Network Engineering of Elastic Data Traffic via Tandem Queueing Network Models

A. W. Berger[1], E. G. Coffman[2] and Y. A. Kogan[1]

1. AT&T Research, Middletown, NJ.
2. New Jersey Institute of Technology, Newark, NJ

Extended Abstract: In this paper, we provide asymptotic analysis for a special class of tandem closed queueing networks (CQNs). Asymptotic results are then used to derive engineering rules for capacity assignment and admission control for elastic data applications in packet-switched communication networks, such as Internet Protocol (IP) or Asynchronous Transfer Mode (ATM) networks. Elastic data applications adapt to time-varying available bandwidth via a feedback control such as the Transmission Control Protocol (TCP) or the Available Bit Rate transfer capability in ATM. Typical elastic data applications are file transfers supporting e-mail or the world wide web.

The tandem CQN model consists of a given number of finite sources (infinite servers) whose customers are sequentially served by two processor sharing (PS) servers. First, there is a designated PS server for each finite source, then all customers share a common downstream PS server. There could be also additional finite sources whose customers are served only at the downstream PS server.

An important practical feature of the CQN model is its insensitivity: the distribution of the underlying random variables influences the performance only via the mean. Thus, given the assumptions of the model, our engineering rules pertain in the topical case when the distribution of file sizes is heavy tailed (though with finite mean), and the superposition of file transfers is long-range dependent.

The tandem CQN models are motivated by the following network engineering scenario. Consider an access network where a set of enterprise LAN switches, or access concentrators, are connected to an edge router, which in turn is connected to a backbone router. For specificity, suppose a number of 10/100 Ethernets are connected to each LAN switch, which in turn are connected to the edge router via T3 lines, and the edge router is connected to the core router via an OC-12c Packet over SONET interface. The set of Ethernets on a given LAN switch corresponds to one of the finite sources. The T3 line from a LAN switch to the edge router corresponds to one of PS servers designated for a finite source, while the OC-12c line from the edge router to core router corresponds to the common downstream PS server. In this scenario, we are interested in the number of Ethernets that can be supported by a T3 line, and the number of T3 lines that can be supported by the OC-12c. Analogously, the capacities of the lines need not be given, but rather are to be determined given assumptions on the finite sources. A topical case would be where the "line" from the edge to core router is actually a segment of a Label Switched Path in Multi-Protocol Label Switching.

The main issue with application of CQN models to packet-switched communication network is the notion of a job. A job in the CQN model represents a file (or a set of files), as opposed to a packet. A job at a bottleneck PS node corresponds to a file (or a set of files) that is being transferred across the network. This modeling assumption is intended for the following network scenario. For dimensioning or admission control at the target link, we are interested in scenarios where the link is heavily loaded. That is, we take the conservative viewpoint that the target link should satisfy a per-flow bandwidth objective even when it is heavily loaded. The feedback controls of TCP and ABR tend to seek out and fill up the available bandwidth. A connection's feedback control, when properly designed and functioning, will attempt to keep at least one packet queued for transmission on the bottleneck link (otherwise the control is needlessly limiting the throughput). We assume that this is the case. (We view as complementary: (1) network design that assumes well-performing closed-loop controls, and (2) ongoing research for control implementations that make good use of the deployed bandwidth.) Thus we obtain a key simplifying assumption for the model: at an arbitrary point in time, the number of sources that are currently transferring a file equals the number of sources that have a packet in queue at a bottleneck, target link. Thus, in the CQN model, the number of jobs at a bottleneck PS node represents the number of sources that are currently transferring a file.

Of particular interest in the present paper is the complication that arises if for a given traffic source, the packet network, and thus the CQN model, contains more than one bottleneck link. In the data network, a file could have packets in queue simultaneously at the multiple bottleneck links, whereas in the CQN model a job moves sequentially from one node to the next. In which case, the equivalence between number of jobs at a bottleneck node and the number of files being transferred begins to break down. However, we show that the CQN can still be used to obtain engineering bounds for the bandwidth needed, or equivalently, for number of sources that can be supported.

The main results of this paper are the complete bottleneck classification of tandem CQNs for the whole range of their parameters and application of asymptotic results to decomposition of the initial engineering problem to engineering of bottleneck links in isolation. Such a decomposition is implied by the following key asymptotic property of tandem CQNs. If one of the first designated PS servers is a bottleneck, it will generate Poisson arrivals to the downstream PS server. That allows us to derive an asymptotic approximation for the marginal distribution for the number of customers at the downstream PS server for those sources for which the first PS server is not a bottleneck. This marginal distribution implies that the available bandwidth at the downstream PS server will be reduced by the bandwidth required for all sources where the first PS server is a bottleneck. The asymptotic results also justify that considering the first bottleneck link in isolation provides a conservative design.

Having this decomposition we can apply our previous results [1-3] on normal and refined normal approximation for bandwidth engineering of a single bottleneck link.

References

1. A. Berger and Y. Kogan, "Dimensioning Bandwidth for Elastic Traffic in High-Speed Data Networks", IEEE/ACM Transactions on Networks (submitted for publication).

2. A. Berger and Y. Kogan, "Multi-class elastic data traffic: Bandwidth engineering via asymptotic approximations", in P. Key and D. Smith (eds.), Teletraffic Engineering in a Competitive World, vol. 3a, pp. 77-86, Elsevier, Amsterdam, 1999.

3. A. Berger and Y. Kogan, "Distribution of processor-sharing customers for a large closed system with multiple classes", SIAM J. Appl. Math. (to appear).