2ND QUARTER 2007, VOLUME 9, NO. 2

IEEE COMMUNICATIONS SURVEYS
The Electronic Magazine of Original Peer-Reviewed Survey Articles

www.comsoc.org/pubs/surveys

ADMISSION CONTROL IN MULTI-SERVICE IP NETWORKS: A TUTORIAL

STEVEN WRIGHT, BELLSOUTH

ABSTRACT

Admission Control (AC) has long been considered as a key mechanism to support Quality of Service objectives in networks. There is a significant base of literature in the area of admission control algorithms, but not all of these algorithms are directly comparable in terms of inputs, outputs or objectives. Current theory does not well describe the when, where, and how to best apply AC in designing network and service infrastructure for multi-service IP networks. This tutorial takes an ontological perspective within which to categorize admission control schemes. Industry-standard architectures (both core and access networks) are used to illustrate some of the key concepts. Two tables are provided to summarize the different approaches for categorization of admission control schemes.

Admission Control (AC) has historically been associated with telephony services. AC has also been proposed, by various parties (see e.g., [1, 2]), as necessary to support a full range of Quality of Service (QoS) based services for ATM and IP networks. The economic value of AC for services is difficult to evaluate using current literature on AC theory which is typically focused on specific algorithms. Implementing AC functions for those services comes with costs (e.g., requirements for support of additional signaling protocols) and scalability constraints, however the benefits are typically stated as meeting a service requirement, rather than justifying that requirement.

This purpose of this article is to provide:

- A simple introduction to the key concepts (particularly the rationale) for those new to AC
- A perspective that enables a higher level of abstraction in considering the objectives and mechanisms of proposed AC schemes
- A particular emphasis on the breadth of considerations that may be involved in AC decisions

The existence and characteristics of an AC scheme are elements of a broader network service definition. In comparing AC schemes, most of the existing literature is focused narrowly on the performance of one or more specific resource management algorithms in the context of a specific protocol or service. The objectives and characteristics of a "good" AC scheme depend on the service definition. In evaluating practical AC systems, we should also consider the effects of errors (or failures) in AC, not just the efficiency metrics. Most existing AC systems consider a single unicast session as an atomic construct, although AC schemes have also been proposed for multicast [3] and anycast [4] network services, and many services require multiple sessions with different types of information flows (e.g., combinations of media and control flows). For this article, we consider primarily the single session case and assume that more complex services can be constructed from this.

AC is not a new subject; there exists a significant base of literature and continuing research in this field. By its nature though, most such literature is typically very specific to the particular network protocols and application scenarios being considered. The main contribution of this article is to summarize (in Table 1 and Table 2) an ontology that collects the major characteristics of AC schemes. This ontology is supported through the examination of existing related taxonomies. I provide additional support for the ontology through consideration of the linkage of AC with the capacity planning processes, examination of the application in typical core and access network architectures, and consideration of the factors involved in scaling up AC as an IP network function for large-scale, multi-service, Wide-Area Networks (WANs). Additional aspects of the ontology including service mix, AC granularity and topological considerations are discussed. This section also provides some consideration of the impact on AC of resource modeling issues and administrative policy aspects. A brief summary is provided as a conclusion, with the major results presented in Table 1 and Table 2.

Who? What? When? Where? Why? Aspects	Dimensions	Examples/Considerations
Who makes the admission decision?	Customer	Calling vs called party
	Network operator(s)	Which one?
What is the admission decision that is made?	Effect on This Service Instance	Accept/ Discard/Remark this instance
	Impact on other services	Pre-emption (of an existing instance), Reduced resource availability May depend on the active Service mix
	Accounting data effects	**Where** is the accounting data collected — network core/network edge (sender/ receiver side)? **How** precise/approximate is the accounting data? **When** is the accounting data retrieved/presented? — real-time/non-real-time **What** are the collected data elements? — duration, volume, node resources (e.g. QoS class, conference bridge ports), distance **Who** pays for the resource usage? — calling party, called party, 3rd party, etc.
	Indication of decision result	Implicit/explicit, with or without parameters
	Topological scope	Node/link, Network (defined by technology or administration), End-end, Mobile (One/Adjacent/All nodes)
	Decision granularity	Packet, Burst, Flow, Session, aggregate
When is the admission decision made?	Reservation type	Instantaneous, advance
	Decision duration	Valid for Session duration, revocable decisions — preemption
	Session timing characteristics	Session duration, session arrival/ departure rates/distribution etc.
	Trigger events	User events, Mobility events, Network events (e.g. faults), time-outs
Where is the admission decision made?	Centralized	Application? network? customer?
	Distributed across	network elements, operators, technologies
Why is the admission decision made?	Network Resource Management	Congestion avoidance, detailed/dynamic resource allocation, service interaction, transient overload avoidance.
	Application Requirement/Service Definition	inelastic application constraints, accounting data events required
	Security	Application integrity, administrative controls on access

Table 1. *The who, what, when, where and why of admission control in IP networks.*

Prior Work on AC Ontologies

Existing AC classification efforts have largely focused on the characteristics and performance of AC schemes as resource management algorithms. I consider the literature regarding admission control in ATM networks. I also consider some of the key differences introduced in IP networks. I introduce the AC concepts associated with mobility. I consider AC in contrast to Media Access Control and congestion control respectively. The relationship between accounting and AC is identified. How these schemes approach the problem of AC is summarized in Table 2, but this article also has a particular focus on the rationale and the "Who? What? When? Where? Why?" aspects of AC are summarized in Table 1.

AC in ATM Networks

ATM signaling and routing (e.g., PNNI) supports a notion of Generic Call AC, with implementation-specific algorithms in various ATM switches. Perros *et al.* [5] classified ATM Call AC schemes into 5 groups based on the underlying aspect of traffic theory utilized in the AC algorithm:

How? Aspects	Dimensions	Examples/Considerations
Service models	Service definitions/objectives	Unicast/multicast/anycast, service specific resources, service instance parameters
	Service Security	Effect of excessive service requests? qualification of service users?
	Service mix interactions	Service mix definition, service isolation vs resource sharing permitted
	Service planning assumptions	Creation/deletion/modification of service definitions, infrastructure capacity planning assumptions
Admission decision input parameters	System Context Assumptions	Implicit/explicit parameters, capacity planning assumptions, protocol layers considered, latency assumptions, price/cost structures, resource type managed (buffers/bandwidth/flow IDs etc.), service models
	(Implicit vs.) Measured vs. Specified parameters	Precision, accuracy, frequency of measurement, measurement of what? — network state or traffic activity
	Traffic source models	Implicit/explicit parameters, Rate limited, Leaky bucket constrained, Stochastic (well characterized distributions, Long Range Dependent, heavy tailed approximations) Traffic source behavioral (e.g. TCP), effective/equivalent capacity models symmetric/asymmetric, elastic/inelastic
	System state/capacity models	Time continuous/sampled, granularity (uniform/non-uniform, aggregate/per user), designed aggregation ratio, rate adaptation., topology considerations, open loop vs. closed loop models of system state, buffering vs. bufferless models
Admission decision mechanisms	Explicit user-network signaling	User provided: Explicit parameters, Acceptance/acknowledgement of charges Network Provided: Results of admission decision
	Signaling internal to a distributed admission control system	Aggregation of results across: Link/node, network, service; Closed loop feedback of system state information
	Decision criteria	Loss, latency, user priority, effective/equivalent bandwidth
Admission decision performance	Network performance metrics	Utilization, Loss, Prob. of excessive denials, network costs (including cost of AC functionality)
	Service performance metrics	Service blocking probability, Fairness of instances within service class, Prob. of invalid admission
	Service mix performance	Size of admissible region, Fairness across services
	Algorithm performance	Robustness, Computational load, Predictability

Table 2. *Aspects of How the Admission Decision is Made.*

- Equivalent Capacity
- Heavy Traffic Approximation
- Upper bounds on Cell Loss Probability
- Fast Buffer/Bandwidth Allocation
- Time Windows

These 5 groups provide different approaches to statistical modeling of the source traffic and/or system state. As such they are summarized in Table 2 as examples of traffic source models or systems state/capacity models that are inputs to the admission decision. Most of the AC schemes in the literature focus on loss as a performance metric, even though ACs are typically considered most relevant for latency sensitive real-time services. Loss as a decision criterion is also listed in Table 2 as an aspect of the admission decision mechanisms. Elsayed *et al.* [6] classified the ATM CAC schemes based on whether the acceptance criteria were based on:

- Cell loss probability
- Cell delay

ATM AC schemes were classified by Shiomoto *et al.* [7] based on whether:

- The switch model was bufferless (Rate Envelope Multiplexing) or not (Rate Sharing Multiplexing)
- The evaluation is of Cell Loss Rate (CLR) or effective Bandwidth
- The method uses a declared traffic descriptor only or measurements (whether of traffic sources or network state) as well

Effective bandwidth as a decision criterion is also listed in Table 2 as an aspect of the admission decision mechanisms, and buffered or bufferless models are reflected there as an aspect of the system state/capacity modeling. The measurement rather than specification of parameters for traffic or sys-

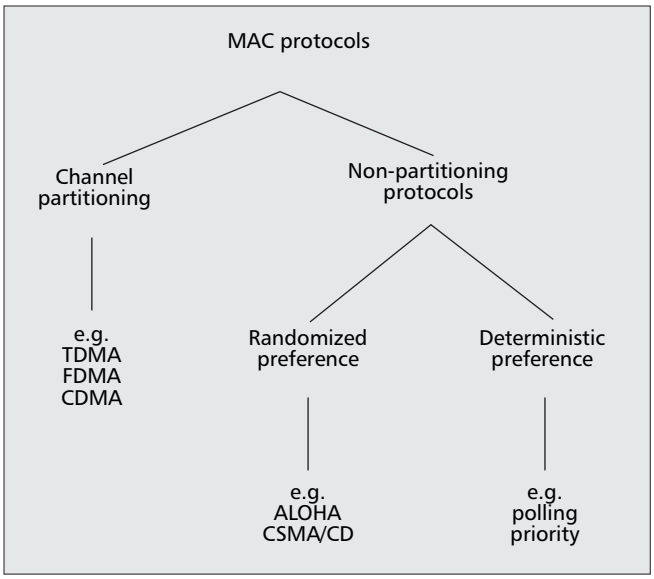

Figure 1. *MAC protocol taxonomy.*

tem state information is also listed in Table 2 as another aspect of the input parameters to the admission decision.

AC in IP Networks

Internet traffic is generally grouped into elastic vs. inelastic applications. While AC is typically considered for inelastic applications it can also be applied to elastic applications. For example, several authors have considered the capacity requirements and AC schemes related to TCP (see e.g., [8–10]). Some further discussion of the relationship between elastic/inelastic applications and AC is provided later.

Early approaches for AC within IP networks [11] provided similar notions of AC to ATM networks, but used RSVP and considered the role of policy-based network control mechanisms. Berg & Mandjes [12] separated AC schemes into Static and Measurement-Based AC (MBAC) schemes. They point out that while the work on static AC algorithms is relatively mature, most of these schemes address only a single service class and few address the problem of integrated networks supporting multiple service classes.

Work on MBAC has can be categorized into schemes resident in the network and managed by the network operator (see e.g., [13]) or to an emerging notion [14] of "End Point" Admission Control (EPAC) schemes, where the measurements are made by the user of the network. In EPAC, the host makes some measurements of network capacity and uses that information to infer network capacity before making a decision to commence an additional session. These measurements are typically achieved by having the end points send probe packets through the network in an effort to infer the current state of the network. Of particular concern with these measurement approaches are the accuracy, precision and temporal stability of the measurements made. Measurements made within the network by the network operator may be improved by dedicated instrumentation and knowledge of the network infrastructure. The EPAC systems trade some potential measurement performance degradation for independence from the network operator. Where earlier work had largely focused on the network operator as providing the AC function, EPAC raises the option of a customer endpoint providing the function and this is reflected in Table 1 under the rubric of "Who makes the admission decision."

AC in a Mobile Context

Single service mobile systems also face the need to re-evaluate admission decisions as the user moves and the session is handed off across different infrastructure. Some mobile systems use a cell-occupancy approach to AC where the arrival and departure and handoff events are characterized on a per cell basis i.e., the system state model is constructed in terms of call arrivals, departures and handoffs into and out of a cell, but irrespective of location. Spatial mobility approaches to AC for mobile systems model the user location as it moves between cells (i.e., the system state model includes concepts of the infrastructure topology and the potential user trajectory). Spatial uniformity considers whether the system state models are homogeneous or heterogeneous. Jain & Knightly [15], proposed the taxonomy of mobile AC algorithms shown below:

- Cell Occupancy Allocation
  - –Uniform
  - –Non-Uniform
    - * Aggregate
    - * Per user
- Spatial Mobility
  - –Non-Uniform
    - * Aggregate
    - * Per user

The cell occupancy vs. spatial mobility is reflected in Table 1 as an example of topological scope — whether the mobile device is admitted to only one fixed infrastructure node, or has a claim on the resources of other nodes as well. The uniform/non-uniform and aggregate/per user dimensions are examples of granularity considerations associated with the system state/capacity model and its associated decision parameters in Table 2. While most of the AC work in mobility is related to handoffs between fixed infrastructure, there are other forms of mobility that could also have implications for an AC scheme (e.g., session handover between different terminals, nomadic access from different fixed locations).

AC vs. MAC

AC and Media Access Control (MAC) protocols appear to be different solutions to related problems of resolving conflicts in resource requests. Like AC mechanisms, MAC protocols are typically evaluated in terms of their fairness, efficiency, and simplicity in supporting decentralized implementations etc. MAC protocols may be classified as channelizing (where the MAC protocols recognize the existence of more than one isolated communications channels) or non-channelizing, where the latter may have either randomized (e.g., CDMA) or deterministic (e.g., polling) sequence preferences. These options are summarized in Fig. 1.

Consider the similarity in the concepts here with the notions of service isolation and sharing implied in the IP network service concepts of service isolation and sharing. The deterministic/randomized preferences are simply different mechanisms to support sharing of the resource. The service isolation vs. sharing concept is captured in Table 2 under the heading of the service definition- an aspect of the service model.

AC vs. Congestion Control

Some AC schemes can be considered a form of congestion control-usually congestion avoidance rather than congestion recovery. Control systems are typically classified as open-loop or closed-loop, where closed-loop control systems have some form of feedback loop providing a measurement of the cur-

rent system state and open-loop control systems do not. Open-loop control systems are simpler, but less sensitive to the dynamics of the system under control. In the closed-loop case, the feedback information can be an explicit message or implied in some other system behavior — e.g., the delay in an Acknowledgment. The feedback information can also be provided globally i.e., from a destination to the source, or locally. The feedback signal may also be persistent (i.e., available continuously) or responsive (i.e., available only in response to some event). MBAC and EPAC systems are closed-loop systems with measurements of network state. Yang *et al.* [16] proposed the taxonomy for congestion control schemes based on a control theory approach shown below:

- Open-Loop Control
  - –Source Control
  - –Destination Control
- Closed-Loop Control
  - –Implicit Feedback (e.g., through other actions)
  - –Explicit feedback (e.g., signaled)
    - * Persistent (global)
    - * Responsive
      - •Global
      - •Local

The open/closed loop algorithm choice is reflected in Table 2 as part of the system state model, and as an internal signaling mechanism for system state information (in the case of closed loop control). The source/destination control is an example of "who makes the admission decision" in Table 1 — where source/destination are types of customers of the network and generally referred to as calling vs. called parties in PSTN nomenclature. The implicit/explicit nature of feedback associated with closed loop control is reflected in Table 2 as options for both an indication of the decision result and as well as options for decision input parameters such as system state models and traffic source models. The persistence of explicit feedback is an example of a time continuous "system state/capacity model" as identified in Table 2.

Accounting and AC

AC functions are part of the service definition at the interface between the user and the network. For commercial networks, AC decisions are typically treated as billable events for external users. Because billing is a use of AC decisions, the accounting model places some requirements on the structure of a viable AC scheme. In addition, the accounting model is relevant to the AC scheme because:

- Pricing mechanisms based on accounting data may involve the user in the AC decision (e.g., Faulkner *et al.* [17])
- Internal resource allocation between services may drive cost allocations
- Accounting information relevant to cost allocation and pricing models may be used in AC decisions.

Accounting policies may be considered as closed (fixed) or open (able to be extended). The accounting taxonomy of Kouadio & Pooch [18] provides a perspective on some of the information flows relevant to admission decisions:

- Interactivity
  - –User input
  - –Provider centric
- Network protocols
  - –Diffserv (no reservation)
    - * Unicast
    - * Multicast
  - –IntServ (reservation)
    - * Unicast
    - * Multicast
- Policy scope
  - –Closed policy
  - –Open policy
- Data collection sources
  - –Core network
  - –Network edge
- Sender side
- Receiver side
- Data Collection Methods
  - –Precise
  - –Approximate
- Collected Data Elements
  - –Duration
  - –Volume
  - –Node Resources (e.g., QoS class, buffer space, media adaptation, conference bridge ports)
  - –Distance

The accounting data provides some reflection of the actual usage of network resources. Some charging models may reflect additional aspects of resource usage such as time of day or network status. A pricing model associates prices with that resource usage. In a multi-service IP network, a pricing model provides a means for comparison between services. The accounting data collected in Table 1 is an example of an admission decision effect, (although it may also be of interest for studies of network and service performance) and includes dimensions such as the collected data elements, the collection methods and the data sources. The interactivity dimension is an aspect of "who makes the admission decision" — the customer or the network in Table 1. The network protocols options are not recorded directly in Table 1, but specific protocols do provide constraints on the admission decision input parameters in Table 2. The accounting policy scope can also be considered as reflected in the admission decision scope in Table 1. The generation of accounting data is also noted in Table 1 as a potential rationale for AC when required as part of the application or service definition.

Network Context for AC Decisions

From the ontology developed in Table 1 and Table 2, it can be seen that there is a very broad range of considerations involved in AC schemes. Evaluation of AC schemes, whether for design or deployment should also consider the network context and service objectives that provide the commercial rationale to justify the realization of the AC functionality, and which therefore provide some perspective to prioritize the relative importance of the dimensions identified in Table 1.

Although Table 1 provides a broad perspective on AC, AC is often associated with the topic of resource management and more particularly with notions of congestion avoidance and capacity management. For large-scale networks, the capacity planning and network upgrade process provides an operational and service context for AC. Protocol specific studies (See e.g., 6], 7]) of AC in WAN protocols such as ATM provide some guidance, but it is not always easy to extract the AC principles from the protocol specifics. I provide an overview of the linkages between capacity planning and AC. Congestion avoidance is typically considered over much shorter timescales than capacity management and is fairly well understood within the context of the admission decision input

[1] *The ITU-T has a standardization effort on AC as part of its QoS work under the Next Generation Network Standardization initiative, see e.g., http://www.itu.int/ITU-T/ngn/index.phtml*

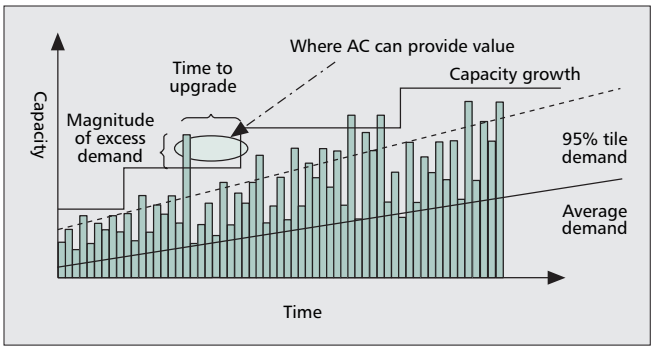

Figure 2. *Capacity planning and AC.*

parameters and admission decision performance dimensions of Table 2.

To further explore the network context of AC, example network architectures drawn from industry standards are introduced, and some options for the application of AC in these examples are then considered.

Many studies on AC focus on AC in the context of a single link or a simple LAN environment (see e.g., [19, 20]). A number of issues arrive in scaling services from LAN to WAN, and some of these impact AC. AC schemes based on resource management must have mechanisms to acquire knowledge of the available resources to be managed. I provide an overview of these practical scaling concerns, and how they might impact the AC considerations.

Capacity Planning and AC

AC is one element of network functionality to be considered in the network capacity planning process. From this perspective, AC is primarily concerned with avoiding the case of network overload. From this perspective, AC can be used to maintain acceptable service levels while deferring upgrades — essentially permitting the network to operate at a capacity below that required for satisfying peak demand. This may be of particular interest to the operator in cases where the time to upgrade the infrastructure is significant, or where the nature of the service is particularly "bursty" — e.g., due to a transient focused overload event.

Figure 2 illustrates this region where the AC scheme has an operational value in protecting against overloads. Operating at a capacity level of average demand would result in denial of service for 50 percent of demand. Providing capacity for peak demand results in excess (underutilized) capacity. Network capacity planning procedures typically target some high percentile of demand to provide a commercial service. Capacity planning has resource modeling issues similar to AC (refer to section 0) albeit on a longer timescale. As a practical matter, the provision of capacity is often only possible in relatively large quanta, resulting in a discrete capacity profile over time.

As an operational consideration in capacity planning, AC is one element in the trade-off between the cost savings associated with the deferral of capacity installation and the risk of demand exceeding capacity. The service is assumed to be such that there is benefit from denying excess demand — generally to avoid service degradation to either incremental demand or to the aggregate demand. Figure 3 illustrates this role of AC in the selection of the operating point for the network. Where the demand is very bursty (i.e., the CDF of demand has a long tail), AC provides a mechanism to support satisfactory services with capacity at levels significantly lower than peak demand.

There are costs associated with the provision of an AC function in the network. There are also costs associated with

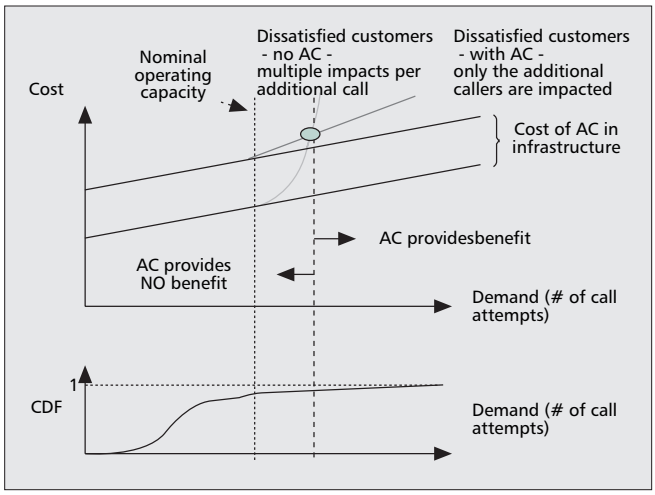

Figure 3. *Cost of AC.*

customer dissatisfaction. With an AC scheme, the number of dissatisfied customers rises linearly above the operating capacity limit as those customers requesting additional network services are denied access. In the absence of an AC scheme, excess demand above the operating capacity limit results in degraded service for the aggregate demand i.e., all customers suffer degraded service.

These costs and benefits of AC are be summarized financially in a business case similar to that shown in Eq. 1 where R_{new} is any new revenue enabled by the AC function, $S_{Capital}$ and $S_{Operations}$ are the cost savings from capital deferments and operations efficiencies, C_{AC} is the cost of the AC function and T is the threshold for the minimum acceptable return from the investment in AC.

$$R_{new} + S_{Capital} + S_{Operations} - C_{AC} > T \quad (1)$$

If we focus on the potential savings from deferring investment (as the other factors depend on other actions well beyond AC) this can be reduced to Eq. 2

$$S_{Capital} - C_{AC} > T \quad (2)$$

For a deferred capital investment, $S_{Capital}$ is the interest on the capital over the period of deferment. Equation 3 captures this for the case of simple interest (i) over a deferment period (t).

$$S_{Capital} = i \cdot t \cdot \text{Capital} \quad (3)$$

Combining Eq. 2 and Eq. 3 and rearranging, we can generate Eq. 4 which provides some insight into a constraint on the ratio of the capital cost of the admission control function to the capital cost of the network equipment being deferred.

$$i \cdot t - \frac{T}{Capital} > \frac{C_{AC}}{Capital} \quad (4)$$

While commercial network operators would likely use more sophisticated financial models, this simple approach is useful to illustrate some of the sensitivities. For most enterprises, the interest rate (i) and the investment hurdle rate (T/$Capital$) are determined by external factors. A deferment-structured business case for AC would then become more viable if the deferment period for the network capital investment is longer. High growth networks would be less likely to support a longer deferment and so would find AC to be less viable. Increased complexity (and therefore cost) in the admission control function would require an ability to defer the investment longer. Network growth, however, is unlikely to be constant; an S-shaped logistic curve a more typical representation of rapid initial deployment eventually constrained as deployment saturates.

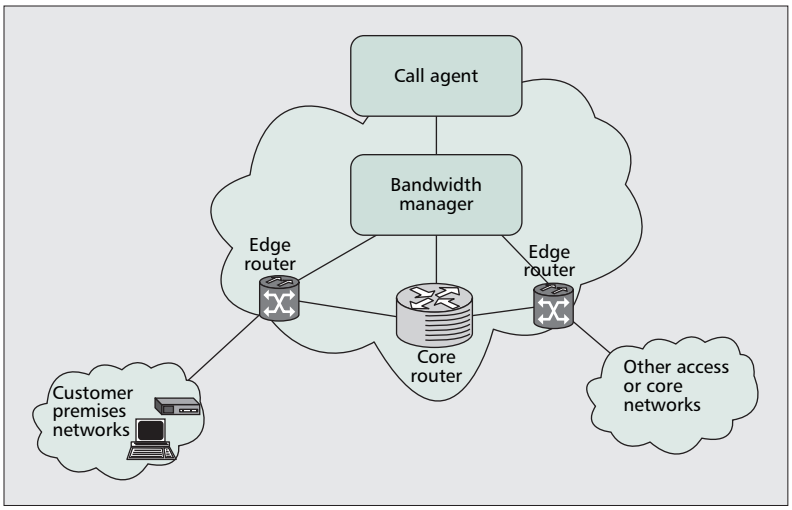

Figure 4. *Bandwidth manager.*

One might assume that postulates (e.g., Reed's "law," Metcalf's "law") of increased value from network size effects might facilitate significant R_{new} to support the deployment of additional functionality, but this is unlikely in typical commercial environments. If we consider a simple subscription based revenue model, then R_{new} only increases with $O(n)$. A usage based model might capture a slightly higher level of revenue and recent work [21] has proposed $O(n \log(n))$ as the appropriate growth curve for the value increase associated with network effects.

While this brief analysis indicates that there may be some ways to value AC from a capacity planning perspective, this is not the only relevant perspective and AC may be required for other purposes as described elsewhere in this article. In particular, it may be required as an inherent aspect of a service definition, or for ensuring service isolation between services in a QoS enabled network.

Example WAN Network Architecture

In WAN networks it is common practice to distinguish between core and access networks, but an end-end AC scheme must be able to operate across both environments. Access networks are typically required to support the full range of the service mix (e.g., both elastic and inelastic traffic types — refer section IV. A) whereas the service may be groomed onto different (isolated) infrastructures in the core.

Example Core Network Architecture — Several emerging industry architectures[1] (e.g., [22]) propose a bandwidth management function to perform the Admission Control decision, and commercial products are emerging to match these architectures (see e.g., [23]). A SIP session in this architecture may be negotiated between the Customer Premise Equipment and the call agent or SIP proxy. The call agent relies on the bandwidth manager to perform the resource management function of deciding whether sufficient capacity exists. The decision may be communicated up to the call agent to respond to the user, and/or be pushed down to the network elements as a policy change.

This creates several challenges, as, for example, the AC scheme must be aware of the resources available in the network (including resources reserved for failure protection). The set of network resources is not entirely static as normal operational actions will remove capacity temporarily for maintenance and add new capacity for growth. Any AC scheme considering the distributed nature of multi-operator Next Generation Networks (NGNs), will require careful design to avoid issues such as deadlock/livelock [24] due to the interactions of the different entities and also considering various fault scenarios that will occur in large scale networks.

Consider the capacity planning/AC linkage for core networks, and assume a core network is designed to be nonblocking for a uniformly distributed traffic load. With the traffic load at the designed capacity, if the traffic routing deviates from a uniform distribution, then "hot spots" or localized overloads may occur. If traffic from multiple sources is focused on a single egress point, that load may exceed the capacity of that egress link. This topology/routing interaction with AC is discussed further, and becomes part of the system model entry in Table 2. Given the scale of public network cores, the scalability and performance of a bandwidth manager is an obvious concern both in terms of the degree of distribution and the latency associated with the signaling response to the user.

Example Access Network Architecture — The TR-59 network architecture [25] is introduced here (refer to Fig. 5) as a baseline, in order to provide a more concrete context for further discussion of classification mechanisms for potential AC schemes in the context of an access or aggregation network.

This is an industry standard access network architecture through which providers may enable IP QoS based services [26] between residential subscribers and Network Service Providers (NSPs) and Application Service Providers (ASPs). Traffic from the subscriber passes through a Residential Gateway (RG) at the customers premises and onto a DSL line where it terminates in a Digital Subscriber Line Access Module (DSLAM) or Remote Terminal (RT). The DSLAMs aggregate traffic from multiple subscribers into an Ethernet or ATM network. The Ethernet or ATM network aggregates the traffic from multiple DSLAMs into a Broadband Remote

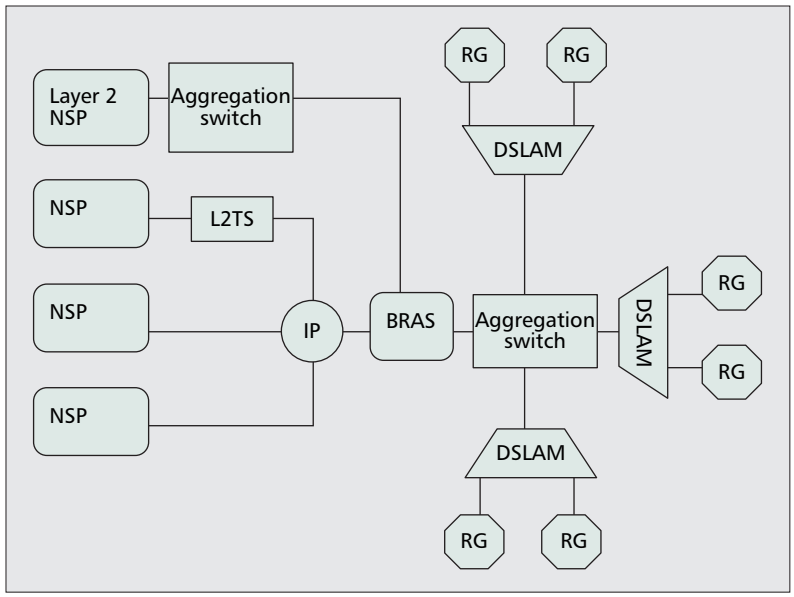

Figure 5. *Example access network architecture.*

Access Server (BRAS) that typically terminates the PPP session from the customer, before handing off the traffic to the appropriate service provider. Other access technologies (e.g., PON, DOCSIS, and Wireless) may differ in detail but typically provide similar asymmetric access bandwidth and perform a similar aggregation function where:

- The sum of the bandwidth available on all the customer facing interfaces typically exceeds the bandwidth available on the interfaces facing the core networks, and,
- Multiple layer 2 devices (e.g., Ethernet or ATM switches) may aggregate the traffic before entering the IP routers.

Access networks are also typically sensitive to the type of consumer; with residential access rates typically lower speed than commercial access rates, resulting in an inherent rate adaptation requirement. Consideration of AC functionality with respect to Fig. 5 leads to some interesting questions on how best to apply AC functionality, since the traffic assumptions of access networks are typically asymmetric and the traffic aggregation mechanisms often span multiple protocol layers (e.g., Layer 2 protocols such as Ethernet or ATM as well as Layer 3 protocols — typically IP), and include additional overheads (e.g., PPP or other tunneling protocols). These design constraints in terms of aggregation ratios, rate adaptation, protocol layering and protocol overheads should be considered within the system state/capacity model of Table 2.

Note that some services (e.g., VOIP) have notions of a service-specific enforcement point for some admission control functions — for example, in a so-called Session Border Controller (S/BC). While AC might be applied independently at a service layer and at a network layer in core networks, this is more problematic in access networks because of the cost implications of deploying additional complexity or equipment in high volumes. Some more recent proposals (e.g., [27]) suggest the integration of these S/BC functions into the access network elements — such as the DSLAM.

Applying AC in the WAN context

Consider the typical assumptions for small-scale networks and how scaling the network may impact AC schemes in these areas:

- *"Reliability is high"* — With a large-scale network, the sheer number of devices leads to the inevitability of failures. AC schemes need practical measures to ensure they fulfill their objectives despite failures in the network, i.e., the AC scheme should not fail because some other network element fails. Core networks typically have redundant interconnection paths (because a failure here impacts multiple users) so that the AC algorithms become concerned with the routing options available. Access networks are typically not redundant, so the AC scheme must consider what to do when resources are not available due to failures.

  Simply adding an additional functionality to the network increases the potential for failures within the network unless some remediation is applied. Where a logically or physically centralized bandwidth manager is used, some redundancy mechanism should be provided to recover from failures in the AC mechanisms. More distributed AC schemes may have more complex redundancy mechanisms. The reliability of an AC scheme is not simply a matter of whether it provides a decision in a timely fashion, but should also consider the validity of that decision. The frequency and impact of failures can impact the perceived quality of the service. Adding additional capacity (at some cost) may improve the reliability of the service as perceived by the end user- because the AC is less likely to reject a resource request, but the additional equipment may actually degrade the overall reliability of the network — resulting from the increased probability of some network element being in a failed state at any given instant because there are now more network elements. From an earlier section, it can be seen that capacity planning and AC are linked, and the perceived reliability of the service is one dimension that can be measured.
- *"Latency is low"* — Both propagation delays and queuing delays increase as the network scale increases. Many closed-loop AC schemes are dependent [28] on the round trip time of the control loop. Latency across core networks may be dominated by propagation delays, whereas latencies across access networks may be dominated by the transmission delays and queuing delays associated with the link speed. AC schemes that consider service latency may need to obtain this information differently for core and access networks. Where the AC scheme relies on signaling, the latencies associated with this can be a significant factor in the perception of responsiveness by the user. The latencies assumptions for signaling, measurement and data transfer can be considered part of the system state/capacity model in Table 2.
- *"Bandwidth is infinite"* — If true, this would only eliminate the congestion control objective for AC, but not the need for administrative controls on network usage. In many networks this is not true for at least some of the links, and so appropriate resource management controls are required. Common design practices typically permit this assumption across the core, but retain finite bandwidth assumptions at the edge. The interconnection of core networks also typically has finite bandwidth constraints. Where AC signaling is required, there should be an appropriate bandwidth allocation and classification defined. By eliminating the resource management argument, this scenario focuses attention on the remaining rationale (security) in Table 1. In the more general case bandwidth is not infinite and some capacity management regime is likely to be in force.
- *"The network is secure"* — AC is one mechanism that can prevent unauthorized traffic from propagating through the network. Where LAN scale network typically maintain physical security by restricting access, WAN scale networks often provide multiple access points that are made physically accessible to other potential consumers to enable rapid provisioning of new or additional service capacity. This creates "administrative" security constraints on which access points are permitted to send traffic to the network. Core networks are typically concerned with aggregate traffic patterns, and fine-grained security concerns may not be feasible for AC schemes here. Access networks have some potential for finer-grained security considerations; however, there are challenges (see e.g., [29]) in describing security constraints, and then recognizing and acting on them in large scale networks. In addition, WAN networks connect other (private) networks that are administered with varying degrees of sophistication and security awareness. A security breach in one of these private end networks may result in traffic that is unauthorized by that private network being sent into the WAN. Administrative constraints enforced by the AC system may provide some mitigation of this unauthorized traffic. A simple example from the PSTN would be denial of long distance or International call requests. Any AC signaling scheme should also be resilient in the presence of attacks against

the signaling infrastructure. While much of the literature discusses AC from a resource management perspective, the security rationale for AC listed in Table 1 is also an area of concern for network designers and operators.

- *"Topology doesn't change"* — If an AC scheme is managing resources at the scale of the whole network, then the AC scheme should consider the impact of these topology changes on new and existing resource commitments. Topology changes may happen automatically due to failures triggering rerouting operations (as mentioned above), or they may happen due to provisioning operations. The AC scheme must track the ongoing provisioning of new and changed resource allocations. While this occurs in core networks, the access can be much more dynamic in terms of the enabling or disabling service endpoints. Topology changes need to be dynamically reflected in the system state/capacity model of Table 1, and in some cases may be the result of capacity planning activities
- *"There is one administrator"* — While a LAN environment may have a single administrator, the Internet has no single administration, but rather each connected network is administered independently and may implement different AC policies and mechanisms. Both access and core networks have different administration than the end user's network. If an explicit signaling mechanism is used in the AC scheme, then a common signaling mechanism across multiple administrations should be advantageous as the multiple administrations must agree on admissibility. For AC based on resource management, this requires a common basis for parameterization of the traffic and resources. For AC based on administrative policy, this requires a common format for expressing that administrative policy. As an example traffic may be admitted to one network with a particular class of network service, but another network may not offer that class of service.

  EPAC schemes (see e.g., [14]) avoid the need to synchronize the AC decision across multiple network administrations by placing the administrative responsibility for AC at the endpoints and conceptually at the application layer, rather than in the network.

  This is captured in Table 1 as "who makes the admission decision. "
- *"Transport cost is zero"* — In a WAN environment, because of the involvement of a network operator in providing the network infrastructure, users of that network infrastructure are typically subject to charges in order for the network operators to recover the costs of that infrastructure (see e.g., [30]). AC may be required as part of that commercial transaction. When provided by commercial operators, both access and core networks require billing mechanisms for the services provided. The granularity of detail may be higher for access networks compared to core networks. The AC scheme may generate billing events and be a point of administrative control in the event of non-payment. The accounting data required to support billing is typically an admission decision effect as shown in Table 1, although price/cost structures may also be inputs considered in the system state/capacity model.
- *"The network is homogeneous"* — Most AC schemes require homogeneity of infrastructure components such as signaling in order to achieve end-to-end service objective. While a core network may be able to assume a homogeneous IP infrastructure, access networks typically have a mix of layer 2 and layer 3 devices that complicate the AC decision based on these mixed resources. AC schemes are emerging (see e.g., [31]) to consider admission control issues between networks that include elements that are fixed or mobile and secure of non-secure. AC in heterogeneous environments, and AC that is not end-to-end, are candidates for further study. The major impact of a heterogeneous network is in the potential variety of factors that must be considered in the admission decision input parameters of Table 1.

Other Categorization Approaches

The AC decision can be also be scoped by the scale of the decision in terms of the:

- Service mix supported
- Discard granularity
- Topology or sequence of resources affected
- Resource modeling
- Administrative policy

Service Definition, Context and Mix

The AC scheme will obviously be impacted by the ranges of services (the service mix) that it is required to support. This service differentiation adds complexity when the services have different dimensions (e.g., loss, delay) for service performance. An additional aspect of AC complexity results from the degree of elasticity that the applications using these network services have.

Elastic vs. Inelastic Applications —While defining a particular service to meet the requirements of a specific application is not trivial, it can be relatively straight forward if the application's QoS requirements are well specified. Many different applications can be supported by a given network service if the traffic characteristics of that application are comparable. Shenker *et al.* [32] provides a taxonomy for applications based on the network services required to support them, which is reflected in Fig. 6. This figure reflects the service model of the IETF int-serv framework. Each network service supported adds complexity and cost to the operation of the network infrastructure. The resource management aspects, in particular, must recognize the different resource requirements of the different service types. A service design objective, therefore, should be to minimize the number of services required, while still supporting the maximum number of applications.

Application Elasticity and AC — Application elasticity is the ability of an application to operate effectively (adapt) in the presence of network defects. Applications based on TCP would typically be considered elastic. In packet networks, defects (loss or delay) increase with load for all services as shown in Fig. 7a and Fig. 7b. Elasticity can be considered as providing a constant (or acceptable) performance despite variations in the level of network defects. Figure 7c illustrates this notion of a service providing a constant level of quality despite variations in network performance.

But is Fig. 7c realistic for known services? That depends on having a service definition of acceptable quality.

For voice services, service quality is typically measured using Mean Opinion Scores (MOS). Standardization efforts (e.g., from ITU-T see [33–35]) have developed the E-Model which provides empirical formulae and data constants enabling the prediction of MOS scores based on network parameters including loss and delay. These models provide a mechanism to assess the elasticity of VOIP applications in terms of loss and delay defects (see e.g., [36]). Even with the trade-offs described in this empirical relationship adequate voice quality

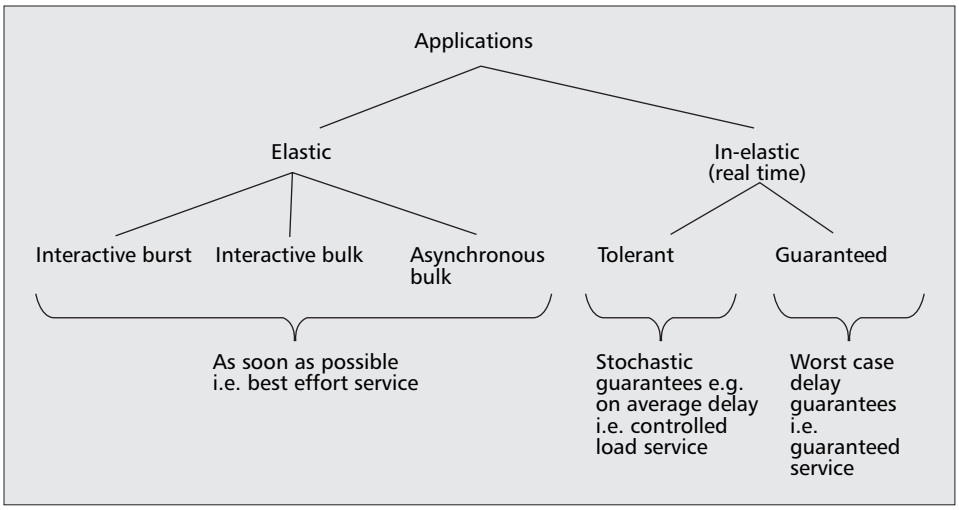

Figure 6. *Mapping application categories to service classes.*

performance cannot be achieved beyond some maximum delay or maximum loss.

For TCP performance, similar empirically derived formulae are available to describe effective TCP performance in terms of throughput (see e.g., [37]) or latency (see e.g., [38]). UDP does not provide intrinsic mechanisms for masking loss or delay, so any elasticity for UDP applications must be embedded in the application layer protocols.

So, it seems reasonable to assume that an application may have a region on the graph of Fig. 7c where an elastic relationship exists between loss and delay. For many applications, that region of elasticity is bounded. Delay beyond some critical amount renders interactive communications unusable. Loss beyond some critical amount is unable to be masked by the application.

Consider the interaction between these loss and delay network defects and the network resources (e.g., bandwidth) required to mask them. To achieve a constant application performance in the presence of increasing errors requires increasing bandwidth (Fig. 7d) to compensate for the errors — e.g., through retransmissions or forward error correction. For a given application quality level there is a minimum bandwidth required without loss, and there is a maximum loss that the application can withstand while delivering that same quality level. Service latency requirements may be achieved by allocating peak bandwidth in the network (i.e., no network buffering) or by allocating a minimum bandwidth in the network and effectively smoothing peak traffic demands (Fig. 7e). There are practical constraints on the range of bandwidth, loss and delay that real networks can accommodate. For a given application, there is a multidimensional region of elasticity relating loss, delay and network bandwidth requirements. For a given application quality level there is a minimum bandwidth required (without delay impacts), and there is a maximum delay that the application can withstand while delivering that same quality level.

Elastic applications can operate with the same application performance over a range of network conditions (Fig. 7f), but this adaptability has implications for AC. Should an AC reserve the maximum or minimum resources required for a service instance? If multiple instances are active in the system, the algorithms that control their adaptation to available network resources must be compatible for stable behavior. This range of elasticity provides for a trade-off in application robustness vs. network efficiency. If the application is elastic, the resource requirements in the traffic or system capacity model will either not be precise, or the AC scheme will become more complex due to the parameterization of the elasticity range (see e.g., [39]). I consider this lack of precision and other resource modeling issues.

Service Mixing and Interaction — The service interaction aspects can be summarized as follows. When two service instances (say A and B) have to compete for the same resources, the possible out comes are:

- A loses
- B loses
- Both lose

The specific interpretation of "losing" may vary depending on the service design. The loss may be in terms of packet loss or packet delay. Further refinements may include identifying specific portions of the service flow as candidate for the loss function (e.g., discriminating on in-profile vs. out of profile uses of a service). The selection of service interaction result is part of the service definition.

Traditional call AC mechanisms were developed for a single service (telephony) network. AC schemes for multi-service networks based on ATM technology and per-VC signaling have been extensively studied in the literature and have had some limited deployment. AC mechanisms in the context of IP network have been studied, but have received less widespread deployment, and there is certainly less agreement on the notion of IP service classes beyond best effort.

Integrated AC schemes provide a single AC algorithm that supports multiple services over a single set of resources. Consider a link AC algorithm that supports two different services that require different amounts of system resources (e.g., link bandwidth). If only voice sessions were in the system, a maximum of *Nvoice* sessions could be supported. Similarly, a maximum of *Nvideo* video sessions could be supported if they were the only service in the system. The system could be configured to support any mix of voice and video sessions that lies within the admissible region illustrated in Fig. 8. The boundary of the admissible region in this figure is shown as linear, but this is only an approximation of the behavior of many real systems. If the number of sessions of each class that can be supported is large, then a continuous function may be a reasonable approximation, but it may be better modeled as curved rather than linear. If there are only a relatively few sessions of any service class that can be supported, then boundary of the admissible region may be a discrete function (e.g., a step staircase). For multiple service classes, the boundary of the admissible regions becomes a surface in a multidimensional space.

The notion of an admissible region is not new and has been used in the literature with different link scheduling mechanisms (see e.g., [39–41]). While the description above implies a deterministic boundary, bounds could also be stochastic in nature — e.g., based on blocking probabilities of the different session types.

In a multi-service network, there will be regions of the network where different AC regimes may need to co-exist. In these regions, the expected behavior needs be properly defined so that AC schemes operating on different timescales or granularities do not interfere with each other's operations. An AC scheme is designed to manage the overload situation for a particular set of resources. With multiple AC schemes operating in parallel, care should be taken that the resources controlled are

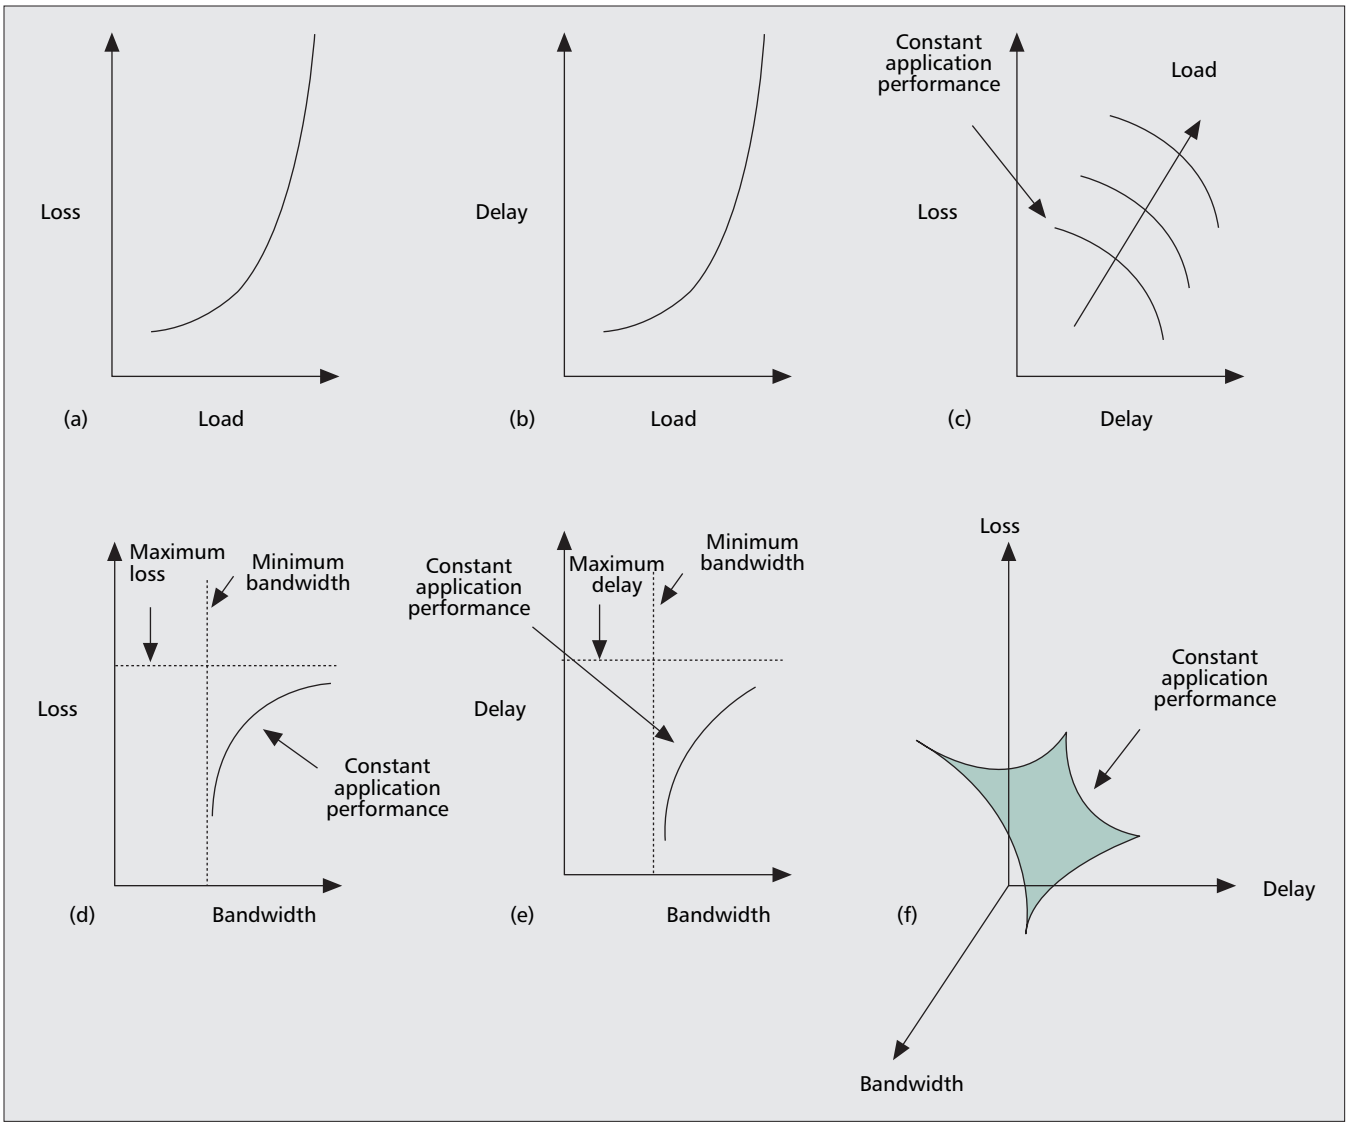

Figure 7. *Service elasticity and QoS dimensions.*

separated. One approach to this at the level of link bandwidth resources is to partition the bandwidth into regions — e.g., a session-based AC bandwidth region and a packet-based AC bandwidth region as illustrated by Fig. 9. This partition could be a "hard" with the overload control regime of each partition having exclusive use of the resources within each partition, or it could be a "soft" partition with "excess" traffic being permitted to use bandwidth of the other partitions when the other traffic class is not available. In terms of link schedulers, this presents a choice between non-work conserving and work-conserving link schedulers. A work-conserving link scheduler is only idle if there is no packet waiting to be sent, whereas a non-work-conserving link scheduler may be idle even if there are packets waiting to be sent. If the service concept permits "soft" resource partitions, (i.e., work-conserving scheduling) then the system behavior during those periods must be appropriately defined. For example, excess bursts may be remarked and discarded elsewhere in the network.

The service mix is considered part of the admission decision effect in Table 1 and part of the service model in Table 2.

AC Granularity

Connection AC (CAC) applies to connections (e.g., telephony calls), but connections can apply at different protocol layers (e.g., PPP connections or TCP sessions). AC for application layer sessions is used [42] for load control on web servers. A similar concept of flow-based AC [43] has been proposed for use in the Internet. Different QoS mechanisms become feasible at different timescales. ACs may be operating at one or more of these levels of traffic granularity. One approach is to consider the AC granularity options based on the entities that may be denied admission (i.e., discarded) such as:

- Packets
- Bursts
- Flows
- sessions (e.g., Calls)

The discarded data at packet or burst granularity is not associated with any session-level semantics; indeed it may be extracted from multiple different sessions. The flow granularity assumes that a single flow corresponds to a single session, but this is not always so. Flow based approaches that do not reflect session semantics suffer from similar problems to the packet and burst scale discard except that now the application must respond appropriately to the failure to establish a particular flow. Call or session-based schemes rely on some indication of call or session start/stop events. This could be explicit signaling of session start/stop, or, the network could implicitly infer this by, for example, snooping on the application level message exchanges [10] that synchronize the start/stop of the

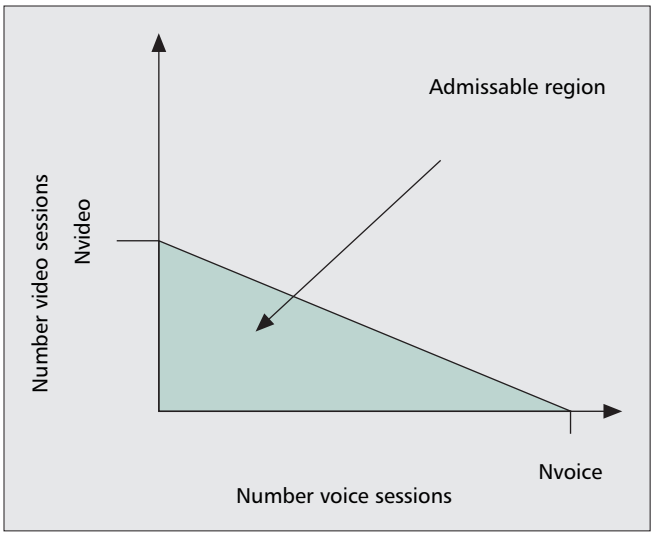

Figure 8. *Service mix trade-off in AC with a shared resource.*

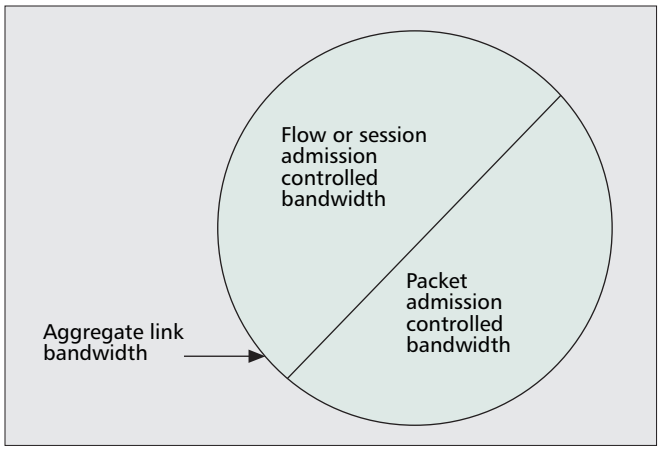

Figure 9. *Partitioning link bandwidth between AC schemes.*

application end points. In order to snoop, the AC needs to understand the specific protocols involved. For infrastructure deployment, this would only be feasible for widely deployed and well understood application layer protocols.

In the context of the access architecture, different discard granularities may be appropriate for different service concepts. The TR-59 NSP PPP session is essentially a permanent connection from subscription time. While packet and burst admission mechanisms can be applied, these are fairly coarse mechanisms. A diffserv-based approach provides for packet or burst granularity AC discard mechanisms that can be applied at the BRAS and RG. Network services that benefit from a flow or session based AC granularity are candidates for further study.

The granularity of the AC scheme is considered an aspect of the admission decision effect in Table 1.

Topological Scope of AC Decisions

The scope of an admission decision can be classified on a basis of topological scale. Admission decisions may be made with respect to:

- A particular point within a single administrative system (e.g., a link)
- An entire administrative system
- Multiple instances of the same type of administrative systems (e.g., multiple networks)
- Multiple types of systems (e.g., hosts and networks)

When considering partitioning of AC by topological scope, consideration should be given to the mechanism for partitioning and re-assembling the decision into subcomponents where operational simplicity may be traded for computational accuracy [44]. This can be complicated by the need to consider the effect of additional options created by the existence of multiple routes. For this reason, some signaling and routing protocols support AC (e.g., GCAC in ATM's PNNI protocol, RSVP in IP networks).

Nodes typically provide AC for internal resources of the node and the egress link. Consider the DSL Access network of Fig. 5. At the link level, the AC problem changes as you move across the network. The link characteristics change from relatively low asymmetric bandwidths at the access line to symmetric higher bandwidths at the BRAS interface. The traffic characteristics remain asymmetric. The service mix that the AC must support also changes from a single (Customer Premises Network) user at the DSL line to an aggregate at the BRAS interface. Statistical approaches that work with large numbers of flows may not be valid when the number of flows is small.

Network Admission Control (NAC) algorithms are sensitive to the underlying topology (network size, node degree, network structure e.g., hierarchy), the traffic matrix used, routing model and network resilience requirements.

Menth *et al.* [45, 46] identified and studied (in an Internet backbone context) four approaches to NAC based on:

- Link Budgets (link by link — similar to ATM AC)
- Ingress Budgets & Egress Budgets
- Border to Border Budgets (based on virtual tunnels)
- Ingress Link Budget and Egress Link Budget

The architectural framework of Fig. 5 identifies several networks that could be used to topologically scope an AC decision. A TR-59 related AC decision could be restricted to be within the networks or sub-components (e.g., Links) of the Customer Premises Network, Access Network, Regional Broadband Network, NSP Network, or ASP Network. These networks are not simply different administrative entities, but, in some cases, different types of systems. A home network and an ASP contain end-systems (hosts) participating in the session, while the other networks do not. At this time, AC mechanisms are not prevalent in any of these networks, and it is unlikely that all parties could move simultaneously to adopt a single AC scheme. If an AC scheme is desired, some sort of phased introduction seems more practical. A simple approach may be to assume that the access network is the only relevant bottleneck in the overall service design, and start there, postponing the problem of aggregating an end-end admission decision for a service. In order to justify deployment of any new service, it must provide significant value for its customers. The initial rollout to the access network alone must provide sufficient service value to justify deployment of the AC feature.

The topology considerations are considered as part of the topological scope in Table 1 as well as part of the system state/capacity model in Table 2.

Resource Modeling Issues and AC

If the models of system capacity or traffic description lack precision and accuracy, this will limit the effectiveness and efficiency of an AC scheme based on these inputs.

Consider an application that can adapt with equal quality results over a bandwidth range of 10 percent of its nominal requirements, and an AC scheme based on equivalent bandwidth decision criteria. If we assume network resources are reserved based on more than the minimum resources required, this provides a margin for transient network defects (e.g., loss events). If there are a large number of service instances, the

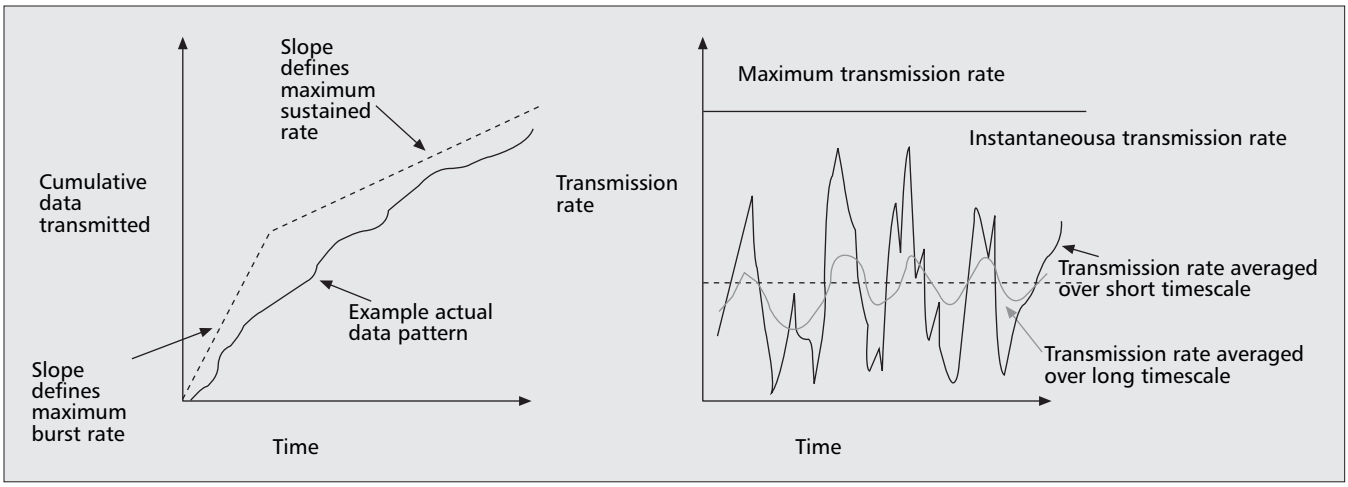

Figure 10. *Source traffic models.*

cumulative network resources allocated to application resiliency can significantly reduce network efficiency.

For every 10 instances of a service accepted with a 10 percent margin, another service instance could have been accepted if they all operated with no such margin.

The following subsections provide some additional discussion of some of the potential approaches to resource modeling and help illustrate why there may be limits to the precision and accuracy of the models that an AC scheme can be based on.

Models of System Capacity — Models of system capacity typically assume a fixed capacity for the system e.g., a fixes bandwidth for a link, although there has been some theoretical work on variable capacity systems. Link-level AC has been studied at length within the theory of queuing and scheduling systems. Typical approaches rely on converting source traffic descriptions (whether inferred through measurements or explicitly supplied via signaling) into some sort of "equivalent bandwidth" (see e.g., [47]). The notion of equivalent bandwidth is useful as a common denominator to rationalize different service classes or QoS grades into a single dimension for the AC decision.

Given a specific capacity constraint, there may be many solutions involving different mixes of traffic. Multidimensional capacity models become very complex and difficult to solve. Blocking probabilities and link capacity calculations in traditional telephony are typically modeled through to a coupled set of nonlinear equations. The Erlang Fixed-Point equation provides such a model (see [48]). Fixed point and behavioral network models (e.g., [49]) may also be used in some packet network contexts if the network is in a steady state and the properties can be assumed independent between network links.

One of the major service classes of interest, however, is that of real-time traffic. In this service class, the primary system requirement is for a delay bound rather than a bandwidth bound. Given information on available buffer space, and other network constraints, and traffic burstiness, it is possible to convert a delay bound into a bandwidth bound.

Aggregating results to larger scale systems in a scalable way is more problematic. For a bandwidth bound, the network scale admission problem is one of assuring sufficient bandwidth at each link along the path. For a delay-bounded service, the aggregate end-to-end delay is of interest. The most common approach is to pin traffic to a particular route and then accumulate the link level results in order to achieve a network level aggregate result. A typical approach is to calculate the worst-case delay at each node and then aggregate these metrics across the network. Such approaches are extremely conservative (see e.g., [44, 50]) resulting in low network utilizations.

End-point admission schemes (e.g., [14]) treat the network as a "black box" and measure its performance before making a decision. Most of these schemes attempt to model the network by a measurement of the bandwidth of the bottleneck link along the path of a particular connection. Such measurement systems are sensitive to assumptions about session durations and inherently involve assumptions regarding the statistical characterization of the traffic flows in the network. While such assumptions may be valid in the context of the applications and services a single user's customer premise network utilizes, they may not be valid when considering the aggregate of all network services supported by the access network. Sustained flows with variable intensity sharing the same queues would appear to be problematic for such endpoint admission regimes as the actual network congestion state may change after the measurement was made.

Models for Traffic Descriptions — Traffic descriptions can be applied to single sources or aggregates [51] of traffic. Analytic models of traffic sources include such things as rate limited sources, or leaky bucket limited sources. These models have the advantage of simplicity and analytic tractability; however, they are typically very poor representations of actual traffic demands. Usually the source model provides one or more measures of bandwidth (e.g., peak bandwidth, average bandwidth etc.). These source model parameters may be used directly, or to derive implied estimators for the AC calculations (see e.g., [52]). Figure 10a illustrates the discrepancy between a typical actual data arrival pattern and an analytic traffic model based on piecewise linear rate limits (e.g., a leaky bucket).

Statistical models of traffic sources are also used (see e.g., [53]), typically based on well-understood random variable distributions e.g., Poisson distributed arrivals for human triggered voice calls. While statistical models of traffic tend to be more representative of real traffic, they are less tractable. The results of admission calculations with statistical parameters are also probabilistic.

The problem of measuring real traffic sources and describing their behavior in a limited set of parameters has proved very difficult. Figure 10b illustrates this effect by showing the comparison between an instantaneous transmission rate, a maximum transmission rate (e.g., a line rate) and average transmission rates calculated over various time periods. Simple analytic source models can identify worst-case limits, but these tend to be so extreme as to be uneconomic in commer-

cial networks. More complex statistical models have been studied for particular types of data sources. For example, Seeling *et al.* [54] provides an overview of the complexity involved in modeling video traffic.

Markov models have traditionally been used for well-characterized traffic streams e.g., voice telephony. Statistical characterization of actual internet traffic has proved more difficult. Several authors have reported that actual internet traffic flows exhibit Self similarity (see e.g., [55]) or Long Range Dependency — the tendency for traffic to have infinite variance and heavy tailed distributions. Consensus on the reasons for self-similarity in traffic patterns appears not yet to have been reached, although some identify a linkage for http traffic to self-similarity in the distribution of file sizes accessible via http and Long Range Dependence effects in user think time.

Specific protocols have end-end control schemes that are complex to characterize with simple statistical distributions. TCP has been analyzed and models are available for the end-end performance of a session, but not a statistical distribution of a TCP source. Statistically characterized source models are assumed to be independent of network state, but TCP behavior is coupled to the network state through its measurement of network congestion. UDP traffic however does not have such a built in control mechanism, although there may be similar functionality at a higher layer. The existence of long-range dependence or end-end control protocols in flows impacts our ability to describe the statistics of the aggregate traffic from particular users.

Administrative Policy Aspects of AC

Another approach to categorization of AC decision schemes is to consider administrative policy. Perhaps the stereotypical example of this approach is based on security considerations (e.g., [28, 56]). In this paradigm, the AC function is resolving a question of whether this instance of access to the network service is authorized.

Another aspect of administrative policy may reflect notions of priority or survivability (e.g., [57]) in the traffic. This is particularly relevant when pre-emption of existing admitted flows is permissible. In a typical scenario, the AC scheme must respond to a decrease in available resources (e.g., due to some sort network failure) by deciding which sessions may be permitted to continue.

Conclusions

The existing ontology related to AC from the literature was reviewed earlier. The value proposition of AC features within the network was described in general terms previously along with some discussion of specific industry based IP network architectures in which AC is under consideration. The linkage between AC (for resource management) and capacity planning is also illustrated, providing support for AC as a mechanism to avoid the effects of transient demand overloads. I provided some consideration of other aspects of the AC decision, including the service mix, granularity, topological considerations, and resource management.

Table 1 is intended to summarize the framework of the discussion on the dimensions for consideration in evaluation of AC schemes from the previous sections. This table provides a summary of different perspectives from which to compare features of AC schemes — it provides some answers to the "Who? What? Where? When? Why?" questions that help position the relevance of an AC scheme to its network context. While AC is well established for circuit based network technologies, it is not currently widely deployed in IP networks. As those networks mature to multi-service infrastructures, AC is expected to have a role to play, but that role is not yet fully defined. Most of the existing AC research has been focused on AC algorithms for resource management rather than other objectives of the AC function (e.g., security/application integrity). If the network does not support AC functions, then the application layer may need to develop them as required. AC functions at the application layer to resolve service interactions require further study.

Table 2 provides some initial insight into how various AC schemes operate. Aspects such as service models, decision parameters, mechanisms such as signaling and the performance of the AC scheme are considered. The purpose here has not been to provide a treatise on resource management, but rather to provide a tutorial introduction to the breadth of issues related to AC. As IP networks evolve beyond "best effort" services into a QoS-enabled, multi-service infrastructure, the rationale for AC also increases with new services that can benefit from AC being developed and the existing services that already rely on AC (e.g., telephony) needing to interwork in some way, with the multi-service IP infrastructure. From this perspective, resource management may be one rationale for admission control functions, but it is by no means the only purpose for the function, and research in various aspects of admission control is expected to continue.

Acknowledgments

I would like to thank my colleagues at BellSouth for useful discussions in refining these concepts, especially Andrew Vernon. The comments of the anonymous reviewers have also helped considerably in strengthening the article.

References

[1] V. Firoiu *et al.*, "Theories and Models for Internet Quality of Service," *Proc. IEEE*, May 2002.

[2] M. Eder, H. Chaskar, and S. Nag, "Considerations from the Service Management research Group (SMRG) on Quality of Service (QoS) in the IP Network," RFC 3387, Sept. 2002.

[3] V. Firoiu and D. Towsley, "Call Admission and Resource Reservation for Multicast Sessions," "Call Admission and Resource Reservation for Multicast Sessions," *Proc. IEEE INFOCOM'96*.

[4] W. Jia, W. Tu, and L. Lin, "Efficient Distributed Admission Control for Anycast Flows," *Proc. ICCNMC-03*, Beijing 2003.

[5] H. Perros and K. Elsayed, "Call Admission Control Schemes: A Review," *IEEE Commun. Mag.*, Nov. 1996, pp. 82–91.

[6] K. Elsayed and H. Perros, "Comparative Performance Analysis of Call Admission Control Schemes in ATM Networks," *Performance Evaluation and Application of ATM Networks*, Edited by D. Kouvatsous, Kluwer Academic Publishers, pp. 113–40, 2000.

[7] K. Shiomoto, N. Yamanaka, and T. Takahashi, "Overview of Measurement-Based Connection Admission Control Methods in ATM networks," *IEEE Commun. Surveys, 1Q*, 1999.

[8] A. Baiocchi, A. DeVendictis, and A. Monticelli, "Simple Models and their limits for TCP/IP network Analysis and Dimensioning" *Proc. IEEE Int'l. Conf. Communications (ICC) 2002*.

[9] A. Kumar *et al.*, "Non-Intrusive TCP Connection Admission Control for Bandwidth Management of an Internet Access Link," *IEEE Commun. Mag.*, vol. 38, no. 5, May 2000, pp. 160–67.

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

[11] R. Yavatkar, D. Pendarakis, and R. Guerin, "A Framework for Policy Based Admission Control," RFC2753, Jan. 2000.

[12] H. van den Berg and M. Mandjes, "Admission Control in Integrated Networks: Overview and Evaluation," *Proc. 8th Int'l. Conf. Telecommun. Systems*, Nashville, 2000, pp. 132–51.

[13] S. Nalatwad and M. Devetsikiotis, "Self-Sizing Networks:

Local vs. Global Control," *Proc. IEEE Intl. Conf. Communications (ICC) 2004*.
[14] L. Breslau *et al.*, "Endpoint Admission Control Architectural Issues and Performance," *Proc. ACM SIGCOMM 2000*.
[15] R. Jain and E. Knightly, "A Framework for Design and Evaluation of Admission Control Algorithms in Multi-Service Mobile Networks," *Proc. IEEE INFOCOM '99*, New York, NY, Mar. 1999.
[16] C. Yang and A. Reddy. "A Taxonomy for Congestion Control Algorithms in Packet Switching Networks," *IEEE Network Mag.*, vol. 9, no. 5, July/Aug. 1995.
[17] M. Falkner, M. Devetsikiotis, and I. Lambadaris, "An Overview of Pricing Concepts for Broadband IP Networks," *IEEE Commun. Surveys, 2Q*, 2000, pp. 2-–13.
[18] M. Kouadio and U. Pooch, "A Taxonomy and Design Considerations for Internet Accounting," *ACM Sigcomm Computer Commun. Review*, vol. 32, no. 5, Nov. 2002, pp 39–48.
[19] D. Gao, J. Cai, and K.–N. Ngan, "Admission Control in IEEE 802. 11e Wireless LANs," *IEEE Network*, vol. 19, no. 4, July–Aug. 2005, pp. 6–13.
[20] R. Yavatkar *et al.*, "SBM (Subnet Bandwidth Manager): A Protocol for RSVP-based Admission Control over IEEE 802-style Networks," RFC 2814, May 2000.
[21] B. Briscoe, A. Odlyzko, and B. Tilly, "Metcalf's Law is Wrong," *IEEE Spectrum*, July 2006, pp. 34–39.
[22] C. Gallon and O. Schelen, "Bandwidth Management in Next Generation Packet Networks," MSF-TR-ARCH-005-FINAL, Multi-Service Forum, Aug. 2005.
[23] U. Bodin, O. Schelen, and C. Vemmervik, "End-to-End QoS Control Architectures in a Wholesale and retail Perspective: Benefits and Challenges" *Proc. 10th Int'l. Conf. Intelligence in service delivery Networks, ICIN2006*, Bordeaux France, May 29–June 1, 2006.
[24] T. Anderson *et al.*, "On the Mechanisms for Real-Time Application Driven Resource Management in Next Generation Networks," *Proc. 10th Int'l. Conf. Intelligence in Service Delivery Networks*, ICIN2006, Bordeaux France, May 29-June 1, 2006.
[25] DSL Forum, "DSL Evolution — Architecture Requirements for the Supports of QoS-Enabled IP Services," TR-59 Rev 1, Sept. 2003.
[26] S. Wright and T. Anschutz, "QoS Requirements in DSL Networks," *Proc. IEEE Globecom 2003*.
[27] B. DeVos, F. Fredericx, and W. van Leekwijck, "Benefits of a Distributed Border Gateway in the Access," *Proc. 10th Int'l. Conf. Intelligence in Service Delivery Networks*, *ICIN2006*, Bordeaux France, May 29–June 1, 2006.
[28] M. Grossglauser and D. Tse, "A Framework for Robust Measurement-Based Admission Control," *IEEE/ACM Trans. Net.*, vol. 7, no. 3, June 1999.
[29] R. Sandhu and P. Samarati, "Access Control: Principles and Practice," *IEEE Computer*, Sept 1994, pp. 40–48.
[30] J. Roberts, "Internet Traffic, QoS, and Pricing," *Proc. IEEE*, vol. 92, no. 9, Sept. 2004.
[31] N. Saxena, G. Tsudik, and J. H. Yi, "Efficient Node Admission for Short-Lived Mobile Ad hoc Networks," *Proc. 13th IEEE Int'l. Conf. Network Protocols (ICNP'05)*, 6–9 Nov. 2005, pp. 269–78.
[32] S. Shenker, L. Zhang, and D. Clark. "A Scheduling Service Model and a Scheduling Architecture for an Integrated Services Packet Networks," Xerox Parc Tech. Report, Available via anonymous ftp from ftp://ftp.parc.xerox.com/pub/archfin.ps, 1993.
[33] ITU-T Recommendation G. 107, "The E-model, a Computational Model for Use in Transmission Planning," Dec. 1998.
[34] ITU-T Recommendation G. 108, "Application of the E-model: A Planning Guide," Sept. 1998.
[35] ITU-T Recommendation G. 113, "Transmission Impairments due to speech processing," Feb. 2001
[36] M. Gardner, V. Frost, and D. Petr, "Using Optimization to Achieve Efficient Quality of Service in Voice over IP Networks," *Proc. IPCCC 2003-The 22nd Int'l. Performance, Computing, and Commun. Conf.*, Phoenix, Apr. 2003.
[37] M. Mathis *et al.*, "The Macroscopic Behaviour of the TCP Congestion Avoidance Algorithm," *Computer Commun. Review*, vol. 27, no. 3, July 1997.
[38] N. Cardwell, S. Savage, and T. Anderson, "Modeling TCP Latency," *Proc. IEEE Infocom 2000*.
[39] G. Butazzo, G. Lipari, and L. Abeni, "Elastic Task Model for Adaptive Rate Control," *Proc. IEEE Real-Time Systems Symp.*, 1998, pp. 286–95.
[40] A. Elwalid and D. Mitra, "Analysis, Approximation and Admission Control of Multi-service Multiplexing System with Priorities," *Proc. IEEE Infocom*, 1995, pp. 463–72.
[41] K. Kumaran *et al.*, "Novel Techniques for the Design and Control of Generalized Processor Sharing Schedulers for Multiple QoS Classes," *Proc. IEEE Infocom 2000*.
[42] J. Carlstrom and R. Rom, "Application Aware Admission Control and Scheduling in Web Servers," *Proc. IEEE Infocom 2002*.
[43] J. Roberts, "Traffic Theory and the Internet," *IEEE Commun. Mag.*, Jan 2001 pp. 94–99.
[44] S. Wright and Y. Viniotis, "ATM Network Procedures Supporting Delay QoS," *Proc. IEEE/IEICE ATM Wksp. '99*, Kochi City, Kochi, Japan, May 24–27, 1999.
[45] M. Menth, S. Kopf, and J. Charzinski, "Impact of Network Topology on the Performance of Budget Based Network Admission Control Methods," *Proc. MIPS 2003*, Napoli, Italy, Nov. 2003.
[46] M. Menth, S. Kopf, and J. Milbrandt, "A Performance Evaluation Framework for Network Admission Control Methods," *Proc. IEEE Network Operations and Management Symp.*, Seoul, Apr. 2004.
[47] Z. Dziong, M. Juda, and L. Mason, "A Framework for Bandwidth Management in ATM Networks — Aggregate Equivalent Bandwidth Estimation Approach," *IEEE/ACM Trans. Net.*, vol. 5, no. 1, Feb. 1997.
[48] A. Girard, *Routing and Dimensioning in Circuit Switched Networks*, Addison-Wesley, 1990, Chapter 4.
[49] R. J. Gibbens *et al.*, "Fixed-point Models for the End-to-End Performance Analysis of IP Networks," *Proc. 13th ITC Specialist Seminar: IP Traffic Measurement, Modeling and Management*, Sept. 2000, Monterey, California.
[50] M. Mandjes *et al.*, "End-to-end Delay Models for Interactive Services on a Large-Scale IP Network," *Proc. 7th Wksp. Performance Modeling and Evaluation of ATM & IP Networks (IFIP99)*, 28–30 June 1999.
[51] M. Buchli *et al.*, "Policing Aggregates of Traffic with the Token Bucket Algorithm," *Proc. IEEE ICC2002*.
[52] R. Gibbens, F. Kelly, and P. Key, "A Decision Theoretic Approach to Call Admission Control in ATM Networks," *IEEE JSAC*, vol. 13, no. 6, Aug. 1995.
[53] J. Qiu and E. Knightly, "Measurement Based Admission Control with Aggregate Traffic Envelopes," *IEEE ACM Trans. Net.*, vol. 9, no. 2, April 2001 .
[54] P. Seeling, M. Reisslein, and B. Kulapala, "Network Performance Evaluation with Frame Size and Quality Traces of Single-Layer and Two-Layer Video: A Tutorial," *IEEE Commun. Surveys and Tutorials*, vol. 6, no. 3, 3rd Quarter 2004, pp. 58–78.
[55] W. Leland *et al.*, "On the Self-Similar Nature of Ethernet Traffic," *IEEE/ACM Trans. Net.*, vol. 2, no. 1, Feb. 1994.
[56] E. Gray *et al.*, "Trust Evolution Policies for Security in Collaborative Ad Hoc Applications," *Proc. 1st Int'l. Wksp. Security and Trust Management*, Milan, Sept. 15th 2005.
[57] C. McCann *et al.*, "A Measurement based Approach for MultiLevel Admission of Heterogeneous Traffic in Wireless Ad-Hoc Networks," *Proc. IEEE Military Commun. Conf. (MILCOM2004)*, 2004.

BIOGRAPHY

STEVEN WRIGHT [S'78, SM'99] (steven.wright@bellsouth.com) holds a B. Eng. (Elec.) degree from the University of Southern Queensland, Toowoomba, Australia, an M. B. A (Marketing) from Arizona State University, an M. Sc. (Computer Information Systems) from Boston University and a Ph. D. (Computer Engineering) from North Carolina State University. For more than 25 years, he has been working in the communications industry, in Australia, UK, Belgium and USA, with companies such as Plessey, GTE, Alcatel, Fujitsu, and, currently, BellSouth. He is a Principal Member of Technical Staff within the Advanced Network Architecture Con-

cepts group at BellSouth Science and Technology. His responsibilities are concerned with the technology evolution of BellSouth's networks and services, particularly for optical and packet technologies, architectures and service interworking arrangements. He has been active in various standards activities with contributions to the ANSI T1, ATM Forum, ATIS IIF and IETF committees; and has been a TPC member for various international conferences including ICC and Globecom. He has published more than 20 refereed articles in conferences and journals, and he holds two US patents.