collection period at the destination
Utilization	Computed by dividing the measured throughput by the contracted throughput expectation	Computed by dividing the measured throughput by the contracted throughput expectation
Packet loss ratio (PLR)	Less than 0.1% for a 5 min sample	<ul style="list-style-type: none"> For assured forwarding (AF) class of service, the PLR should be $\leq 0.1\%$ for 95% of the collected sample For expedited forwarding (EF) class of service, the PLR should be $\leq 2.0\%$ for 95% of the collected sample
Packet delay (PD)	Less than 150 ms for a 5 min sample. This implies an upper limit of 70 ms for IP packet delay, since the codec to audio and audio to codec delay are approximately 40 ms each	<ul style="list-style-type: none"> For AF class of service, the PD should be ≤ 50 ms for 95% of the collected sample For EF class of service, the PD should be ≤ 25 ms for 95% of the collected sample
Availability	$\geq 99.9\%$	$\geq 99.9\%$
Packet jitter	$\leq \pm 45$ ms	NA

■ Table 1. Examples of SLA parameters.

include detailed service performance parameters and traffic profiles such as token bucket parameters. Other parameters are general service characteristics, such as availability, reliability, encryption service, and pricing and billing mechanisms. SLA contracts can also specify sharing of excess bandwidth in an overprovisioned network, as well as state how bandwidth would be distributed in an under-provisioned network. Some examples of SLA parameters are described in Table 1.

SLA MEASUREMENT

The measurement mechanisms of SLA are classified into passive probe, active probe, polling MIB and classical tools. The characteristics of these mechanisms are summarized in Table 2.

There are numerous research projects related to the measurement of SLA [3–5]. These are summarized in Table 3.

SLA/QoS MANAGEMENT

Bhoj [6] presented a Web-based SLA management framework to allow easy interdomain communication. She demonstrated how service

providers could offer SLA monitoring capabilities to their customers for a number of services, including e-mail and network access services. Park's research [8] supported QoS management using the SLA concept, which was measured, monitored, and controlled systematically in a multomain environment.

The work by the Integrated Services Working Group of IETF has made a significant contribution to providing controlled QoS for multimedia applications over the Internet. The group has defined a comprehensive integrated service architecture [9] and QoS framework to specify the functionality of the Internet system elements which could make multiple dynamically selectable QoS available to applications.

In addition to this work, several management information bases (MIBs) have appeared which are related to the performance management of Internet services. These include the Host Resources MIB, Network Services Monitoring MIB, Application MIB, and SLA Performance Monitoring MIB. The Host Resources MIB defines a uniform set of

Method	Advantages	Constraints	Comments
Passive probe	<ul style="list-style-type: none"> Independent of NE MIBs Collects a lot of statistics No additional traffic 	<ul style="list-style-type: none"> Instrumentation required Cost concerns Needs synchronization to measure delay 	<ul style="list-style-type: none"> Needs to be strategically placed in the network (e.g., OC-nMON, RMON, RMON2)
Active probe	<ul style="list-style-type: none"> Independent of NE MIBs Good for measuring delays 	<ul style="list-style-type: none"> Instrumentation required Cost concerns Generates additional traffic Test packets processed differently Intrusive 	<ul style="list-style-type: none"> Needs to be strategically placed in the network (e.g., Cisco's RTR)
Polling MIB	<ul style="list-style-type: none"> Scalable, simple 	<ul style="list-style-type: none"> Some required data not always available Not good for measuring delay 	<ul style="list-style-type: none"> Polling MIBs necessary for network performance as well
Classical tools	<ul style="list-style-type: none"> Simple 	<ul style="list-style-type: none"> Only static information Traceroute Netstat 	<ul style="list-style-type: none"> Ping

■ Table 2. SLA measurement mechanisms.

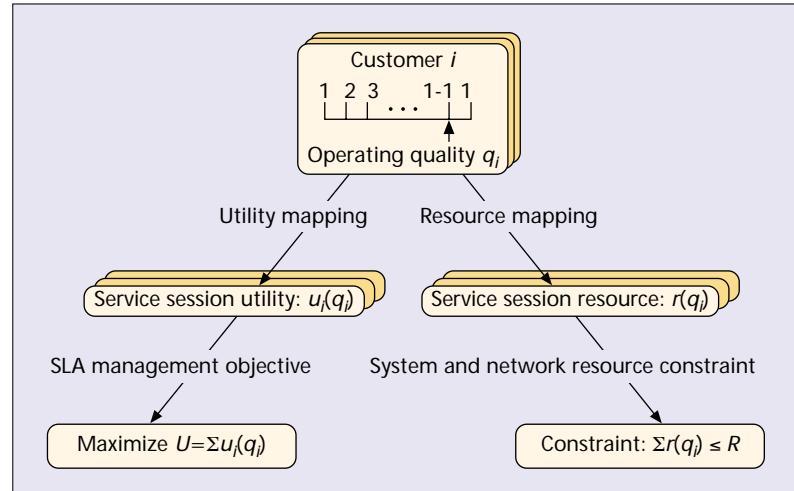
objects useful for the management of host computers independent of the operating system, network services, or any software application. The Application MIB has been defined to represent installed and running applications and their components. The SLA Performance Monitoring MIB defines a set of objects for monitoring SLAs.

A UTILITY MODEL FOR SLA MANAGEMENT

While the related work presented earlier offer a good start for SLA management, unsolved problems still remain. That work only focused on the monitoring and reporting mechanisms of SLAs, and do not adequately address how to manage and control the QoS level provided to customers while efficiently utilizing the system and network resources of a service provider. None of the research presents a formalized solution to the problem. To solve the problem, we propose a utility model to capture the management and control aspects of SLAs for multimedia Internet service.

In this section we present a utility model which is a mathematical model designed to capture the management and control aspects of SLAs for multimedia Internet service. This utility concept has been used in microeconomic theory [7], but here we apply it to SLA management. In microeconomics, utility is defined as the satisfaction of wants and needs obtained from the use or consumption of goods and services. In SLA management, utility is defined as the satisfaction obtained from a service provider of the consumed system and network resources.

The utility model formulates the adaptive SLA management problem as integer programming. This model provides a unified and computationally feasible approach to make session admission control, quality adaptation, and resource allocation decisions of an SLA



■ Figure 1. The main concept of the utility model.

management system for multimedia Internet service. The utility model is based on the concepts of quality profile, quality-to-resource mapping, resource constraints, and utility function. The main concepts of the utility model for SLA management are illustrated in Fig. 1.

The quality profile specifies the quality preferences of customers. It is a set of acceptable operating qualities for the service and is specified via a contract between the customer and service provider. We assume the existence of a mapping from an operating quality to the resources required to provide that quality. For each resource related to the service, the sum of the quantities of the resource allocated to all the customers cannot exceed the total available quantities of the resource.

The utility function is one that maps a customer's operating quality to the amount of utility. For example, while contracting an SLA, the customer and service provider specify the quality

Project	Description	Organization
Skitter	Measures forward IP paths (each hop) from a source to many destinations	DARPA/NSF/CAIDA
AMP	Intended to improve the understanding of how high performance networks perform and to help in problem diagnosis for both the network's users and its providers	NSF/NLANR/Internet 2
NIMI	Measures the global Internet; NIMI probes can be delegated to administration managers for configuration information and measurement coordination	DARPA/NSF
PingER	Provides active monitoring of end-to-end performance of Internet links	DOE/ESnet/HENP/XIWT
Surveyor	Provides active tests of one-way delay and packet loss along paths between measurement machines	CSG/Advanced
RIPE Test Traffic	Provides independent measurements of connectivity parameters such as delays and routing-vectors	RIPE
IPMA	Studies the performance of networks and networking protocols in local and wide area networks; also collects routing and network statistics	Merit

■ Table 3. Research projects related to SLA measurement.

The heuristic solution starts with an operating quality that has the smallest utility in each customer and iteratively improves the solution by gradually replacing it with an operating quality that has a larger utility as long as the solution is feasible.

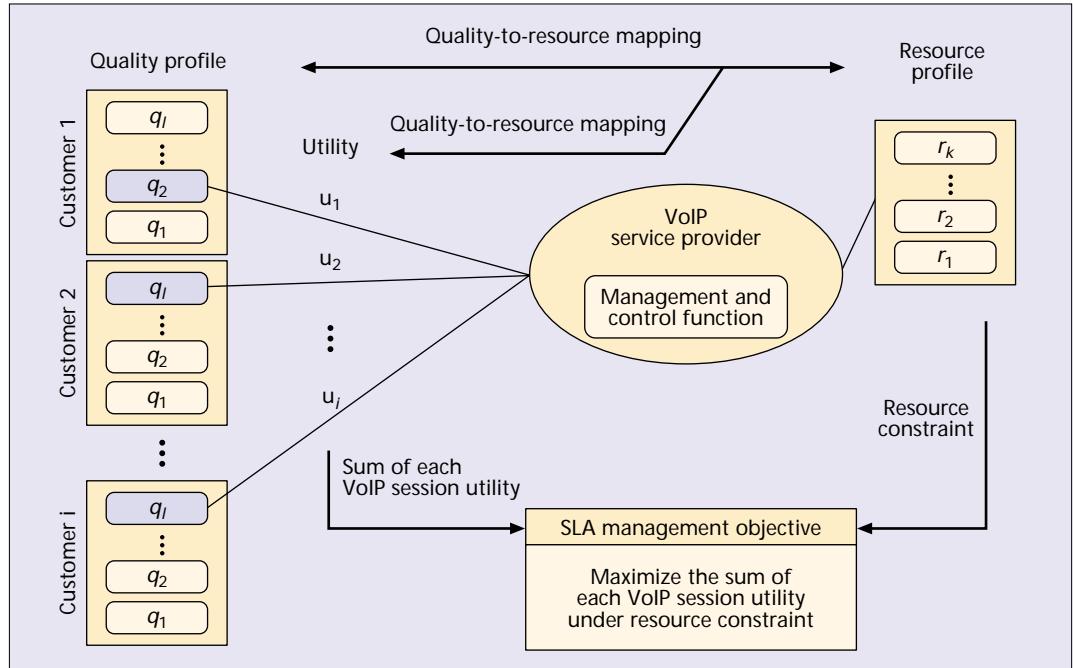


Figure 2. The utility model for VoIP service.

profile and agree on a dollar value paid for each operating quality. The service provider can map this value to the utility value. This utility function may be a simple table, a linear function, or a more complex function. The extraction of a utility function is a nontrivial problem. These complex questions may be highly dependent on the application and the service provision environment. Generally, the utility function can be empirically estimated and established. An intelligent utility agent may be used dynamically to derive such a mapping.

Suppose that there exist several quality levels which can be assigned to customers by a service provider. The total amount of available resources, denoted R , is limited. In this case, the objective of SLA management for a multimedia service provider is to maximize the service utility objective function U such that

$$U = \sum_i \sum_j x_{ij} u_{ij}, \quad (1)$$

where

$$\sum_i \sum_j x_{ij} r_{ij} \leq R, \quad \sum_j x_{ij} = 1 \text{ and } x_{ij} \in \{0,1\}.$$

Equation 1 means that the problem of SLA management of a multimedia service provider is to find the operating quality of each customer, while maximizing the service utility under the resource constraint. Here, only one operating quality must be chosen among the acceptable operating qualities of each customer.

Since the problem is known to be NP-hard [10], the computation time exponentially increases in the worst case. This is not suitable for time-critical control and management, such as dynamic resource allocation and admission control for multimedia Internet services. To cope with time-critical management and control, we developed a heuristic solution to the problem, which is shown below.

PROCEDURE HEURISTIC()

BEGIN

- 1 Initialize a solution with an operating quality that has the smallest utility in each customer
- 2 loop
 - 3 for customer $\leftarrow 1$ to n do
 - 4 for operating quality \leftarrow initial solution + 1 to the number of the quality level do
 - 5 if the current solution exceeds the available resources then continue
 - 6 if calculate Δr and it is greater than zero
 - 7 then determine whether it is the maximum resource savings
 - 8 else calculate Δp and determine whether it is the maximum utility gain
 - 9 repeat
 - 10 repeat
 - 11 if there is no feasible solution then return the solution
 - 12 Update the solution with the maximum Δr or Δp
 - 13 repeat

END HEURISTIC

The heuristic solution starts with an operating quality that has the smallest utility in each customer and iteratively improves the solution by gradually replacing it with an operating quality that has a larger utility as long as the solution is feasible. The decision criterion of upgrading an operating quality is to maximize the extra resource savings (Δr) and utility gain (Δu) per unit of the extra resource savings (Δp). The main loop of the procedure HEURISTIC attempts to find a feasible upgrade (lines 3–10). If an upgrade is found, the solution is updated and the heuristic starts another iteration. The process of locating a feasible upgrade involves the following steps:

- Find the extra aggregate resources (Δr) for all feasible upgrades.
- If there exists at least one feasible upgrade which provides savings in aggregate

resources, procedure *HEURISTIC* chooses the upgrade that maximizes the savings in aggregate resources.

- However, if a feasible upgrade which provides savings in aggregate resources does not exist, procedure *HEURISTIC* chooses the upgrade that maximizes the utility gain per unit of extra aggregate resources (i.e., $\Delta p = \Delta u/\Delta r$).
- If the heuristic fails to find a feasible upgrade in an iteration, it returns the current solution and terminates.

AN SLA MANAGEMENT EXAMPLE USING THE UTILITY MODEL

In this section we present how SLAs for VoIP service can be managed using the proposed utility model presented earlier.

A UTILITY MODEL FOR VOIP SERVICE

Figure 2 shows how the proposed utility model is applied to SLA management for VoIP service.

A VoIP service provider makes contracts with customers and makes the quality profile of each customer. The quality profile is a sequence of acceptable operating qualities from the lowest acceptable to the highest acceptable operating quality. The service provider maps each operating quality level to the appropriate resource profile. In addition, it also maps it to a utility value. The utility value of a quality level can be obtained using a utility function, which is determined by the service provider's operating policy. A utility value is a real number and represents the amount of satisfaction the service provider obtains from a VoIP service. The details of the procedure are explained as follows:

- **Getting the quality profile:** A VoIP service provider must be able to specify a quality profile that expresses QoS requirements. This can be achieved using a static table of acceptable qualities. For instance, a simple quality profile for a VoIP session may have three discrete qualities: bronze, silver, and gold. The VoIP session's minimum acceptable quality is bronze, and its maximum desired quality is gold.
- **Quality-to-resource mapping:** The utility model assumes the existence of an operating-quality-to-required-resource mapping. However, the extent of mapping to be obtained is another research issue. Resource allocation can be done by profit maximization, fair share policy, or priority policy. Resource allocation may be obtained using offline experimental evaluation, but is dependent on the service provider's platform.
- **Quality-to-utility mapping:** If the utility of each VoIP session represents a customer's bill, the quality of the VoIP session can be mapped to a utility value using the utility function, $u(x) = 1 - e^{-x/\tau}$, where x is the price of an operating quality and τ is a constant. The above function is an increasing concave function of the price x . This is a standard assumption, indicating diminishing marginal utility as the price x increases. In this case, the sum of each VoIP session utility means

the service provider's profit. The management system uses these values for admission control and runtime quality adaptation.

A QOS MANAGEMENT FUNCTION FOR SLA CONFORMANCE

Here we describe the admission control and QoS adaptation functions for SLA conformance to VoIP service.

• Admission control of new VoIP sessions:

Admission control is necessary for service with quality guarantees because the system must ensure that sufficient resources are available at runtime to meet the minimum quality guarantee. Suppose the VoIP service provider currently has n sessions and the current total utility is U_n . When a customer requests a new session, the utility-driven admission control can be processed as follows:

Step 1: The management system checks for a feasible solution of Eq. 1, where the $n + 1$ VoIP sessions can share the currently available resources. If such a solution does not exist, the new session must be rejected.

Step 2: If a feasible solution with $n + 1$ sessions exists, suppose the maximum service utility of $n + 1$ VoIP sessions is U_{n+1} . If $U_{n+1} < U_n$ (i.e., no increase in service utility), the new session should be rejected as unprofitable; otherwise, the session should be accepted.

• QoS Adaptation of VoIP Sessions: Customers are more satisfied as voice quality improves. However, a service provider should guarantee a minimum service quality for all customers, although the service provider's system and network conditions are changing. In other words, the service provider must be able to dynamically adapt the operating quality of each VoIP session when the quality is degraded or the network resource status is changed.

Figure 3 describes a QoS adaptation function based on the proposed utility model. It is composed of a monitoring function, an assessment function, and a control function. The monitoring function plays the role of monitoring the performance of VoIP sessions and the network resource status. The assessment function decides whether a QoS violation occurs or QoS restoration is required. If required, the control function determines the new operating quality of each VoIP session and reallocates the resources to guarantee the QoS of the VoIP session within SLA.

The QoS adaptation function in Fig. 3 can be processed as follows:

Step 1: Initialization

- a) Get quality profile and resource profile
- b) Get utility function
- c) Determine operating qualities for each customer

- d) Allocate resources as a result of Step 1(c)

Step 2: Performance monitoring

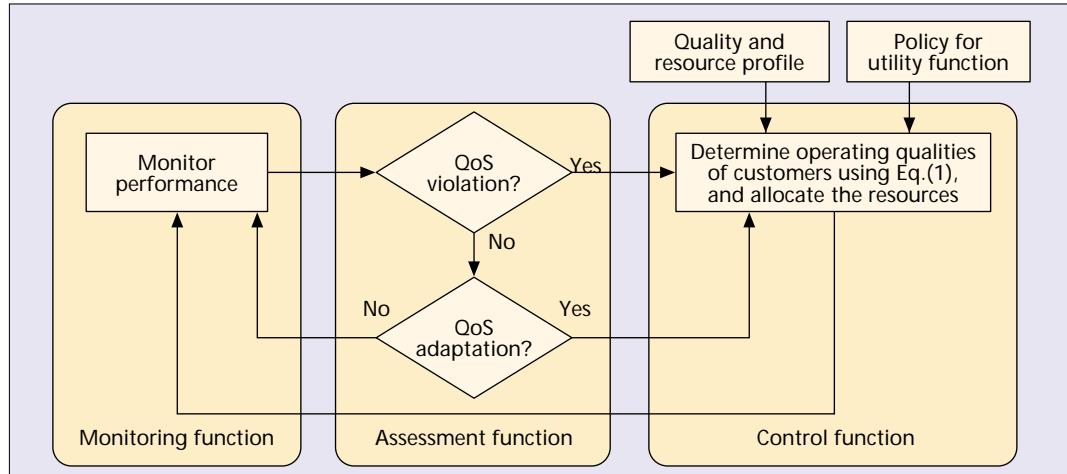
Step 3: Current state assessment

- a) If QoS violation occurs, then go to Step 4
- b) If QoS Adaptation is required, then go to Step 4
- c) Go to Step 2

The quality profile is a sequence of acceptable operating qualities from the lowest acceptable operating quality to the highest acceptable operating quality.

The service provider maps each operating quality level to the appropriate resource profile.

Further research is required for mapping operating qualities to required resources. The utility model assumes that the mapping of an operating-quality to required-resource mapping is available *a priori*.



■ Figure 3. QoS adaptation functional description.

Step 4: QoS adaptation

- Determine new operating qualities for each customer
- Reallocate resources as a result of Step 4a
- Go to Step 2

CONCLUSION

In this article we present a high-level conceptual SLA management framework for multimedia Internet service using a utility model. The various measurement mechanisms of SLA parameters are comparatively reviewed, and other research on SLA/QoS management problems is also investigated. Most previous work focused only on the monitoring and reporting mechanisms of SLA, and do not adequately address how to manage and control the quality level of service provided to customers. None of the research presents a formalized solution to the problem. We present a utility model to solve the SLA management problem. The utility model can be used not only for resource allocation decisions, but also for quality adaptation and admission control for multimedia Internet service. We have also presented an example of SLA management for VoIP service which applies the utility model.

Further research is required for mapping operating qualities to required resources. The utility model assumes that the mapping of an operating quality to required resource is available *a priori*.

REFERENCES

- [1] F. Barillaud, L. Deri, and M. Feridun, "Network Management Using Internet Technologies," *Proc. IEEE/IFIP Int'l Symp. Integrated Network Mgmt.*, San Diego CA, May 1997, pp. 61-70.
- [2] NMF, "Performance Reporting Definitions Document," NMF 701, June 1998.
- [3] S. Kalidindi and M. J. Zekauskas, "Surveyor: An Infrastructure for Internet Performance Measurements," *Proc. 9th INET Conf.*, San Jose, CA, June 1999.
- [4] H. Uijterwaal and O. Kolkman, "Internet Delay Measurements Using Test Traffic," RIPE NCC, 1997.
- [5] G. R. Malan and F. Jahanian, "An Extensible Probe Architecture for Network Protocol Performance Measurement," *Proc. ACM SIGCOMM '98*, Sept. 1998,
- [6] P. Bhoj, S. Singhal, and S. Chutani, "SLA Management in Federated Environments," *Comp. Networks*, vol. 35, no. 1, Jan. 2001.
- [7] A. Mas-Colell, *Microeconomics*, Oxford Univ. Press, 1995.

- [8] J. T. Park and J. W. Baek, "Web-Based Intranet/Internet Service Management with QoS Support," *IEICE Trans. Commun.*, vol. E82-B, no. 11, Nov. 1999, pp. 1808-16.
- [9] R. Braden, D. Clark, and S. Shenker, "Integrated Services in the Internet Architecture: An Overview," RFC 1633, 1994.
- [10] E. Horowitz and S. Sahni, *Fundamentals of Computer Algorithms*, Comp. Sci. Press, 1989.

BIOGRAPHIES

JONG-TAE PARK [SM] (park@ee.knu.ac.kr) received B.S. and M.S. degrees from Kyungpook National University and Seoul National University, respectively, in Korea, and a Ph.D. degree in computer science and engineering from the University of Michigan, in 1987. After graduation, he was at AT&T Bell Labs, working on network management and service provisioning. Since 1989 he has been teaching at the School of Electrical Engineering and Computer Science at Kyungpook National University, Korea, where he is now a professor. He is also chair of the Information and Communication Department and director of the Center for Next Generation Information Technology. He serves as chair of the Technical Committee of Information Infrastructure of IEEE Communication Society. He was a general chair for APNOMS '97. He serves as editor for several journals, including *International Journal of Communications and Networks* and *Journal of Network and Systems Management*, and has published more than 100 journals and articles in the areas of computer communication networks, network management, databases, and distributed systems. His current research interests include IP-based multimedia service management, Internet and telephony integration, and wireless network control and management. He is a member of KICS, KISS, and KNOM.

JONG-WOOK BAEK (jwback@ain.knu.ac.kr) received B.Sc., M.Sc., and Ph.D. degrees in electronic engineering from Kyungpook National University in 1995, 1997, and 2001, respectively. As of 2001, he is teaching at the School of Internet Engineering at Dongseo University, Korea, where he is an assistant professor. His current research interests include the Internet and its service control and management, SLA management, and telecommunication network management.

JAMES WON-KI HONG [M] (jwkong@postech.ac.kr) is an associate professor in the Department of Computer Science and Engineering, POSTECH, Pohang, Korea. He has been with POSTECH since May 1995. Prior to that he was a research professor in the Department of Computer Science, University of Western Ontario, Canada. He received B.Sc. and M.Sc. degrees from the University of Western Ontario in 1983 and 1985, respectively, and a Ph.D. degree from the University of Waterloo, Canada in 1991. He has served as a Technical Chair for IEEE CNOM from 1998 to 2000. He was a Technical Co-Chair for 2000 IEEE/IFIP Network Operations and Management Symposium (NOMS 2000). He has been a standing committee member for AsiaPacific Network Operations and Management Symposium(APNOMS) since 1997. His research interests include network and systems management, distributed computing, and multimedia systems. He is a member of KICS, KNOM, and KISS.