Traffic Specification and Service Models

Abstract

A brief overview of traffic specificatin and service models for integrated-services networks is presented. QoS parameters are discussed and services models from IETF and ATM Forum are compared.

1 Traffic Contract Model

Integrated-services networks, as the name suggests, support different services over a common infra-structure. Each service should be well defined by a set of parameters describing what the user can expect from the service provider. Conversely, in order to be able to provide some kind of guarantee to the user about the characteristics of a service, the service provider may require that the user characterize his usage of the service, the traffic he is going to generate. Therefore, integrated-service network operation relies in a contract between user and service provider, the former committing to generate traffic between specifc limits and the latter committing to provide a service with the characteristics advertised as long as the user keeps his side of the contract.

In order for that contract to work, the characterization of both committing sides must be done precisely, through formal description methods. In this context, traffic specification models characterize the user committment side, while service models describe the quality of service (QoS) for each service offered by the service providers.

2 Traffic Specification Models

Traffic specification models are general concepts in the sense that they are not directly connected to specific service providers or service models, they just try to capture the user behavior. However, as mentioned above, traffic specifications can be used by service providers or service models to limit the behavior of users.

From the service provider point of view, important characteristics of traffic specifications are how well it describes the user behavior and also implementation issues. A complex and precise model may lose to a simpler model that gives enough bounds of the traffic. In the next sections, some simple and useful models are presented.

2.1 xmin, xave, I, s

xmin is minimum packet inter-arrival time, *xave* is the maximum average packet inter-arrival time over any time interval of length *I*, *s* is the maximum packet size.

2.2 leaky bucket/token bucket (s,r)

s is the maximal bucket depth, *r* is the token filling rate. For any time interval of length *t*, the traffic from the connection is less than $s + rt$.

2.3 {(s_1, r_1), (s_2, r_2), ..., (s_n, r_n)}

For any time interval of length *t*, the traffic from the connection is less than *min_i (s_i + r_i * t)*.

3 Service Models

In general we can say that there are three classes of services: best-effort, guaranteed and predicted. While best-effort doesn't provide specific guarantees, and the guaranteed services have hard bounds based on worst case values, predicted services work in the between, with more assurance than best-effort and less restrictive bounds than guaranteed.

Service models defined by IETF and ATM Forum are presented on Table 1 and Table 2. An interpretation of the positioning of those services in the context of application sensitivity to delay and bandwidth variations is then presented.

4 Implementation Issues

4.1 Scheduling algorithms for guaranteed traffic

A traffic filter and a scheduling algorithm together provide the guaranteed service. The bounds given in guaranteed service must be worst-case bounds, the primary function of a scheduling algorithms designed to deliver guaranteed service is to isolate flows from each other, so that a flow can have only a limited negative effect on other flows.

The traffic filter: It is a token bucket filter (r, b) with token filling rate r and bucket maximal depth b. A network client uses this to characterize its traffic load to the network, so that the network can commit bandwidth and manage queues in a way that realizes the service commitment.

The scheduling algorithm: Several time-stamp based algorithms have been developed, one of them is WFQ (weighted fair queueing algorithm). The WFQ scheme isolates each source from the others by providing it a specified share of the bandwidth under overload conditions.

4.2 Scheduling algorithms for predicted traffic

The traffic filter is same as that in guaranteed service.

The scheduling algorithm: The nature of applications allows the scheduling algorithm to delay all packets up to the play-back point without adversely affecting the application's performance. So a standard earliest-deadline-first scheduling algorithm can be used. FIFO and Priority are examples. The delays are shared in these algorithms, the post facto jitter bounds are smaller than when the sources are isolated from each other as in WFQ. FIFO and Priority allow a number of sources aggregating their traffic to obtain a lower overall delay jitter.

Isolation provides guaranteed service for well-behaved clients and quarantines misbehaving sources. Sharing performs the very different goal of mixing traffic from different sources in a way that is beneficial to all, bursts are multiplexed so that the post facto jitter is smaller for everyone.

FIFO is and effective sharing scheme and WFQ is an effective method for isolation. Priority is consisted of isolation mechanism and sharing mechanism.

5 References

References 1 - 12 from "Reading List on Resource Management for Application-aware Networks"

ATM Forum Traffic Management Specification Version 4.0

Table 1: IETF Services Model

Service	End-to-End Behavior	Motivation	TSpec	RSpec	Policing	Implementation Ex.
Best-Effort	• no guarantee • no admission control	• elastic applications	• none	• none	• none	• FIFO
Committed Rate	• committed transmission rate • minimal loss • no absolute delay guarantee • delay comparable to dedicated • admission control • relaxed guaranteed service • simpler admission control than controlled load	• approximates dedicated circuit • applications tolerant to delay	• token bucket (b,r) • peak rate (p) • minimum unit (m) • maximum packet (M)	• desired service rate (R)	• edge Non-conforming packets: • transfer to best effort • marked Inside the network: • reshaping	• local admission • circuit emulation by WFQ, Virtual Clock...
Controlled Load	• approximates best effort service under unloaded conditions • admission control • no absolute guarantees • delay comparable to unloaded • loss comparable to unloaded	• adaptive real-time applications	• token bucket (b,r) • peak rate (p) [1] • minimum unit (m) • maximum packet (M)	• none	• edge • each network element Non-conforming packets: • transfer to best effort with low priority[2] • reshaping distortions	• conservative approach: sum of the token buckets • two priority levels • traffic classes (WFQ...) • local admission
Controlled Delay	• three levels of delay control • admission control • no absolute guarantees • delays no worse than best effort • small losses • average bandwidth • level n has more tightly delay control than m if n<m • tails of delay distributions are reduced compared to best effort	• service and delay adaptive applications	• token bucket (b,r) • minimum unit (m) • maximum packet (M)	• service level (1,2,3) • exported info: maximal packet delay for T = 1 s, 60 s, and 3600 s for each level	• edge • heterogeneous source branch points • source merge points Non-conforming packets: • transfer to best effort[3] • marked	• three priority levels • end-to-end admission
Guaranteed	• an assured level of bandwidth • a delay-bounded service with no queueing loss • conforms to the fluid model • end-to-end queueing delay and bandwidth will be stable along a path • does not control the minimal or average delay of datagrams	• guarantees the datagrams will arrive within the guaranteed delivery time • not attempt to minimize the jitter, only maximal queueing delay	• token bucket (b,r) • peak rate (p) • minimum unit (m) • maximum packet (M)	• a rate R greater than or equal to r • a slack term S nonnegative • RSpec can be bigger than TSpec	• two forms: simple policing and reshaping • policing is done at the edge • reshaping at branch points and merge points Non-conforming packets: • transfer to best effort • marked	• WFQ • Jitter-EDD • Virtual Clock

1. Present for compatibility (?) only, not required.
2. Introduces reordering.
3. Introduces reordering.

Table 2: ATM Forum Services Model

Service	End-to-End Behavior	Motivation	TSpec	RSpec	Policing	Implementation Ex.
Constant Bit Rate (CBR)	• admission control • fixed amount of bandwidth available • statistical guarantees on delay and loss rate	• real-time applications	• Peak Cell Rate • Cell Delay Variation Tolerance	• peak-to-peak Cell Delay Variation • maximum Cell Transfer Delay • Cell Loss Ratio	• PCR and CLR on 0 and 1 flows • GCRA(PCR, CDVT)	
Real-Time Variable Bit Rate (rt-VBR)	• admission control • statistical guarantees on delay and loss rate	• real-time applications with bursty traffic	• Peak Cell Rate • Cell Delay Variation Tolerance • Sustainable Cell Rate • Maximum Burst Size	• peak-to-peak Cell Delay Variation • maximum Cell Transfer Delay • Cell Loss Ratio	• PCR, SCR, and CLR on 0 and 1 flows or • PCR on 0 and 1 flows, SCR and CLR on 0 flows (non-conforming packets may be switched to 1 flow) • GCRA(PCR, CDVT) • GCRA(SCR, BT+CDVT)	
Non-Real-Time Variable Bit Rate (nrt-VBR)	• admission control • statistical guarantees on loss rate	• non-real-time applications with bursty traffic	• Peak Cell Rate • Cell Delay Variation Tolerance • Sustainable Cell Rate • Maximum Burst Size	• Cell Loss Ratio	• PCR, SCR, and CLR on 0 and 1 flows or • PCR on 0 and 1 flows, SCR and CLR on 0 flows (non-conforming packets may be switched to 1 flow) • GCRA(PCR, CDVT) • GCRA(SCR, BT+CDVT)	
Unspecified Bit Rate (UBR)	• no admission control • no guarantees	• non-real-time applications such as file transfer or email	• optional Peak Cell Rate and Cell Delay Variation Tolerance[1]	• none	• optional PCR on 0 and 1 flows (non-conforming packets may be switched to 1 flow)	
Available Bit Rate (ABR)	• admission control[2] • low cell loss ratio • adapts to available network resources	• use flow control mechanism to allow source to change its rate in response to changing resources	• Peak Cell Rate • Minimum Cell Rate	• network promises to provide compliant connections a low Cell Loss Ratio	• complicated • ex: DGCRA(MCR, PCR, ICR, τ_1, τ_2, τ_3) • min: GCRA(PCR, τ_1)	• source and destination behaviors • switch behaviors: EPRCA, ERICA, ...

1. Optional admission control and UPC.
2. Connections specifying MCR = 0 shouldn't have admission denied.

IP Service Spectrum

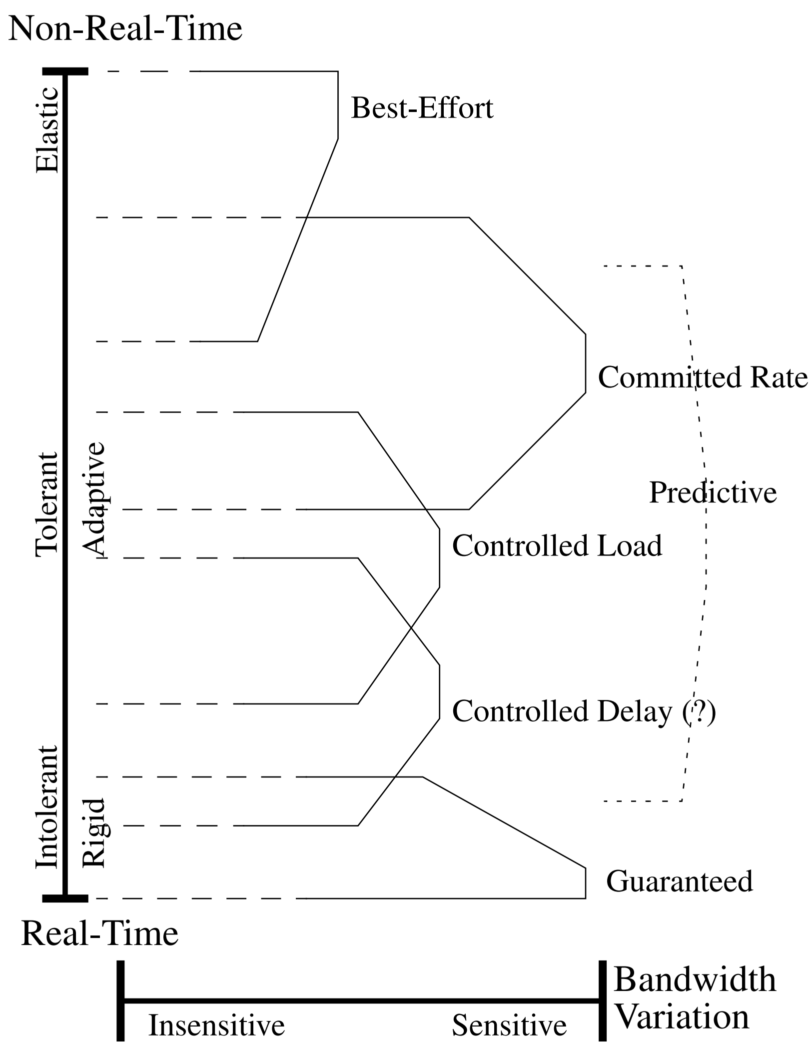

ATM Service Spectrum