Admission Control for Providing QoS in DiffServ IP Networks: The TEQUILA Approach

Eleni Mykoniati and Charalampos Charalampous, National Technical University of Athens
Panos Georgatsos and Takis Damilatis, Algonet S.A.
Danny Goderis, Alcatel Network Strategy Group
Panos Trimintzios and George Pavlou, University of Surrey
David Griffin, University College London

ABSTRACT

Admission control is a key component for QoS delivery in IP networks because it determines the extent to which network resources are utilized and whether the contracted QoS characteristics are actually delivered. Unlike most of the existing approaches, we adopt a holistic view and position the admission control in a unified architecture, where traffic engineering and service management interact in different levels of abstraction and timescale. Distinguishing between service subscription and invocation epochs, the applied admission logic is based on a feedback model, the operation of which can be tuned by operational policies and business strategies. The validity of the approach is supported by a scalability analysis while it has also been verified by a prototype system.

INTRODUCTION

The enlargement of the Internet user community has generated the need for IP-based applications requiring guaranteed quality of service (QoS) characteristics. The integrated services (IntServ) and differentiated services (DiffServ) frameworks have been proposed to address QoS. While IntServ operates on a per-flow basis and hence provides a strong service model that enables strong per-flow QoS guarantees, it suffers from scalability problems. On the other hand, DiffServ [1] keeps per-flow information only at the edge of a domain and aggregates flows into a limited set of traffic classes within the network, resolving the scalability problem at the expense of looser QoS guarantees.

Beyond the standardized functionality at the IP layer, a large body of work has been devoted to architectures and functions necessary to deliver end-to-end QoS. These functions can be categorized into traffic engineering (TE) functions and service management (SrvMgt) functions [2]. TE functions are mainly concerned with the management of network resources with the purpose of accommodating offered traffic in an optimal fashion. SrvMgt functions deal with the handling of customer service requests, trying to maximize incoming traffic, in terms of number of contracts and throughput, while respecting the service provider's (SP's) commitments on the agreed QoS guarantees. SrvMgt mechanisms for service offering, agreement, and activation need to be in place. In addition, in order to guarantee the agreed QoS requirements, SrvMgt needs to avoid overloading the network beyond loads it can gracefully sustain. SrvMgt functions that deal with the latter task are referred to as *admission control* and are the key focus of this article.

In [3] a bandwidth broker (BB) architecture implies that admission control decisions are made at a central location for each administrative domain. Although the cost of handling service requests is significantly reduced, it is unlikely that the approach scales for large networks. In order to cope with scalability most relevant studies adopt distributed admission control schemes, which are further distinguished into model-based and measurement-based approaches. Both approaches assess QoS deterioration probability upon service request arrivals; model-based approaches [4] maintain state information for active services and employ mathematical models, whereas measurement-based approaches rely on either passive [5] or active [6] aggregate measurements.

This article presents the admission control approach of the TEQUILA[1] (Traffic Engineering for Quality of Service in the Internet at Large Scale) project. TEQUILA designed and implemented a complete solution for QoS deliv-

---

[1] *See: www.ist-tequila.org*

0163-6804/03/$17.00 © 2003 IEEE

ery in IP DiffServ networks. A unified architecture has been proposed [2], combining SrvMgt and TE functions at various levels of timescale and abstraction. Proposed algorithms have been validated through simulation, while a prototype system has also been implemented and related tests have been conducted to assess the validity of the approach.

Considering that providing hard QoS guarantees is prohibitively costly and that probabilistic QoS guarantees cannot be safely given at service request times, the proposed admission control scheme adopts a feedback-based model for asserting the risk of QoS deterioration as a function of admitted traffic. To increase efficiency, admission control logic is applied at both service subscription and invocation epochs. Feedback information is used at different levels of abstraction: input from offline TE functions on the ability of the engineered network to deliver QoS given subscription admission control decisions, and input from measurements on the actual status of the network given invocation admission control decisions. Furthermore, our approach allows for policy-driven operation based on a best-practice paradigm.

The rest of this article is organized as follows. We position the proposed admission control scheme by placing it into the context of a general service management framework. We lay down the basic dimensions of our admission control approach. We also outline the admission logic at the subscription and invocation epochs. We discuss the scalability of the proposed scheme, and finally we present the summary and conclusions.

SERVICE MANAGEMENT FRAMEWORK

Admission control is an integral part of an IP QoS delivery solution. Hence:

- It is dependent on the informational model used to describe QoS-based services. From the perspective of admission control, service models should include specification of the QoS parameters as well as user traffic conformance for receiving the specified QoS.
- Its operation is bound to the QoS capabilities available in the network.

Furthermore, admission control should be policy-driven in order to adjust to the specific service provisioning strategies of the particular SP.

The above issues should be tackled in the context of a general SrvMgt framework, since they are strongly coupled with the environment where admission control is to operate. In our approach, we adopt the TEQUILA SrvMgt framework presented below.

A HIERARCHICAL SERVICE MODEL

A hierarchical service model is adopted (Fig. 1), which spans from service level agreements (SLAs) to per hop behaviors (PHBs), the basic QoS building blocks in IP DiffServ networks [7]. SLAs describe all aspects of a service contract. The technical aspects of a service contract are described by the so-called service level specifications (SLSs). For QoS-based IP connectivity services SLSs are modeled on the basis of standard templates [8]. The TEQUILA service hierarchy introduces the notion of *QoS classes* to link SLSs with PHBs. QoS classes depict the elementary QoS transfer capabilities of an SP domain, consisting of an ordered aggregate (OA) and associated QoS parameters such as one-way delay and packet loss. Each service corresponds to a number of SLSs, and each SLS corresponds to a number of QoS classes. Therefore, given a service, its QoS is completely defined through the QoS classes of its constituent SLSs.

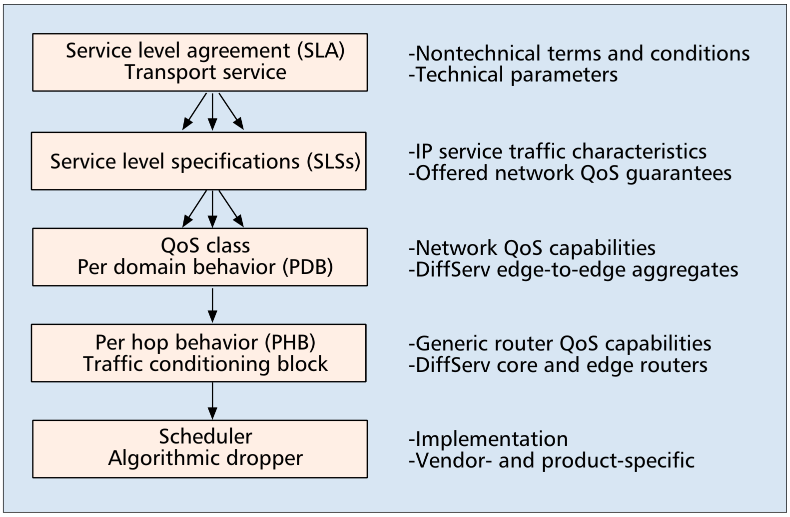

Figure 1. *The hierarchical service model.*

SERVICE SUBSCRIPTION AND INVOCATION

We distinguish between service *subscriptions* and *invocations* [2, 7]. This distinction is mainly necessitated by authentication, authorization, and accounting (AAA) concerns and is in line with current business practices. The term invocation implies both implicit and explicit invocations of admitted subscriptions. In implicit service invocations users may directly inject traffic into the network (e.g., for virtual private network, VPN, services), whereas in explicit service invocations a signaling protocol interaction is required with the SP domain before traffic is actually injected into the network.

INTERACTION WITH TRAFFIC ENGINEERING

SrvMgt and TE functions do not act in isolation (Fig. 2). TE functions provide the grounds on which the SrvMgt functions operate, while SrvMgt functions set the traffic-related objectives for the TE functions to fulfill. Specifically, SrvMgt establishes subscriptions based on which a Traffic Forecast function produces the *traffic matrix* (TM), which specifies anticipated QoS traffic demand between network edges. Traffic demand is forecast from historical data and/or SP expectations (e.g., sales targets). Based on the forecasted traffic demand, the network is appropriately dimensioned by the TE functions, in terms of PHB configuration parameters and QoS route constraints. In turn, TE functions produce the *resource availability matrix (RAM)*, which specifies estimates of the availability of the engineered network to accommodate QoS traffic between network edges. Based on the availability estimates, SrvMgt functions, also uti-

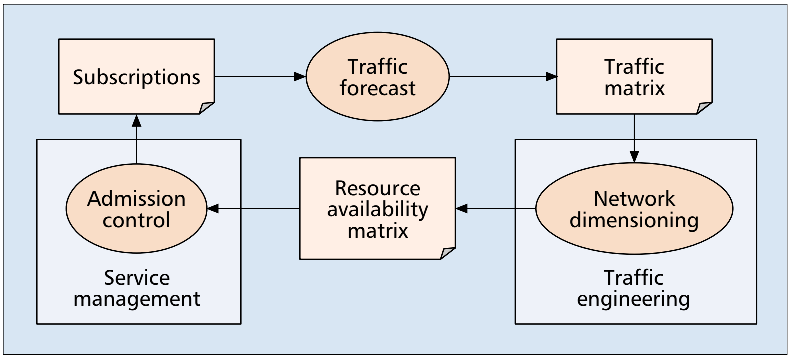

Figure 2. *The resource provisioning cycle.*

lizing network state information, handle the admission of service requests so as not to overload the network.

The above interactions occur in resource provisioning cycle (RPC) epochs [7]. Should anticipated traffic demand significantly change, a new TM is produced, the network is appropriately redimensioned, the RAM is in turn produced, and a new RPC starts.

ADMISSION CONTROL PRINCIPLES AND NOTIONS

TWO-LEVEL ADMISSION CONTROL

The effectiveness of admission control largely depends on the offered load, which in turn depends on the subscribed population. Therefore, applying logic at subscription epochs (i.e., regulating the subscribed population) increases the level of control and flexibility in resolving the traditional admission control trade-off between network utilization and QoS deterioration. In our design, subscription and invocation logic are not bound together, but operate autonomously and can be suppressed if required. Invocation logic is distributed, whereas subscription logic is centralized as its dynamics, in terms of frequency of requests and required network information, is regarded less stringent than the dynamics of invocation logic.

SERVICE SATISFACTION — QOS GUARANTEES

Inherent to any admission control scheme, as well as central to the design of TE-based QoS delivery solutions, is the nature of guarantees to be provided for the agreed QoS.

The provisioning of hard QoS guarantees implies the adoption of worst-case models for characterizing the behavior of traffic sources and the aggregation of traffic streams, and peak allocation-based reservation schemes for the allocation of network resources. However, for admission control, even peak allocation schemes cannot provide hard guarantees, unless complete partitioning of network resources in the core and hard edge-to-edge reservations are applied, which are regarded to be prohibitively costly, especially for bursty traffic. Moreover, statistical/stochastic models cannot provide hard QoS guarantees, because of the assumptions and approximations pertinent in modeling and aggregating the largely unpredictable user traffic.

It is our view that:

- QoS guarantees can only be given within statistical bounds. For this, the adopted TEQUILA SLS template [8] specifies QoS performance targets not as fixed bounds but as bounds over time intervals, or percentiles.
- The statistical error in providing QoS guarantees cannot be determined accurately a priori. Reliably assessing the deviation from the hard QoS guarantees entails the same statistical errors as ensuring hard guarantees. Deviations from agreed QoS can only be assessed *a posteriori* based on measurements.

Rather than developing yet another model to a priori assess the eventual QoS deterioration and determine its acceptable levels, our admission control approach opts for mechanisms ensuring that QoS is kept at acceptable levels as set by policies, independent of the inherent statistical errors of the underlying models.

It is should be noted that some applications may tolerate throughout their lifetime less stringent QoS than the required targets specified in the SLSs [9]. The degree to which occasional QoS deterioration may be acceptable varies depending on how adaptive the application is. In line with this view, it is the distinction between the expedited forwarding and assured forwarding PHB groups. We also take the view that acceptable QoS levels are subject to policies for service provisioning (e.g., favoring specific customer types and/or types of services).

The notion of *service satisfaction* is introduced to enable reasoning on acceptable QoS. The *satisfaction* of a particular service is defined as the percentage of the traffic admitted into the network over the contractual service rate, given that the QoS targets are met. Furthermore, to capture the different degrees of tolerance in QoS deterioration, we introduce the notions of *almost satisfied* and *fully satisfied* services. Services are deemed *fully satisfied* when they enjoy QoS at their contractual rate and *almost satisfied* when they enjoy QoS at rates lower than their contractual rates but above what is considered acceptable for the particular type of service. The rate-orientated definition is in line with [10], arguing that, in times of congestion, due to TCP retransmissions, it is good practice to reduce the traffic admitted to the network to keep the loss ratio low.

In trying to achieve the best of service satisfaction, our admission logic is parameterized in terms of appropriate operational parameters. Because QoS guarantees cannot be a priori assessed, we opt for a relative setting of these parameters, following a “tuning” paradigm, where operational policies set these parameters on the basis of best practice performance (Fig. 3).

ADMISSION MODEL

The majority of the admission control approaches utilize information on preallocated network capacity. Based on this information, their model- or measurement-based logic assesses the current availability of the network to either admit or reject incoming service requests. Alternatively, probe-based schemes [6] deduce the ability of the network to sustain the offered load directly, without relying on preallocated network capacity information. However, the latter introduce significant latency in response times, and have inherent problems caused by probes stealing bandwidth from established flows and denial of service when simultaneous attempts congest the network and none is accepted although resources are available.

Our admission control scheme uses the resource availability estimates calculated by the offline TE functions in the RAM per *traffic trunk* (TT). A TT is a QoS class in a certain topological scope. That is, TTs are aggregates of traffic having the transfer characteristics of the associated QoS class between specific network edges. Resource availability estimates are not hard-reserved in the network. Furthermore they are shared among QoS classes in order to compensate for the inevitable traffic fluctuations. Hence, admission control utilizes resource availability estimates as guidelines rather than as stringent directives.

Such guidelines are interpreted through the abstraction of a resource availability buffer, deduced from RAM, that both subscription and invocation admission control may utilize, each from its own perspective. The risk of overloading the network is deduced on the basis of admissible/noncritical area in that buffer, which is determined by interpreting appropriate operational parameters set by policies.

It is worth mentioning that although both functions make use of the same buffer, there is no correlation of their decisions since they affect different aspects of the offered load to be finally injected into the network and are applied at different timescales.

SUBSCRIPTION LOGIC

OBJECTIVES

The objective of the subscription logic is to determine whether service subscription requests can be accepted, with the purpose of not eventually overloading the network, while maximizing subscribed traffic.

RATIONALE

Topological information is considered to be uavailable; therefore, the overlapping of TTs in terms of competition for the same network resources is not taken into account.

In line with the principles presented in the previous section, we introduce the *satisfaction level* policy parameter, denoting the degree of service satisfaction, as defined earlier, the SP would opt to provide to its subscriptions. Its values are relatively set in the range from –1 to 1; 1 corresponding to the *fully satisfied*, 0 to the *almost satisfied*, and –1 to less than the almost satisfied level of service satisfaction.

OPERATION

Given a subscription request, its implied demand is derived and aggregated with the demand of already established subscriptions. Total anticipated demand is estimated in terms of a minimum and maximum value corresponding to the almost and fully satisfied service rates of the subscribed services, respectively.

The satisfaction level determines the portion of minimum and maximum demands to be *safely accommodated* in the network. Positive values imply that, in addition to the minimum demand, a portion of maximum demand needs to be safely accommodated. Negative values imply that only a portion of minimum demand needs to be safely accommodated.

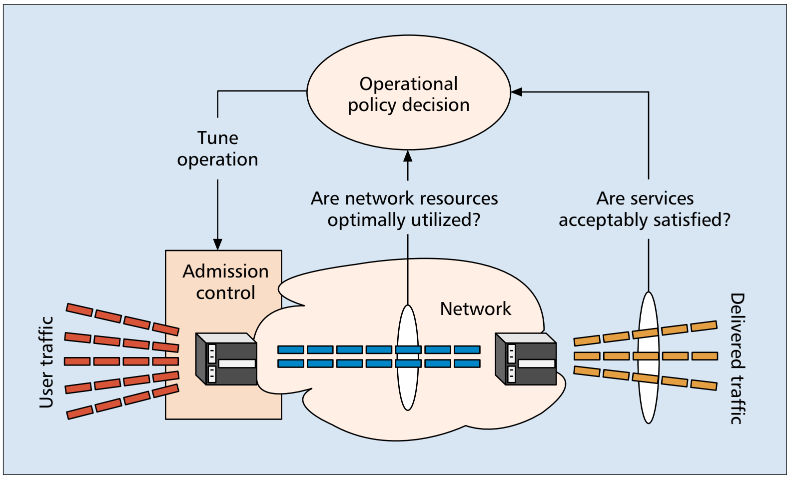

Figure 3. *Best-practice operational policies setting.*

Safe accommodation is deduced by checking whether the determined portions of minimum and maximum demand can fit in the area of the resource availability buffer, dedicated to this traffic even at congestion times. If they do fit, the subscription request is accepted; otherwise, the subscription is either rejected or alternatives of lower QoS and/or lower service rates are offered.

It is evident from the above that our subscription logic can accept subscriptions either with no traffic considerations at all, on the grounds of their expected demand, or on a worst-case reasoning, depending on the particular setting of the satisfaction level. This is left to the policy of the SP domain.

INVOCATION LOGIC

OBJECTIVES

Invocation logic encompasses two different aspects:

- Controlling the number and type of active services
- Controlling the volume of the traffic injected by active services

The objectives of the invocation logic are:

- To maximize the number of admitted services and the QoS they enjoy, thus maximizing network utilization while preventing QoS degradation caused by overloading the network
- To effectively resolve eventual congestion
- To provide fair treatment to services networkwide

OVERVIEW

Invocation logic is based on a feedback model, driven by monitoring events. We distinguish between the *dynamic admission management* and *admission strategy enforcement* functions (Fig. 4).

An admission control strategy specifies the

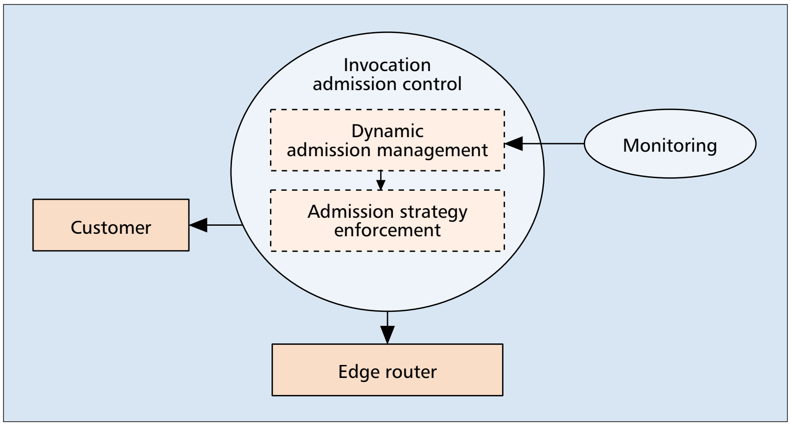

Figure 4. *Invocation admission control decomposition.*

explicit service invocation admission probability and the service rate percentage to be offered to active services in a range between their fully and almost satisfied service rates.

The dynamic admission management function, driven by monitoring events, deduces the appropriate admission control strategy to be realized by the admission strategy enforcement function. The latter interacts with the customer to handle explicit service invocation requests, and with the edge router to configure the traffic conditioning settings that determine the admitted service rate and consequently the degree of satisfaction for active services.

Separating the logic from the enforcement of the admission strategy allows decoupling from the underlying technologies for invocation protocols, AAA, and service activation, and, most important, relieves the service invocation process of expensive calculations. In the following paragraphs we concentrate on the description of the dynamic admission management function.

DYNAMIC ADMISSION MANAGEMENT RATIONALE

Monitoring events received by the dynamic admission management are in the form of edge-to-edge network status alarms and threshold crossing events on aggregate traffic injected into the network by the local edge.

Network status is expressed by a binary flag per TT, indicating whether the delay and loss requirements associated with the corresponding QoS class are (close to being) violated or not; the network status per TT is said to be *red* or *green*, respectively. On red status the invocation admission logic has to react to immediately resolve the congestion.

Locally injected traffic is projected to the resource availability buffer per TT, to estimate the risk in QoS deterioration as a result of admitting traffic. Since congestion cannot be accurately predicted, no hard limits are imposed in the utilized resources. Instead, when the injected traffic is considered to reach a critical point, admission strategies need to proactively protect the network and the users conforming to their contractual traffic levels from greedy users tending to exhaust the network resources.

Determining the critical area in the resource availability buffer influences how conservative the admission logic will be. Due to the highly uncontrollable environment, we argue that a universal optimum solution to resolve the related trade-off cannot be given.

Following the admission control principles, we introduce the *precaution level* policy parameter, denoting the level of precaution the SP considers acceptable to take in order to avoid network overloading and consequently service satisfaction degradation at the expense of network utilization. The precaution level designates the critical area in the resource availability buffer. In addition, the admission strategies are parameterized to be more/less conservative in terms of the precaution level.

DYNAMIC ADMISSION MANAGEMENT OPERATION

Dynamic admission management operates per TT.

Admission strategies are finite, predetermined, and classified into standard, mild, and severe depending on the incurred service satisfaction penalties; services are *fully satisfied*, at least *almost satisfied*, and less than *almost satisfied*, respectively.

The dynamic admission management algorithm specifies the finite state machine of traveling across admission strategies, triggered by monitoring events. The algorithm operates in three distinct states: normal, reactive, and normalization (Fig. 5).

In the normal state the network status is green, and the target is to maximize service satisfaction while avoiding building congestion. Standard or mild strategies may apply, depending on the volume of the locally injected traffic with respect to the critical area.

In the reactive state the network status is red, and congestion needs to be resolved in a fair way networkwide. The algorithm determines strategies of increased severity should this state persist. As a result, service satisfaction degradation may for some active services drop below their almost satisfied rates. It should be noted that convergence in resolving congestion is guaranteed thanks to the centralized nature of the offline TE functions determining the resource availability buffer.

Normalization is the transit state from reactive to normal state, when the network status becomes green again. Its target is to provide for a smooth and effective mediation as dictated by the precaution level, avoiding oscillations in network status. The algorithm relaxes in a randomized fashion the severity of the applied strategies.

SCALABILITY

The entities of the environment affecting the scalability of the proposed solution are:

- The number of network edges
- The number of QoS classes
- The number of subscriptions
- The total number of SLSs included in the subscriptions

Note that as SLSs denote the legs of a connectivity service, their number for a particular subscription is bound by the square of the number of network edges. Moreover, since an SLS may imply traffic for a number of QoS classes, the number of TTs per subscription is bound by the square of the number of network edges multiplied by the number of the supported QoS classes.

SUBSCRIPTION ADMISSION CONTROL

A subscription request is analyzed to its constituent SLSs. For each SLS the implied demand is derived per TT and aggregated with the corresponding demand of the established subscriptions. The admission decision involves a comparison of the new total traffic demand with the threshold designating the admissible area. Therefore, the complexity of the subscription admission process primarily depends on the number of TTs per subscription; hence, the process scales.

As subscription requests are processed sequentially for consistency reasons, the number of simultaneous subscription sessions depends entirely on the implementation and processing capabilities of the infrastructure.

In case a subscription request is not admitted, negotiations may initiate. As the adopted Service Negotiation Protocol (SrNP) [7] is session-oriented and dialog-based, and provides a means to terminate the session, the messages per session can be bound and denial of service on initiating new sessions prevented.

INVOCATION ADMISSION CONTROL

Since the algorithm operates per TT, the admission strategy to apply is deterministically deduced by the current state of the algorithm and the received monitoring event. This task involves looking-up into directly indexed data per TT. Therefore, the complexity of dynamic admission management grows linearly with the number of affected TTs, hence it scales.

Dynamic admission management may result in adjusting traffic conditioners in the edges, the number of which networkwide is bound by the number of subscribed SLSs. Therefore, it scales as long as the number of subscribed SLSs is scalable. Note that the set of traffic conditioners is nonetheless limited by the vendor-specific capabilities of the edge routers; this very fact poses requirements on the number of SLSs and therefore on the number of subscriptions established in the network.

Important to the scalability of dynamic admission management are the monitoring requirements. Dynamic admission management requires load measurements of local traffic injected into the network and network ability to deliver QoS per TT. The first measurements scale since they are performed locally through passive monitoring. For the second type of measurements, efficient monitoring mechanisms and techniques for aggregate and/or hop-by-hop measurements need to be provided. It is important to note that from the perspective of invocation logic, the inherent errors in the measurements are not deemed critical.

At explicit service invocation times, authentication and administrative- and resource-based authorization are performed. These operations involve looking up subscription data and probability-based admission per involved TT. The admission probability has been determined per TT by dynamic admission management. Therefore, the complexity of the explicit service invocation process primarily depends on the number of TTs per invocation; hence, it scales. The actual response times depend on the efficiency of implementation and deployment.

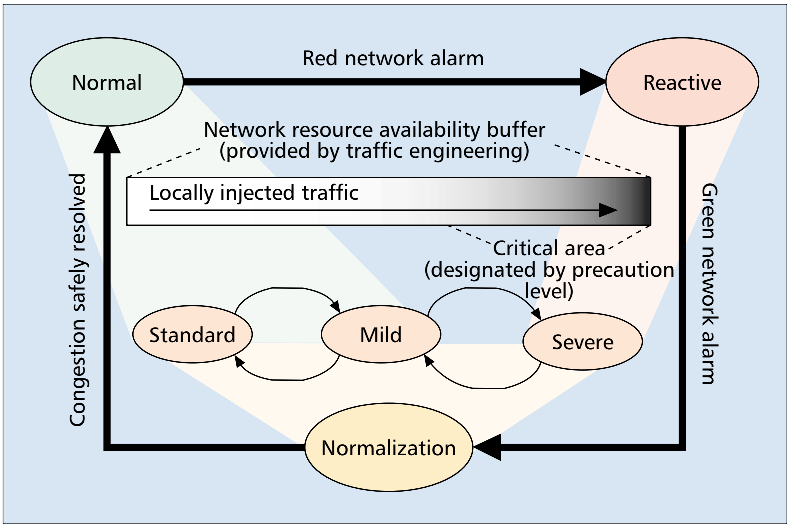

Figure 5. *Dynamic admission management operation.*

SUMMARY AND CONCLUSIONS

This article presents the principles and basic operation of a service admission control scheme for QoS delivery in IP DiffServ networks. The innovative aspects of the proposed approach can be summarized as follows.

We adopted a feedback-based model for asserting the risk of QoS deterioration. This feedback is at two levels of abstraction: on the ability of the network to deliver QoS, determined through the offline TE functions that dimension the network on the basis of anticipated demand; and on the actual status of the network to deliver QoS, provided through measurements. Based on this feedback, our scheme assesses the risk of QoS deterioration and accordingly adjusts the parameters for admitting service requests. This two-level feedback scheme allows service request response times to be minimized. Our approach differs from the majority of the studies reported in the literature, which are based on statistical and/or mathematical models for asserting on QoS deterioration before deciding whether to admit traffic.

Another aspect of our approach is its strong policy-driven nature; our approach is highly parameterized to allow it to be tuned to the special needs and features of the operational environment. Because we believe that QoS cannot be safely guaranteed a priori, we adopt "relative parameterization" based on operational policies following a best-practice paradigm; these operational policies should also be in conjunction with business policies for service provisioning.

As opposed to the majority of the schemes in the literature, our approach also applies admission control logic at service subscription epochs. We believe that through this approach the potential to efficiently resolve the traditional trade-off between QoS deterioration and network utilization is increased, since control may be exerted at different levels.

We believe that through this approach the potential of efficiently resolving the traditional trade-off between QoS deterioration and network utilization is increased, as control may be exerted at different levels.

The proposed admission scheme scales because its complexity grows in a polynomial fashion with the number of entities of the external environment influencing its behavior.

The validity of the approach has also been verified through prototype implementation in a real network environment. The implementation verified that the interactions of the proposed admission control solution with the external systems (customers, routers, traffic engineering, and monitoring systems) as well as its internal operational mechanisms and algorithms can be feasibly realized.

Experimentation showed that the response times measured at subscription and invocation epochs are in line with the expectations of the scalability analysis. Experimentation for quantitatively asserting the efficiency of the proposed approach is currently underway.

Finally, oscillations are inherent to our approach because of its feedback nature, since admission control logic is driven by network state, and network state is affected by admission control logic. Although appropriate oscillation identification and management mechanisms have been defined, they are not presented in this article due to space limitations. Their validation through experimentation is the main topic of further work.

ACKNOWLEDGMENTS

This work was undertaken in the Information Society Technologies (IST) TEQUILA project, which is partially funded by the Commission of the European Union. We would also like to thank the rest of our TEQUILA colleagues who have also contributed to the ideas presented here.

REFERENCES

[1] S. Blake *et al.*, "An Architecture for Differentiated Services," RFC 2475, Dec. 1998.

[2] P. Trimintzios *et al.*, "A Management and Control Architecture for Providing IP Differentiated Services in MPLS-based Networks," *IEEE Commun. Mag.*, vol. 39, no. 5, May 2001.

[3] Z. Zhang *et al.*, "Decoupling QoS Control from Core Routers: A Novel Bandwidth Broker Architecture for Scalable Support of Guaranteed Services," *Proc. ACM SIGCOMM 2000*, Aug./Sept. 2000.

[4] E. Knightly and N. Shroff, "Admission Control for Statistical QoS: Theory and Practice," *IEEE Network*, vol. 13, no. 2, Mar. 1999, pp. 20–29.

[5] C. Cetinkaya and E. Knightly, "Egress Admission Control," *Proc. IEEE INFOCOM 2000*, Mar. 2000.

[6] I. M. Ivars and G. Karlsson, "PBAC: Probe-Based Admission Control," *Proc. QoFIS 2001*, 2001, pp. 97–109.

[7] D. Goderis *et al.*, "A Scalable Service-Centric IP Quality of Service Architecture for Next Generation Networks," *Proc. NOMS 2002*, Florence, Italy, Apr. 2002.

[8] D. Goderis *et al.*, "Service Level Specification Semantics, Parameters and Negotiation Requirements," Internet draft, <draft-tequila-sls-01.txt>, June 2001; http://www.ist-tequila.org/sls.

[9] B. Devalla *et al.*, "Adaptive Connection Admission Control for Mission Critical Real-Time Communication Networks," *Proc. IEEE MILCOM*, Oct. 1998, pp. 614–20.

[10] R. Mortier *et al.*, "Implicit Admission Control," *IEEE JSAC*, vol. 18, no. 12, 2000, pp. 2629–39.

BIOGRAPHIES

ELENI MYKONIATI (mykel@telecom.ntua.gr) received a B.Sc. degree in computer science from the University of Piraeus, Greece, in 1996. Since 1997 she works as a research associate in the Telecommunications Laboratory at the National Technical University of Athens. In 1998 she became a Ph.D. candidate. Over the last years she has also worked for the Knowledge and Database Systems Laboratory of NTUA and Telscom AG in Bern, Switzerland. She has been involved in the REFORM, VIKING ACTS, and TEQUILA IST research projects. Her main research interest is in QoS-aware admission control.

CHARALAMPOS CHARALAMPOUS (charalam@telecom.ntua.gr) received a B.Sc. degree in computer engineering and informatics from the University of Patras in 1999. That year he became a Ph.D. candidate at the National Technical University of Athens. Currently, he works as a research associate at the Telecommunication Laboratory of NTUA. He has been involved in the TEQUILA IST research project.

PANOS GEORGATSOS (Pgeorgat@egreta.gr) received a B.Sc. degree in mathematics from the National University of Athens, Greece, in 1985, and a Ph.D. degree in computer science, with specialization in network routing and performance analysis, from Bradford University, United Kingdom, in 1989. He is working for Algonet S.A., Athens, Greece, as head of the R&D group. His research interests are in the areas of service management and network performance evaluation.

TAKIS DAMILATIS (pdamil@egreta.gr) received a B.Sc. degree in electrical and computer engineering in 1997 and an M.Sc. degree in telematics in 1999 from the National Technical University of Athens. He is working for Algonet S.A. since 2001 as a research fellow in the Telecommunications R&D Group. He has participated in the IST TEQUILA project, working on XML-based modeling of QoS-based IP connectivity services and service management issues.

DANNY GODERIS (Danny.Goderis@alcatel.be) received an M.S. degree in physics and a Ph.D. in mathematical physics from the Catholic University Leuven, Belgium, in 1986 and 1990. After some more years at the university he joined the Alcatel Corporate Research Center in 1998. He is currently responsible for IP QoS technologies in the Alcatel Corporate Network Strategy Group and is leading the IP traffic performance studies. He is project manager of the European IST project TEQUILA.

PANOS TRIMINTZIOS (P.Trimintzios@eim.surrey.ac.uk) holds a B.Sc. in computer science and an M.Sc. in computer networks, both from the Computer Science Department, University of Crete, Greece. From 1995 to 1998 he was a research associate at ICS-FORTH, Greece, working on research projects involving high-speed network management and charging. He is currently a research fellow at the Center for Communication Systems Research (CCSR), University of Surrey, United Kingdom, where he is also about to finish his Ph.D. His main research interests include traffic engineering, QoS provisioning, policy-based networking, service management, network performance control, and network monitoring.

GEORGE PAVLOU (G.Pavlou@eim.surrey.ac.uk) is professor of communication and information systems at the Center for Communication Systems Research, Department of Electronic Engineering, University of Surrey, United Kingdom, where he leads the activities of the Networks Research Group. He holds a diploma in engineering from the National Technical University of Athens, Greece, and M.Sc. and Ph.D. degrees in computer science from University College London, United Kingdom. His research interests include network planning and dimensioning, traffic engineering, policy-based networking, programmable/active networks, and multimedia service control. He has contributed to standardization activities in ISO, ITU-T, TMF, OMG, and IETF.

DAVID GRIFFIN (D.Griffin@ee.ucl.ac.uk) received a B.Sc. degree in electrical engineering from Loughborough University, United Kingdom, in 1988. He has been involved in communications networks research for the last 14 years, initially at Marconi Communications, United Kingdom, then at ICS-FORTH, Greece, joining University College London in 1996. He is currently a senior research fellow in the Department of Electronic and Electrical Engineering, where he is also completing a part-time PhD. His main research interest is in the planning and dynamic management of QoS in multiservice networks, especially in the areas of traffic engineering and routing in the Internet.