Admission Control in Multiservice IP Networks: Architectural Issues and Trends

Solange Rito Lima, Paulo Carvalho and Vasco Freitas
University of Minho, Department of Informatics, 4710-057 Braga, Portugal
{solange,pmc,vf}@di.uminho.pt

Abstract— **The trend toward the integration of current and emerging applications and services in the Internet has launched new challenges regarding service deployment and management. Within service management, admission control (AC) has been recognized as a convenient mechanism to keep services under controlled load and assure the required QoS levels, bringing consistency to the services offered.**

In this context, this article discusses the role of AC in multiservice IP networks and surveys current and representative AC approaches. We address and compare the architectural principles of these AC approaches and their main features, virtues and limitations that impact on the quality control of network services. We identify important design aspects that contribute to the successful deployment of flexible and scalable AC solutions in multiservice networks.

I. Introduction

Supporting today's Internet service heterogeneity and integration while at the same time assuring consistent quality of service (QoS) levels requires new service management paradigms, protocols and control mechanisms. In multiservice networks, admission control (AC) is viewed as a convenient means to assure high quality communication by safeguarding enough resources availability for customer traffic. This can be particularly useful for services such as IP telephony and video conferencing [1].

However, introducing specific QoS control mechanisms into IP networks has been a controversial issue. Overprovisioning of communication resources is a method often used to provide QoS in network backbones so as to avoid or reduce network control complexity. Although overprovisioning might be an attainable solution for some service providers, it should not be assumed to be a widely available and permanent answer. In fact, it is likely that the number of users and the demands of their applications will continue to outgrow the availability of resources. Thus, there is the need for additional service and traffic control mechanisms to guarantee that QoS commitments can be precisely specified and honored. Despite this need a key aspect and a major objective in the deployment of such mechanisms in real networks should be that the network control plane is kept as simple as possible.

In general, the QoS guarantees and predictability required by a service determine the control complexity inherent to an AC strategy. Reaching a good compromise between service guarantees, complexity and efficient resource utilization is a major challenge. The challenge increases in multiservice networks because service classes have distinct characteristics and demand different QoS assurance levels. Adding to the challenge is the case of end-to-end QoS delivery where multiple and heterogeneous domains need to agree and fulfil Service Level Specifications (SLSs) established among themselves.

The main purpose of this paper is to discuss existing AC proposals with an emphasis on their main features, advantages and limitations in controlling multiple service levels. The analysis aims to understand, identify and compare representative AC approaches and to point out strategic directions toward improving AC. In this discussion, we highlight relevant design principles for the scalable management of QoS and SLSs in multiclass networks.

The remainder of this article is organized as follows: the main characteristics of current AC proposals are identified in Section II; representative AC approaches are discussed in Section III; important architectural principles to achieve deployable AC solutions are highlighted in Section IV; finally, conclusions are provided in Section V.

II. Relevant characteristics of AC approaches

Important high-level characteristics distinguishing AC approaches have been identified as follows:

(i) *the underlying network paradigm* - the network model under which AC operates ranges from single service (best-effort) to multiservice architectures, following flow or class-based paradigms. Their scope as regards targeting an intradomain, interdomain and/or end-to-end solution also varies;

(ii) *the type of service to control* - the type of service usually depends upon the applications' characteristics, whether they are rigid or adaptive and whether they have quantitative or qualitative QoS targets that determine the service level guarantees to be provided. Common and similar terminology includes guaranteed vs. predictive, guaranteed vs. controlled load or hard vs. soft real-time services;

(iii) *the signaling supported* - signaling is closely related to the applications' ability to explicitly inform the network of their needs. This is commonly expressed in terms of a traffic profile and/or QoS requirements, using soft or hard state signaling for that purpose. Signaling may also occur at higher level, for instance between specific nodes in distinct network domains or directly between end-systems;

(iv) *the location of the AC decision* - this aspect is related to the centralized or distributed nature of AC. This can be further detailed depending on which nodes (e.g., all nodes or specific nodes) are involved and how they participate in the AC process. For instance, a node can make an AC decision or only gather information for some other control entity to use. The amount and type (per-flow or per-class) of state information kept in those nodes and the need for coordination among them are also important factors to consider;

(v) *the characteristics of admission decision criteria* - these can be determined by: the nature of the algorithm, i.e., whether it is parameter-based, measurement-based or hybrid[1]; the information used for AC, which can be based on keeping track of resources' usage (usually bandwidth) or on congestion indicators (e.g., explicit congestion notification (ECN) marks); and the concrete AC equations, which can be based on more or less intricate theoretical concepts involving distinct control parameters, whose tuning will, in turn, influence the conservative nature of AC decisions.

Having discussed these points, the overall performance of an AC approach can be characterized by several related aspects, namely: the ability to fulfill the QoS commitments; the efficiency of resource utilization for the service levels provided; the overhead introduced in the network data and control planes that influence scalability; and the latency regarding the time it takes to make an AC decision. The ease of migration and implementation in real environments is another key point as it provides a practical perspective on the real usefulness of the AC approach.

III. Surveying AC approaches

The following discussion covers AC in IP networks under distinct service paradigms, with emphasis on proposals for multiservice networks. In more detail, centralized and distributed AC oriented to flow, hybrid or class-based QoS models, as well as active/passive measurement-based AC proposals are discussed.

A. Intserv and RSVP aggregation

Although independent from the Intserv architecture, RSVP is cited as a convenient explicit setup mechanism to signal per-flow resource requirements in order to sustain node-by-node AC and resource reservation aiming at a guaranteed end-to-end QoS delivery[2].

The impairments of large scale deployment of Intserv/RSVP have motivated the aggregation of individual flow requests (see RFC 3175) . This aggregation process aims at avoiding per-flow signaling and per-flow state information in the core, at the cost of reducing the isolation among flows. In this process, interior nodes only maintain a reservation state for aggregates; their state only changes when the corresponding aggregate reservation needs to be updated (increased or reduced) with a new bandwidth bulk. The level of aggregation or bulk size influences the admission of the flows, the utilization and the demand for signaling in the core. While large bulks negatively influence the acceptance of flows and utilization, small bulks influence these aspects positively at expense of more signaling. The need for signaling the aggregation region also depends on traffic load variability and, ultimately, under high variability and low aggregation, the process tends to become a per-flow reservation process [2].

B. Intserv/Diffserv integrated solutions

According to the framework for Intserv/Diffserv operation (see RFC 2998), Intserv, RSVP and Diffserv are complementary technologies which can facilitate pursuing the objective of a scalable end-to-end quantitative QoS solution. While Intserv/RSVP allows per-flow request signaling quantifying the resources needed and obtaining a corresponding AC feedback, Diffserv enables scalability in large networks. In this framework, end-to-end RSVP signaling requires at least that RSVP messages are carried out across the Diffserv region. But depending on the specific realization of the framework, none, some (e.g. border routers) or all routers in the Diffserv region may process these messages. The coexistence between the two architectures assumes the control of the amount of traffic submitted to the Diffserv region, which must be able to support Intserv-like services through proper Per-hop Behaviors (PHBs) and Intserv/Diffserv parameters mapping (see Figure 1). The options for resource management in the Diffserv region may include: (i) static provisioning; (ii) dynamic provisioning using RSVP; and (iii) dynamic provisioning resorting to other means, such as Bandwidth Brokers (see Section III-D). In practice, despite the guarantees provided in the Intserv/RSVP region, end-to-end services guarantees depend on the resource management policies and supported services within Diffserv regions. Ongoing IETF work on this topic [3] is covered in Section III-E.1.

C. DPS and SCORE architecture

The service architecture proposed in [4] aims to offer per-flow delay and bandwidth guarantees similar to the Intserv guaranteed service, but in a more scalable way. This architecture, called Scalable Core (SCORE), is based on: (i) bringing per-flow management to edge nodes; (ii) a stateless core, i.e., a core where no per-flow information is maintained; and (iii) a dynamic packet state (DPS) technique, which uses specific fields of the IP packet header to embed per-flow state. Core nodes process each packet based on its state, update it and, eventually, their own state before forwarding the packet.

DPS technique is the key concept of SCORE architecture allowing to coordinate routers' actions and implement distributed algorithms. The packet state inserted at ingress nodes and removed at egress nodes is used by each core node to perform scheduling (based on the concept of packet eligible time and deadline) and to support per-hop AC.

Despite avoiding per-flow state in the core and providing a guaranteed service, this architecture requires that all routers in the flow's path participate in the AC process, implement the same scheduling mechanism and update packet headers. The proposal for packet state insertion in packet headers

[1] *Parameter-based* AC algorithms take into account the network resources already in use (reserved) by accepted flows and the resources the new flow will consume, according to its explicit traffic descriptor. *Measurement-based* AC algorithms take into account measures reflecting the impact of existing flows on the network load and/or QoS before deciding about a new admission.

[2] The framework Next Steps in Signaling (NSIS) (see RFC 4080), proposed recently, contemplates a wider variety of possible signaling scenarios, being more versatile and flexible than RSVP.

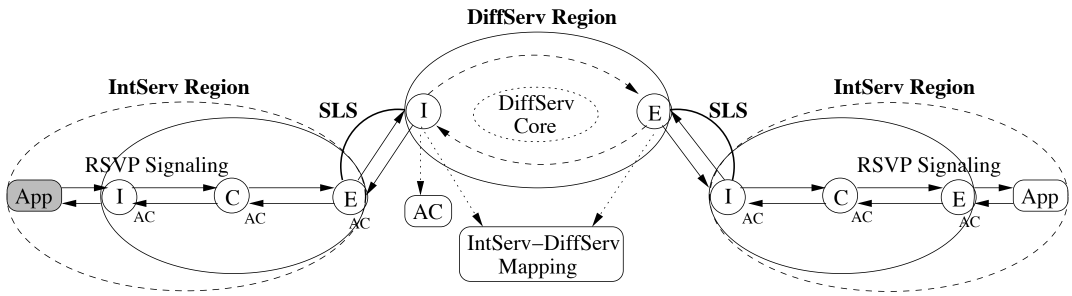

Fig. 1. Intserv/Diffserv integrated solution

may reveal itself incompatible with existing protocols and mechanisms such as Diffserv marking, headers compression and encryption. Although presented as an end-to-end solution, the operation crossing multiple domains is not covered in [4].

D. Centralized approaches based on BBs

One of the first approaches to perform resource management and AC in a Diffserv domain suggests the use of a central entity called Bandwidth Broker (BB) [5]. The principle behind BB architecture is to introduce several service management tasks in a network domain in order to provide consistent QoS, without complicating the control plane inside the network. This is achieved by centralizing information concerning network resources and their usage, domain topology, service policies and negotiated SLSs, all of which are required to perform control tasks such as AC, removing these tasks and the corresponding state information from the network core.

As far as AC is concerned, at an intradomain level, when a new flow requires admission, a signaling message is sent to the BB specifying the flow profile and QoS requirements. The BB, after authenticating and authorizing the request, makes a decision, taking into consideration the domain service policies, the corresponding SLS usage and the available resources along the path. If the destination is outside the domain, the AC decision may involve interdomain signaling with downstream BBs, thus extending the AC process and resource reservation end-to-end. According to the resulting AC decision, each BB updates its state information databases and configures the involved edge nodes consistently. For scalability reasons, the AC requests to BBs, the reservations and the interdomain communication should consider flow aggregation. Figure 2 illustrates the operation described above; a more detailed description of the BB functionality and operation is available in [5].

In [6], a BB's architecture based on a core stateless Virtual Time Reference System and DPS technique is suggested in order to achieve a scalable solution to provide guaranteed services without requiring per-flow state in core routers.

The main advantage of centralized AC approaches is that centralizing state information and control tasks allows a global vision of the domain's QoS and operation, relieving the control plane inside the network. Centralization also facilitates creating and changing service policies and control mechanisms such as AC algorithms. However, the cost of centralized approaches is high. BBs need to store and manage large amounts of information, which in large and highly dynamic networks with many signaling messages and information state updates that need to be processed in real-time is either difficult or prohibitive. The congestion and functional dependence on a single entity is another well-known problem of centralization.

To improve reliability and scalability in large network domains, several approaches suggest the use of a distributed or hierarchical architecture involving multiple BBs in the domain instead of a single centralized BB [5], [6]. A single BB strategy is considered more suitable to small and less dynamic environments involving long-lived flows. In the case of large and more dynamic domains, the use of multiple BBs improves reliability, BB congestion avoidance and scalability, but at an eventual cost in coordination among BBs and in fragmentation of resources.

E. Measurement-based AC approaches

AC approaches based on network measurements performed node-by-node, edge-to-edge or end-to-end erupted within the context of providing predictive service guarantees. Measurement-based AC aims to solve or reduce the disadvantages of the AC approaches described above, in particular, regarding the state information and control overhead, at an eventual cost of QoS degradation. Measuring network utilization and congestion can be expressed by the estimation and control of parameters such as bandwidth, delay, loss or ECN marks during a given measurement period.

1) Passive measurement-based AC: The term *passive* stems from the fact that the measurement process resorts to real traffic within the network for parameters' estimation.

In the context of Intserv, MBAC has been proposed to assist the provision of a predictive service for tolerant applications able to accommodate occasional delay-bound violations [7]. As the behavior of existing flows is determined by measurements rather than by their rate reservations (e.g., worst-case parameters), the service provided is less reliable due to traffic fluctuations. However, this allows improved AC flexibility and takes advantage of statistical multiplexing, which may lead to significant utilization gains. In [7], AC is distributed node-by-node, using rate and/or delay-based equations. Other relevant MBAC algorithms are presented in [8].

Taking into account the burden of performing AC in all network nodes in relation to the changes and overhead introduced in those nodes, a different type of passive MBAC considers the measuring of edge-to-edge network status without requiring

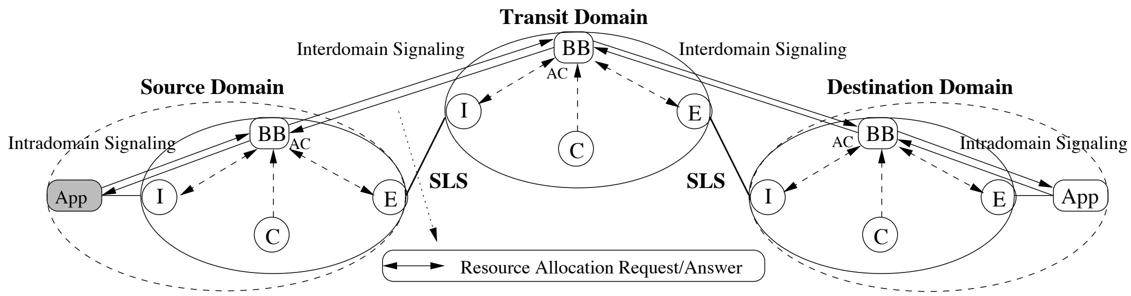

Fig. 2. Centralized AC using BBs

additional processing in the network core. AC is then left for network edges such as ingress nodes, egress nodes or both.

In the context of multiclass networks, an AC solution based on the theory of traffic envelopes is proposed in [9]. In this proposal, egress nodes assume a preponderant role as they perform both edge-to-edge aggregate traffic measurements and AC. The measurements passively and continuously assess available service on a path between ingress-egress pairs, without involving per-flow state in any network node and ignoring core details. Despite the scalability resulting from not involving the network core, the need for continuous ingress-egress measurements and updates in all real packets makes the solution more suitable to a single domain, rather than to end-to-end. Moreover, the problem of controlling SLSs is neither covered in this approach nor in the active AC proposals discussed below.

Ongoing work within IETF targets the deployment of end-to-end Controlled Load service. It considers the access to large Diffserv regions based on edge-to-edge distributed measurement-based admission control and flow preemption [3]. This proposal, in accordance with the Intserv-Diffserv framework (see Section III-B), is based on the pre-congestion notification (PCN) concept and requires per-flow admission state at edge routers only and PCN markers operating on the aggregate traffic at all routers.

2) Active measurement-based AC: In opposition to passive measurement, the designation *active* measurement is adopted when specific traffic, called probing traffic, is injected into the network for measurement purposes.

As regards AC, this technique is designed to overcome the overhead associated with signaling and AC processing in network nodes, leaving the responsibility of inferring the network congestion status and deciding on flow admission only to endpoints (end systems or edge routers). This inference process resorts to per-flow probing traffic in order to obtain measures of delay, jitter, loss or ECN marks reflecting the congestion along the corresponding path. It also simultaneously assesses the ability of the path to support the new flow. AC approaches based on this technique are commonly called Endpoint Admission Control, Probe-based Admission Control or End-to-end Measurement-based Admission Control (EMBAC) [10]. Generically, in EMBAC, the admission of a new flow is preceded by a *probing phase* for congestion inference. Upon receiving AC feedback, the sender endpoint either enters into a *data phase*, where flow packets are sent, or aborts the sending process. In order to increase robustness, the sender implements a timeout mechanism associated with the start of the probing phase to deal with missing feedback. Figure 3 illustrates this behavior.

Existing EMBAC approaches differ in several aspects: (i) the measured parameter involved in the AC decision, e.g., packet loss ratio; (ii) the characteristics of the probing phase such as its duration and/or rate; and (iii) the underlying network model. As regards the type of service, EMBAC solutions are intended to have the same applicability as other MBAC solutions, i.e., soft real-time services. A detailed discussion of EMBAC performance for a simplified network model with two priority service levels is available in [10]. Despite the simplicity and scalability of EMBAC solutions, requiring few or no changes in the networks, several disadvantages are commonly cited, namely: (i) the significant initial latency or setup delay which may limit its attractiveness for certain applications; (ii) the overhead of per-flow probing traffic which, depending on the weight and overlapping degree of the current probing phases, may lead to bandwidth stealing and thrashing regimes [10]; and (iii) the measurements' dependency on instantaneous network congestion. As regards fairness, a common concern of MBAC and EMBAC solutions is that both usually imply a single decision policy that tends to favor small flows, flows with more relaxed QoS objectives and flows that traverse smaller path lengths [8], [10].

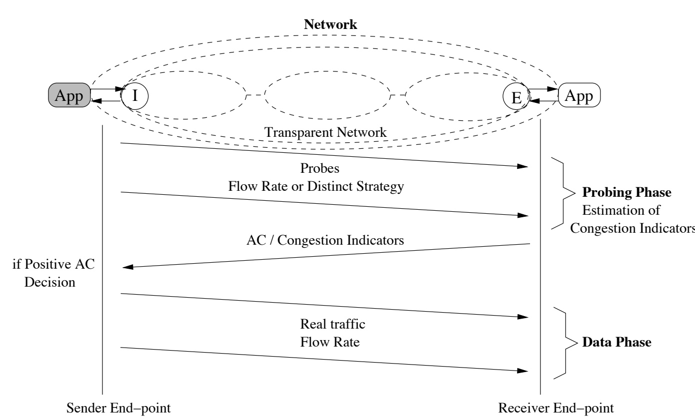

Fig. 3. EMBAC approach

F. AC proposals to control elastic traffic

Although there is consensus on performing AC for real-time traffic, the need to control the admission of elastic TCP traffic is controversial. Some argue that once TCP is adaptive, controlling the number of flows sharing the available bandwidth is unnecessary. Others argue that controlling the overload is required in order to preserve an acceptable throughput per active flow, thus ensuring the QoS offered to users. In fact, a minimum TCP bandwidth is required to achieve a minimal session level user utility [11]. The use of AC will assure this user utility and avoid wasting network resources on retransmissions and incomplete transfers.

Due to the large number of TCP flow arrivals and their eventual short duration, controlling individual flows using explicit signaling and reservations is not feasible. Therefore, a measurement-based AC approach for elastic traffic is usually proposed to assure that the solution reacts and scales properly. Without per-flow signaling, the detection and acceptance/rejection of a new flow is implicit. Common implicit AC criteria use the estimation of current load, available bandwidth or packet loss probability and compare the obtained estimation with a pre-defined threshold. The threshold may depend on an upper limit of admitted flows. These estimates can relate to a link or path; however, path estimations are preferable when considering AC performed only at ingress nodes.

Within implicit AC, the simple discard of initial flow packets is usually enough to inform the source of a rejection decision. Otherwise the packets will proceed. Possible solutions to support detection and subsequent AC decision are: (i) to intercept packets initializing the TCP connection (TCP SYN and/or SYN ACK); and (ii) to maintain a list of accepted and active flows based on the corresponding flow identifiers. The former solution is easy to implement; the latter is more flexible but may be critical for high-speed interfaces due to its potential overhead.

Implicit AC can also be applied to traffic other than TCP. For instance, it can be applied to UDP traffic from real-time applications that do not send explicit signaling to the network.

G. AC proposals for QoS and SLS Control

A relevant issue for the deployment and management of multiservice networks is the control of QoS and SLSs, both intra and interdomain, in a flexible and simple way.

In this context, the distributed AC model proposed in [12] aims to: (i) support multiple services with distinct assurance levels; (ii) control the QoS levels inside each domain and the existing SLSs between domains; (iii) operate intra and interdomain, providing an unified end-to-end solution; and (iv) be simple, flexible, efficient, scalable and easy to deploy in real environments. This proposal considers a service-dependent degree of overprovisioning in order to achieve a simple and manageable multiservice AC solution. Such overprovisioning allows for the simplification of the AC process while providing the required service level guarantees.

In the model operation, illustrated in Figure 4, only edge nodes are involved, leaving the network core unchanged. While ingress nodes perform explicit or implicit AC depending on the application type and corresponding traffic class, egress nodes perform on-line QoS and SLS monitoring. QoS and SLS monitoring statistics are sent to the corresponding ingress routers to update an ingress-egress service matrix used for distributed AC and active service management.

More specifically, the *Ingress-Egress QoS Monitoring* task measures relevant parameters for each service (service metrics) using appropriate time scales and methodologies. The resulting measures are expected to reflect the service available from each ingress node and are used by a QoS Control rule to drive AC decisions. This rule checks the controlled parameters of each service class against pre-defined thresholds to determine an AC status for the measurement time interval. The *SLS Control* task monitors the usage of downstream SLSs at each egress, to ensure that traffic to other domains does not exceed the negotiated profiles and packet drop will not occur due to a simple and indiscriminate traffic conditioning process. An SLS Rate Control rule checks if the SLS can accommodate the traffic profile of the new flow, which complements the AC decision process. In implicit AC, as flows are unable to describe a rate profile, AC is restricted to QoS control. Thus, flows are accepted or rejected implicitly according to the current AC status, which is computed once for each measurement interval.

In this model, the control of the negotiated QoS parameters of the SLSs is embedded in the QoS control of the corresponding service classes, reducing the amount of SLSs' dynamic state information and control overhead. It is of paramount importance that the expected SLS and flows' parameters values be embedded into their respective class parameter target values so as to simplify QoS and SLS control in the domain.

The *end-to-end* operation is viewed as a repetitive and cumulative process of AC and available service computation performed at ingress nodes. At each domain, the ingress node decides if a flow can be accepted and, if so, the service metric values in the domain are added to the flow request to inform the downstream domain of the service available so far. Thus, each domain performs AC using both the incoming and its own measures. Note that a cumulative process for end-to-end QoS computation is consistent with the cascade approach for the support of interoperator IP-based services, which is in conformance with the Internet structure and operation, and more scalable than the source-based approach [13].

A cumulative approach for end-to-end QoS support has also been proposed in the Mescal project, where AC is performed at ingress nodes based on algorithms that consider the available bandwidth on those nodes. This bandwidth reflects the minimum available bandwidth in any of the possible paths [14]. The multidomain approach is based on peak rate control at the ingress nodes. Although the AC rules in the Mescal project do not address the control of SLSs, its deliverables provide an in-depth analysis of and proposals for the problem of intra and interdomain SLS management.

H. Comparison summary

In order to better identify and compare the several AC solutions discussed above, Table 5 summarizes their main features,

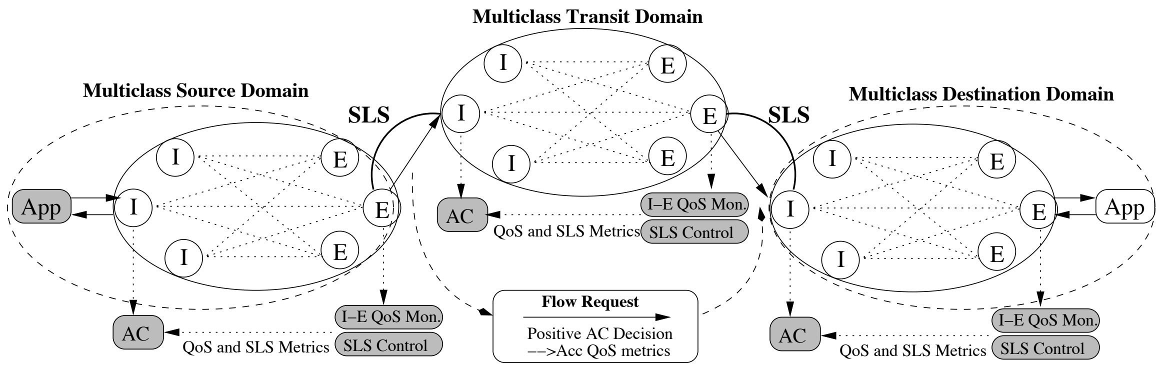

Fig. 4. Distributed monitoring-based AC approach

advantages and disadvantages. The aspects discussed take into consideration the high-level characteristics that distinguish AC approaches addressed in Section II.

IV. Design principles for scalable AC

The discussion above allows to identify several important issues as main design principles in the pursuit of scalable and deployable AC solutions to the management of multiple service levels in today's networks. These issues include: relieving the network core from control tasks, reducing state information and control overhead, sensing and adapting to network dynamics through measurements and supporting AC irrespective of the ability of applications to explicit requirements and signaling the network.

The relevant AC design principles are the following:

(i) Control at network edges - this avoids hop-by-hop AC overhead and provides a convenient level of abstraction and independence from network core complexity and heterogeneity;

(ii) State information at service class level - this follows the aggregation principle, leading to reduced state information. In addition, maintaining control on a ingress-to-egress basis is particularly suitable for SLS auditing;

(iii) per-Class systematic on-line monitoring - considering that monitoring is an essential management task in multiservice networks, AC can be designed to take advantage of network performance feedback. Measuring traffic load and QoS metrics per-class systematically allows for self-adaptive service management, avoids per-application intrusive traffic to obtain measures and reduces AC latency as measures are available on-line. Furthermore, systematic measurements have an intrinsic auto-corrective nature, allowing the detection of short- or long-term traffic fluctuations depending on the measurement time interval. These measurements implicitly take into account the effect of cross-traffic and other internally generated traffic (e.g., routing, multicast and management);

(iv) Flexible AC criteria - the rules upon which AC is based should be flexible enough so that they can be adjusted to the semantics and needs of each service. This flexibility should also accommodate the technological evolution of services and applications. Both the service-dependent nature of AC rules and the conceptual modular independence between AC and monitoring tasks increase the ability to integrate new improvements and developments;

(v) Controlled degree of overprovisioning - considering a controlled service-dependent degree of overprovisioning is a relevant aspect in achieving a simple and manageable multiservice AC solution. This allows the relaxation of the AC process and widens the range and assurance of the service types it supports.

V. Conclusions

This paper has discussed and compared representative AC approaches oriented to multiservice IP networks. The discussion clearly illustrates the compromise between the level of QoS guarantees and the complexity introduced in the network control plane. The above analysis of the evolution of AC approaches and several AC perspectives argues for the adoption of solutions based on measurements of network usage and performance rather than solutions that bring too much state information about reserved resources into the network. In this context, important design principles to achieve deployable and scalable AC solutions for multiservice networks have been identified.

References

[1] J. Babiarz, K. Chan, and F. Baker, "Configuration Guidelines for Diffserv Service Classes," RFC 4594, Aug. 2006, informational. [Online]. Available: http://www.ietf.org/rfc/rfc4594.txt

[2] H. Fu and E. Knightly, "Aggregation and Scalable QoS: A Performance Study," in *IWQoS'01*, June 2001.

[3] B. B. et al., "An edge-to-edge Deployment Model for Pre-Congestion Notification: Admission Control over a DiffServ Region (work in progress," draft-briscoe-tsvwg-cl-architecture-04.txt, Oct. 2006.

[4] I. Stoica and H. Zhang, "Providing Guaranteed Services Without Per Flow Management," in *ACM SIGCOMM'99*, Oct. 1999.

[5] K. Nichols, V. Jacobson, and L. Zhang, "A Two-bit Differentiated Services Architecture for the Internet," RFC 2638 (Informational), July 1999. [Online]. Available: http://www.ietf.org/rfc/rfc2638.txt

[6] Z. Zhang, Z. Duan, Y. Hou, and L. Gao, "Decoupling QoS Control from Core Routers: A Novel Bandwidth Broker Architecture for Scalable Support of Guaranteed Services," in *ACM SIGCOMM'00*, 2000.

[7] S. Jamin, P. Danzig, S. Shenker, and L. Zhang, "A Measurement-Based Call Admission Control Algorithm for Integrated Services Packet Networks (Extended Version)," *IEEE/ACM Transactions on Networking*, pp. 56–70, Feb. 1997.

AC Approach	Architecture	State Information	Signaling	AC criteria	Net. Paradigm / Service Type	Main Advantages	Main Disadvantages
Intserv RSVP/AggRSVP	Distributed AC in all nodes AC in AggRSVP nodes	At all nodes Per flow or aggregate	Explicit	Flow profile Available resources (bw) Reserved or Measured (MBAC)	Inserv Guaranteed or Predictive	Service guarantees Tight resources control AggRSVP–Aggregated state information	Scalability problems Control overhead State information Signaling Difficult to deploy widely
Intserv/Diffserv (IS/DS)	Distributed IS – AC all nodes/Agg. DS – AC at boundaries with dist. or cent. control	IS – At all nodes DS – At edges / BB / PDP Per flow or aggregate	Explicit	Flow profile Available resources Reserved or Measured	Intserv/Diffserv Guaranteed or Predictive Best–effort	Tradeoff service guarantees vs. scalability dependent on DS control strategy, IS region dimension and control, RSVP request type Mapping functions needed	
SCORE/DPS	Distributed AC in all nodes	Per flow at edges Per aggregate in core Embedded in packets	Explicit 2 phases: End–systems – Edges Internal signaling	Flow profile Estimated upper–bound of aggreg. reserved bw	Proprietary (IS–like) Adjustable to Diffserv Guaranteed Service	Service guarantees Low state information at core Scalable core	Proprietary Difficult to deploy widely Changes in all nodes Packets updating Route enforcement
BB	Centralized Distributed BBs	At BBs Net. resources usage Topology, SLSs, policies Per flow or aggregate Per flow at edges for TC	Explicit End–system/ingress and BBs	Flow profile Available resources SLS status Reserved or Measured	CoS –Diffserv or proprietary Guaranteed or Predictive Best–effort	Edge and core simplified Service guarantees Global network vision Flexibility to improve control rules SLSs management	Dependence on a single entity Single point of failure and congestion Large amount of info. to manage Signaling overhead (intra and interdomain) Latency and Scalability Distributed BBs: resources fragmentation BBs coordenation
Passive Mon. MBAC	Distributed AC in all nodes	At all nodes Usually per aggregate	Explicit	Flow profile Available resources (bw) Rate and/or delay Measured	Intserv–like Predictive	Less conservative Better resources usage Take advantage of stat. multiplexing Flexibility to adjust to traffic fluctuations	Difficult to deploy widely QoS degradation may occur Fairness
Passive Mon. Envelopes	Distributed AC at egress	At egress nodes Per class and Ingress–Egress pair	Explicit End–systems – Egress	Flow profile Measured service Service envelopes Arrival envelopes	CoS – Proprietary Adjustable to Diffserv Predictive	Involves only edge nodes Black–box core Low state information Service guarantees	Difficult to deploy end–to–end number of I– E combinations Changes in all packets Measurement overhead
Active Mon. EMBAC	Distributed AC at endpoints	At endpoints Per flow in probing phase No state in the network	Between endpoints Transparent to network	Congestion indicators Measured Thresholds	CoS – Proprietary or Diffserv Predictive Best–effort	Easy deployment Scalability End–to–end operation Low state information	Initial latency Per flow probing intrusion Stealing and thrashing regimes QoS estimation based on one sample QoS instability and Fairness
Elastic traffic	Distributed AC at boundaries	At ingress nodes Measures for AC Eventual per–flow tables No state in the network	Implicit Flow detection	Current load Available bandwidth Congestion indicators Measured	CoS or Best–effort Best–effort	Easy deployment Low complexity	Difficulties in high–speed interfaces Heavy flow tables (if used)
QoS and SLS Control	Distributed AC at ingress/egress	At edge nodes Per Class / Service Per Ingress–Egress pair	Explicit or Implicit	Flow profile QoS measures SLS utilization	CoS – Generic Guaranteed or Predictive Best–effort	QoS and SLS management Reduced state information and overhead Flexibility and self–adaptability Easy to deploy widely	I–E monitoring overhead

Fig. 5. Comparison of representative AC approaches

[8] L. Breslau and S. Jamin, "Comments on the Performance of Measurement-Based Admission Control Algorithms," in *IEEE INFOCOM'00*, Mar. 2000.
[9] C. Cetinkaya, V. Kanodia, and E. Knightly, "Scalable Services via Egress Admission Control," *IEEE Transactions on Multimedia*, vol. 3, no. 1, pp. 69–81, Mar. 2001.
[10] L. Breslau, E. Knightly, S. Shenker, I. Stoica, and H. Zhang, "Endpoint Admission Control: Architectural Issues and Performance," in *ACM SIGCOMM'00*, 2000.
[11] R. Mortier, I. Pratt, C. Clark, and S. Crosby, "Implicit Admission Control," *IEEE Journal on Selected Areas in Communication*, vol. 18, no. 12, pp. 2629–2639, Dec. 2000.
[12] S. R. Lima, P. Carvalho, and V. Freitas, "Self-adaptive Distributed Management of QoS and SLSs in Multiservice Networks," in *IEEE/IFIP International Conference on Integrated Management (IM 2005)*. Nice, France: IEEE Press, May 2005.
[13] P. Georgatsos, J. Spencer, D. Griffin, P. Damilatis, H. Asgari, J. Griem, G. Pavlou, and P. Morand, "Provider-level Service Agreements for Inter-domain QoS delivery," *Fourth International Workshop on Advanced Internet Charging and QoS Technologies (ICQT04)*, Sept. 2004.
[14] S. Georgoulas, P. Trimintzios, and G. Pavlou, "Joint Measurement- and Traffic Descriptor-based Admission Control at Real-Time Traffic Aggregation Points," *IEEE International Conference on Communications (ICC2004), QoS and Performance Symposium*, June 2004.