Modelling the Performance of Distributed Admission Control for Adaptive Applications

Alan Bain* Peter Key†

1 Framework

We outline a methodology for the performance analysis of distributed admission control for adaptive applications, or long-lived flows, where each flow can transmit at one of a discrete set of rates, and can switch between these rates. Specific examples are certain voice over IP or conferencing applications. In the Internet, UDP based protocols such as RTP have the potential to allow sources to alter their sending rates using periodic feedback, where again the rate can take values from a limited set.

In *distributed admission control* [1, 4, 3], the end-system or edge-device plays a key role in the admission decision: it probes the network with a small number of packets and decides to enter or not depending on the fate of these probe packets (where the network drops or marks packets to signal congestion). These packets are marked at resources (routers) if these resources are nearing congestion, and the probe packets are fed back to the source. The source then decides whether to enter, and what rate to choose, according to the fraction of probing packets that have been marked. In addition, the application can occasionally reprobe the network and alter its rate.

2 Connections, packet-level models and marking

We concentrate on the simplest case where the resource load is solely generated by the class of calls under consideration. This is an appropriate model if such applications are segregated, for example in a separate DiffServ class.

We shall assume that there is a sufficiently great timescale difference between the packet level and the call level that individual acceptance/switching decisions may be considered as independent. Then at the call timescale, the end-user and marking behaviour may be taken as implicitly defining an acceptance probability $a^i(\rho)$ for a connection at level i, where ρ is the load on the network. For example, a particularly simple user policy is the following: send n packets into the network, enter at a high rate if no packets are marked, and otherwise enter at a low rate. In-call probing then uses the same policy, using packets of the data stream, hence when in-call probing occurs, $a^i(\rho)$ is also the probability of switching to level i. We assume that in-call probing is a relatively infrequent process, occurring a few times during a call.

We assume that routers mark packets according to some policy. Packet marks are fed back to the application. The marking can be done either by simply dropping the packet, or additionally by explicit marking, such as using the Congestion Experienced (CE) bit defined in the ECN proposal [7].

We suppose that calls of type r arrive as a Poisson process of rate ν_r, and when admitted last for an exponentially distributed holding time. While the connection is in progress, probing occurs as a random (Poisson) process of rate λ_r. We assume that (mean) packet generation rates are determined by the level i of the connection, in which case the load ρ on a resource is a linear function of number of calls at level i currently in progress.

Notice that when we have more than one distinct level, the load on the resource is not proportional to the number of connections in progress, hence we lose the reversibility which gave the product form solution of [4].

In a network with several resources, a user or route r may use a subset of these resources. Provided acceptance decisions across resources may be considered conditionally independent given the load, the acceptance functions have a product form.

3 Analysis and Performance

In order to understand the behaviour in a network of fixed topology, carrying large amounts of traffic (for example a backbone network), we consider a sequence of networks where the mean arrival rates of connections on each route grow proportionally to N, as does the route capacity, while the mean length of connections remains unchanged. Under this scaling we can find a fluid limit system of ordinary differential equations for $\mathbf{x}$, the vector of normalised number of calls at the different levels, $\frac{d}{dt}\mathbf{x} = f(\mathbf{x})$.

The fluid limit enables us to look at trajectories, or transient behaviour, and also examine the convergence and stability of the equilibrium point $\bar{\mathbf{x}}$. Under natural restrictions on the functions $a^i(\cdot)$, the equilibrium point is unique and stable.

We also show stability for the delay differential equations

*Statistical Laboratory, Cambridge, U.K.
†Microsoft Research Limited, Cambridge, U.K.

(caused by the round-trip-time delay between sending probe packets and receiving congestion indication signals). When there is no such delay, the fixed point is stable for all values of the probing rate λ. In contrast, when a delay is present there is a maximum value of λ for which the fixed point is stable.

To assess the performance of the network, we need to look at a finer level of detail. Let $\mathbf{U}(t)$ be the vector of differences of the normalised traffic vector from the equilibrium $\bar{\mathbf{x}}$. Then using the weak convergence results of Kurtz [2, 6, 5], under mild restrictions on the functions $a^i(\cdot)$, the following central limit theorem holds,

Theorem 1 *If in the limit as* $N \to \infty$, $\sqrt{N}\mathbf{U}(0) \stackrel{\mathcal{D}}{=} \mathbf{A}$ *for some random variable* $\mathbf{A}$*, then as* $N \to \infty$*, the deviations from the equilibrium fluid limit satisfy* $\lim_{N\to\infty} \sqrt{N}\mathbf{U}(t) \stackrel{\mathcal{D}}{=} \mathbf{R}_t$*, where* $\mathbf{R}_t$ *is the unique solution of the stochastic differential equation*

$$\mathrm{d}\mathbf{R}_t = H\mathbf{R}_t \mathrm{d}t + F\mathrm{d}\mathbf{B}_t,$$

where H *and* F *are matrices, and* $\mathbf{B}_t$ *is vector Brownian motion. This equation describes an Ornstein-Uhlenbeck process, where* H *is the matrix formed from linearization about the fixed-point of the fluid limit. The convergence described is weak convergence in the space of càdlàg paths, endowed with the Skorohod* J_1 *topology.*

We can express this as a functional central limit theorem:

Corollary 2 *Let the initial condition* $\mathbf{A}$ *be such that the stationary solution to the SDE is obtained, then at any fixed time* t*, the scaled difference vector* $\sqrt{N}\mathbf{U}_t$ *is distributed according to a multivariate normal distribution with mean zero and covariance matrix* Σ *given by*

$$\Sigma = \int_{-\infty}^{0} e^{-uH} F F^T (e^{-uH})^T \mathrm{d}u.$$

Integration by parts shows that Σ satisfies a Lyapunov stability equation $\Sigma H^T + H\Sigma = -FF^T$. Efficient and stable numerical techniques exists to solve this system.

Using these results, for given acceptance functions, the covariance matrix can be expressed in terms of the properties of the original network. We can calculate quantities such as the traffic variance, $v(\lambda)$, and examine their dependence on λ. The traffic variance can be also be used in a packet-loss calculation.

We examine in detail a single link, and certain specific network topologies. For concreteness, we use the Virtual Queue marking scheme [3, 4], although similar results hold for other marking schemes.

Analytic results for the single link show how the traffic variance, $v(\lambda)$ decreases sharply with λ. As a specific example, a value of $\lambda = 1$, corresponding to just one in-call probe per call on average, typically reduces the variance of the carried traffic by about 30%; while for $\lambda = 7$ the reduction is about 50%. Little is gained by having much larger values of λ.

These theoretical results were compared with a series of *ns* simulations of the same link, with a packet level implementation of the Virtual Queue scheme. The analyses and simulations show good agreement, even when some of the modelling assumptions were relaxed.

As another specific example, we consider a star network topology, where calls originate at edge nodes, to other randomly chosen edge nodes. This is motivated by an Internet core architecture with a transparent core switch connected to routers, which are represented by the edge nodes. Again we explicitly calculate the total traffic variance analytically, using the functional central limit theorem, and explore the performance of the system.

4 Concluding remarks

We have outlined a technique for modelling and analysing a large class of adaptive control problems. Once the fixed-point is found, the technique reduces to matrix algebra, hence it is widely applicable to arbitrary networks and control strategies. The performance estimates are produced through simple calculations rather than extensive simulations, hence they are relatively fast to calculate, and also enable the effect of altering a parameter to be seen easily.

By way of example, we have demonstrated in the specific case of a scheme which adapts between two rates that lightweight probing works, and only a very low reprobing rate is needed to achieve a significant benefit to the system; the further benefit gained by rapid reprobing is negligible.

References

[1] L. Breslau, E. Knightly, S. Shenker, I. Stoica, and H. Zhang. Endpoint admission control: Architectural issues and performance. In *Sigcomm 2000*, August 2000.

[2] S. Ethier and T. Kurtz. *Markov Processes – Characterization and Convergence*. Wiley, 1986.

[3] R. J. Gibbens and F. P. Kelly. Distributed connection acceptance control for a connectionless network. In P. Key and D. Smith, editors, *Teletraffic Engineering in a Competitive World, Proceedings ITC16*, pages 941–952. Elsevier, June 1999.

[4] F. P. Kelly, P. B. Key, and S. Zachary. Distributed admission control. *IEEE J. Sel. Areas in Comm.*, 18(12), December 2000.

[5] T. G. Kurtz. Strong approximation theorems for density dependent Markov chains. *Stochastic Process. Appl.*, 6:223–240, 1978.

[6] A. Mandelbaum, W. A. Massey, and M. I. Reiman. Strong approximations for Markovian service networks. *Queueing Systems – Theory and Applications*, 30:149–201, 1998.

[7] K. K. Ramakrishnan, S. Floyd, and D. Black. The addition of Explicit Congestion Notification (ECN) to IP. RFC 3168, IETF, September 2001.