Instability in stochastic and fluid queueing networks

John Hasenbein
Graduate Program in OR/IE
Department of Mechanical Engineering
The University of Texas at Austin
1 University Station C2200
Austin, TX 78712-0292, USA

Abstract

In a series of papers, starting in the early 1990's, researchers established a strong connection between the stability of a queueing network and the stability of the corresponding fluid model. Initiated by Rybko and Stolyar and generalized by Dai, Stolyar, and Chen, among others, it has been demonstrated that the stability of a fluid model implies stability of a corresponding queueing network. The stability results in the aforementioned papers were established both for classes of policies, e.g. the set of non-idling policies, and specific policies, e.g. First-In-First-Out.

Since the stability behavior of the fluid model is often significantly easier to analyze than that of the stochastic model, the results above have led to sweeping advances in understanding the stability of queueing networks via the fluid model. However, a major element needed for a satisfactory theory of stability via fluid models is a converse to the aforementioned stability results. Specifically: if the fluid model is *not* stable in some sense, does this imply instability of the corresponding queueing network? Unfortunately, it turns out that formulating an appropriate converse is a delicate matter. Partial converses which appear in the literature include important results by Dai, Meyn, Puhalskii, and Rybko. In all of the papers which prove a converse result to the original stability theorems, some uniform requirement over a set of

fluid trajectories, or more precisely a set of fluid limits, (sometimes restricted to fluid limits starting from a particular type of state) is needed for the result to be applicable. Recall that the original stability results of Dai and Chen require the *all* fluid trajectories are stable in some sense. Hence, we use the term “partial converse” above because there is some gap between the stability and instability results.

In a 1995 paper, Chen shows that a multiclass queueing network is globally rate stable if the corresponding fluid network is globally weakly stable. In this talk, we present a result which is a full converse to Chen¡s stability result. It is a full converse in that for some networks, in particular two station networks, the result implies that the stochastic network is globally rate stable if and only if the corresponding fluid network is globally weakly stable. In particular, this implies that if there is just one linearly divergent fluid trajectory, then the stochastic network is not rate stable under some non-idling policy. Combining our main result with a result of Dai and Vande Vate we show that a certain computable condition of the form $\rho^* \leq 1$ is a necessary and sufficient condition for rate stability in networks with two stations. This is the first tight condition for stability for such a general class of networks.

Our proof uses a series of large deviations estimates to establish the result and the only restriction in the stochastic network is that the estimates are applicable to the primitive stochastic processes defining the network. The result also relies on a property of fluid networks called the *finite decomposition property*, which holds for all two station networks, and can be shown to hold for some three station networks.

Joint work with David Gamarnik of IBM T. J. Watson Research Labs..