Forty-Seventh Annual Allerton Conference
Allerton House, UIUC, Illinois, USA
September 30 - October 2, 2009

Stability of Networks under General File Size Distribution with Alpha Fair Rate Allocation

Fernando Paganini*, Ao Tang†, Andrés Ferragut* and Lachlan L. H. Andrew‡
* Universidad ORT, Montevideo, Uruguay.
† Cornell University, Ithaca, NY.
‡ Swinburne University of Technology, Australia.

***Abstract*—Convex optimization has been widely used to model bandwidth allocation policies among TCP flows in the Internet. When the offered load is less than capacity, stochastic stability of networks using such policies has been established for exponentially distributed file sizes. The problem has remained open for general file size distributions, which is very relevant as it is well known that Internet file sizes follow heavy-tailed distributions. In this paper, building on existing results on the fluid model of the system, we use a partial differential equation to characterize the dynamics. The equation keeps track of residual file size and therefore is suitable to study networks under general file size distributions. For alpha fair bandwidth allocation, with any positive alpha parameter, a Lyapunov function is constructed to prove stability when the offered load does not exceed capacity. The result gives an affirmative answer to this open problem and sets a foundation for further analysis of such systems.**

I. Introduction and prior work

A fundamental step in the understanding of resource allocation mechanisms in the Internet has been the formulation by Kelly et al. [8] of congestion control in terms of network utility maximization. In a scenario of a fixed number of connections across routes in a network, this approach characterizes an equilibrium and leads to the formulation of dynamic, distributed methods to achieve it. An interesting class of utility functions is the "α-fair" family of Mo and Walrand [14]; by varying the α parameter it encompasses various notions of flow-level fairness, in particular proportional fairness ($\alpha = 1$) and max-min fairness ($\alpha \to \infty$).

However, this analysis with fixed numbers of connections does not capture the reality that flows come and go in the network, a situation better modeled through stochastic processes. This issue was identified by Roberts and Massoulié [15], who studied queueing systems with random arrivals and workloads, and a processor sharing discipline where service rates depend on bandwidth allocation, assumed to occur at a faster time-scale. This leads to a basic *stability* question, first posed by De Veciana et al. [5]: under which connection level demands (job arrival rate and mean workload) is the resulting queueing process stable? The answers given in [5] apply to Poisson arrivals and exponentially distributed job sizes, and max-min fair or proportionally fair bandwidth allocation. In this case the numbers of connections per route form a Markov chain, which is shown to be stable (i.e. ergodic) under the natural stability condition: namely, that the *mean load in each link of the network is strictly less than the link capacity.* In a subsequent paper by Bonald and Massoulié [1], these results were generalized to the α-fair case. More recent work where the time-scale separation assumption is relaxed is reported in [11].

The above results all rely on the critical restriction of exponentially distributed inter-arrival times and file sizes. The latter is particularly unsatisfactory, since it has been observed that file sizes in the Internet follow a heavy tailed Pareto-type distribution [3]. This has motivated recent efforts in extending the stability results for general file-size distributions. Removing the exponential file-size assumption is "*well-known to be a difficult problem*" [11]; without it, the number of ongoing connections is no longer a Markov state. Some existing partial results are the following: [12] has showed that for $\alpha = 1$, the stability result can be generalized to an appropriate Jackson-type routing scheme, thus providing the tool to establish the condition for phase-type file distributions; [10] gives a result for phase-type file distribution and general α, in two particular network topologies through Lyapunov functions obtained numerically.

A strategy that has proven relevant for this problem is the use of *fluid limits*, already invoked in [1] and further studied in [9], for the exponential case. The extension to general distributions is developed by Gromoll and Williams [6], [7], as briefly reviewed in Section II. Based on these models, [2] obtains a stability result for general file size distributions of bounded support, for a sufficiently small α. In [6] it is also established that the stability condition guarantees the fluid model is stable in the special cases of linear and tree networks.

In this paper, we build on the results of [6], [7] and establish in Section II a fluid model of the process in terms of a Partial Differential Equation for the distribution of remaining job workloads. We then show in Section III that the fluid model is stable under the natural stability condition regardless of the job size distribution and the underlying network topology, and for any value of α, by constructing a Lyapunov function that is shown to converge to zero. In Section IV we comment on the stochastic implication of this result.

II. Problem formulation and Partial Differential Equation model

The problem under consideration is the stability of a queueing system where flows arrive at various routes, and are served according to rates allocated by a congestion control algorithm. The latter is modeled through the network utility maximization problem

$$\max \sum_m z_m U_m\left(\frac{\varphi_m}{z_m}\right), \text{ subject to} \sum_m R_{lm}\varphi_m \leq c_l. \tag{1}$$

Here m denotes the route, z_m the number of connections in the route, and φ_m the total allocated rate of all connections of route m. R is the routing matrix ($R_{lm} = 1$ iff route m uses link l) and $c = (c_l)$ the vector of link capacity constraints. The utility function U_m is assigned to each connection as a function of its per-flow rate φ_m/z_m. In this paper we focus on the "α-fair" utility functions introduced in [14], where $U'_m(x) = x^{-\alpha}$.

Given $z = (z_m)$, it is assumed the congestion control sets $\varphi = (\varphi_m)$ to the optimum of (1); we assume separation of time-scales, i.e. the mapping $z \mapsto \varphi$ is instantaneous.

Remark 1. *A property of α-fair utility functions (with common α across routes) is that the resource allocation is invariant under scaling: namely, if all z_m are scaled by a common factor r, the resulting φ_m do not change.*

Consider now a network where flows arrive at route m through a stochastic process of mean intensity $\lambda_m > 0$, and a general distribution of the file sizes: let $G_m(\sigma)$ be the probability that the file size is greater than σ, and

$$\frac{1}{\mu_m} = \int_0^\infty G_m(\sigma)d\sigma$$

the mean file size, assumed to be finite. At any given time, currently active flows are served with the rate φ_m/z_m that results from (1).

The aim is to prove that if the loads $\rho_m := \lambda_m/\mu_m$ strictly satisfy the network capacity constraints,

$$\sum_m R_{lm}\rho_m < c_l \quad \forall l,$$

then the stochastic system is stable.

The classically studied case [5], [1] is when the arrival process is Poisson, and the file distribution exponential. In that case the process is a Markov chain with state $z = (z_m)$, and stability means positive recurrence. For the general case, the system state requires substantially more information, as discussed later on. We now turn our attention to fluid models for this problem, which are the basis of our stability studies.

A. Fluid model

We first recall the fluid model for the M/M case, following [1], [9]. This is obtained in the limit by scaling time and the initial condition of the process, leaving fixed the network capacity and the external load. Let r be the scaling parameter, and define $z^r(0) = rz(0)$, where $\|z(0)\| = 1$ in a suitable vector norm. If $z^r(t)$ is the resulting stochastic process as described above, the fluid limit is defined by

$$Z(t) = \lim_{r\to\infty} \frac{z^r(rt)}{r}.$$

Invoking the strong law of large numbers, [1] writes the following ordinary differential equation model:

$$\frac{dZ_m}{dt} = \lambda_m - \varphi_m(t)\mu_m \tag{2}$$

for each m. Here $\varphi_m(t)$ corresponds to the service rate with re-scaled time. Due to the scale invariance of the resource allocation (see Remark 1) the fluid versions of φ and Z are still related by the the analog of (1), i.e. $\varphi(t)$ is the maximizer of

$$\max \sum_m Z_m U_m\left(\frac{\varphi_m}{Z_m}\right), \text{ s.t.} \sum_m R_{lm}\varphi_m \leq c_l. \tag{3}$$

We refer to [9] for more details on this type of scaling.

We now state a basic inequality that will characterize the resource allocation. It follows directly from the fact that at the optimal point of (3), the feasible set must be inside a negative half-space defined by the gradient vector.

Lemma 1. *Let (φ_m) be the vector of rates that optimizes (3), and (ψ_m) another vector of rates satisfying the constraints $\sum_m R_{lm}\psi_m \leq c_l$. Then:*

$$\sum_m U'_m\left(\frac{\varphi_m}{Z_m}\right)(\psi_m - \varphi_m) \leq 0. \tag{4}$$

B. PDE Model

In the general distribution case, bandwidth allocation is still a function only of the numbers of flows $Z = (Z_m)$. However, once we remove the memoryless property of the exponential distribution, characterizing the network state requires keeping track of residual file-sizes, not just their number. Furthermore, the resource allocation per route is a processor sharing discipline, where all flows present receive equal service. This complicates the description since we must keep track of residual file sizes of *all* flows. In order to proceed, we look at the problem in more detail by modeling, in a fluid setting, the evolution of the residual file distribution.

Let $F_m(t,\sigma)$ ($t \geq 0$, and $\sigma \geq 0$) be the number of class (route) m files at time t with residual file size larger than σ, in the fluid limit. $F_m(t,\sigma)$ is a finer descriptor of the system than $Z_m(t)$, indeed the definition implies

$$Z_m(t) = F_m(t, 0). \tag{5}$$

We now model the evolution of $F_m(t,\sigma)$ through the following partial differential equation:

$$\frac{\partial F_m(t,\sigma)}{\partial t} = \frac{\partial F_m(t,\sigma)}{\partial \sigma}\frac{\varphi_m(t)}{Z_m(t)} + \lambda_m G_m(\sigma). \quad (6)$$

The above equation holds under the assumption that $Z_m(t) > 0$; it must be suitably complemented for $Z_m = 0$, as described below.

Note that (6) reduces to (2) in the exponential file size distribution case, $G_m(\sigma) = \exp(-\mu_m \sigma)$. This can be readily checked by using $F_m(t,\sigma) = Z_m(t)\exp(-\mu_m\sigma)$ in (6), which reduces it to (2).

C. An intuitive derivation of the PDE

At time $t+dt$, jobs that have residual file size at least σ come from two sources:

- New arrivals between t and $t+dt$ of size greater than σ. With arrival rate λ_m, we have $\lambda_m dt\, G_m(\sigma)$ such jobs, in the fluid limit.
- Files already present at time t, which had at that time a residual size of at least $\sigma + \frac{\varphi_m(t)}{Z_m(t)}dt$. Note each file receives a service rate $\varphi_m(t)/Z_m(t)$.

Therefore

$$F_m(t+dt,\sigma) = F_m\left(t, \sigma + \frac{\varphi_m(t)}{Z_m(t)}dt\right) + \lambda_m G_m(\sigma)dt \quad (7)$$

Subtracting $F_m(t,\sigma)$ from both sides and dividing by dt gives

$$\frac{F_m(t+dt,\sigma) - F_m(t,\sigma)}{dt} = \frac{F_m\left(t,\sigma + \frac{\varphi_m(t)}{Z_m(t)}dt\right) - F_m(t,\sigma)}{dt} + \lambda_m G_m(\sigma).$$

In the limit when $dt \to 0$ we obtain (6).

D. A formal justification based on [7]

We now explain how to relate (6) to the rigorous fluid limit set up in Gromoll and Williams [7].

In this formulation, the system state is characterized by a time-dependent, positive *measure* $\zeta_m(t)$ for each class (route) m. The measure is defined over the positive real numbers, representing distribution of residual workload. In particular, in the stochastic model $\zeta_m(t)$ at any given time is a finite sum of Dirac deltas, located at the sizes of remaining workloads for currently active jobs. The integral of this measure is the number of active jobs $Z_m(t)$.

In the fluid limit under appropriate scaling, the limiting measure $\zeta_m(t)$ satisfies (for all t except a set of Lebesgue measure zero) the following:

$$\frac{d}{dt}\langle f, \zeta_m(t)\rangle = \begin{cases} -\frac{\varphi_m(t)}{Z_m(t)}\langle f', \zeta_m(t)\rangle + \lambda_m \langle f, \nu_m\rangle, & \text{for } Z_m > 0; \\ 0, & \text{for } Z_m = 0. \end{cases} \quad (8)$$

This equation coincides with (5.62) in [7], modulo notational changes. Here the measure ν_m represents the probability distribution of arriving jobs; $f(\sigma)$ is an arbitrary bounded and continuously differentiable test function in the class

$$\mathcal{C} = \{f \in C_b^1(\mathbb{R}_+), f(0) = f'(0) = 0\};$$

and $\langle f, \nu\rangle := \int_0^\infty f(\sigma)d\nu$.

In this model, the probability of an arriving job being larger than σ, and the number of jobs at time t with residual workload greater than σ, are represented by

$$G_m(\sigma) := \int_\sigma^\infty d\nu_m; \qquad F_m(t,\sigma) := \int_\sigma^\infty d\zeta_m(t).$$

To derive the PDE we assume that the measures ν_m and $\zeta_m(t)$ are absolutely continuous with respect to Lebesgue measure. In particular,

$$d\nu_m = -G'_m(\sigma)d\sigma, \qquad d\zeta_m(t) = -\frac{\partial F_m(t,\sigma)}{\partial\sigma}d\sigma.$$

By integration by parts we have the following identities:

$$\langle f, \nu_m\rangle = -\int_0^\infty f(\sigma)G'_m(\sigma)d\sigma = -f(\sigma)G_m(\sigma)\Big|_{\sigma=0}^\infty + \int_0^\infty f'(\sigma)G_m(\sigma)d\sigma. \quad (9)$$

$$\langle f, \zeta_m(t)\rangle = -\int_0^\infty f(\sigma)\frac{\partial F_m(t,\sigma)}{\partial\sigma}d\sigma = -f(\sigma)F_m(t,\sigma)\Big|_{\sigma=0}^\infty + \int_0^\infty f'(\sigma)F_m(t,\sigma)d\sigma. \quad (10)$$

Due to the definition of the class $\mathcal{C}$, the incremental terms above vanish, which turns (8) into

$$\frac{d}{dt}\int_0^\infty f'(\sigma)F_m(t,\sigma)d\sigma = \frac{\varphi_m}{Z_m}\int_0^\infty f'(\sigma)\frac{\partial F_m(t,\sigma)}{\partial\sigma}d\sigma + \lambda_m\int_0^\infty f'(\sigma)G_m(\sigma)d\sigma, \quad (11)$$

for the case $Z_m > 0$. Assuming the differentiation with respect to t on the left can be interchanged with integration, the above yields

$$\int_0^\infty f'(\sigma)\mathcal{D}[F_m(t,\sigma)]d\sigma = 0, \quad (12)$$

where $\mathcal{D}[\cdot]$ is the differential operator given by

$$\mathcal{D}[F_m] := \frac{\partial F_m}{\partial t} - \frac{\partial F_m}{\partial\sigma}\frac{\varphi_m(t)}{Z_m(t)} - \lambda_m G_m(\sigma).$$

Since $f'(\sigma)$ is a free continuous function, we must have $\mathcal{D}[F_m] \equiv 0$, i.e. (6).

Remarks:

- Going beyond the above assumptions (absolute continuity, differentiation under the sign) involves dealing with integral equations, in particular equation

(3.1) from [7], the integral version of (8); this can be interpreted as defining (6) in the distributional sense. We will not pursue this here, and assume enough smoothness in $F_m(t,\sigma)$ for a classical treatment of the PDE, and to interchange differentiation with respect to t with integration over σ.

- If $Z_m = 0$, (8) leads to the condition $\frac{\partial F_m}{\partial t} = 0$, i.e. the system can stay at $F_m = 0$. It may seem awkward that the arrivals term is turned off here, as remarked in [7]. In intuitive terms, having $Z_m = 0$ for an open interval of time represents the "chattering" of the state around zero, when service rate is exceeding arrival rate.

III. Stability result

We will show that if the loads $\rho_m = \frac{\lambda_m}{\mu_m}$ strictly satisfy the capacity constraints, $\sum_m R_{lm}\rho_m < c_l$, then the solutions to our fluid model asymptotically converge to zero. Before proceeding we recapitulate the partial differential equation model

$$\frac{\partial F_m(t,\sigma)}{\partial t} = \begin{cases} \dfrac{\partial F_m(t,\sigma)}{\partial \sigma}\dfrac{\varphi_m(t)}{Z_m(t)} + \lambda_m G_m(\sigma) & Z_m > 0, \\ 0 & Z_m = 0, \end{cases} \tag{13}$$

and establish some basic facts involving the *residual workload function* $W_m(t)$ for each route m. This measures the total residual workload at time t, in the fluid limit, and can be expressed as

$$W_m(t) = \int_0^\infty \sigma d\zeta_m(\sigma) = \int_0^\infty F_m(t,\sigma)d\sigma. \tag{14}$$

Here the second step follows by integration by parts. The following Lemma (analogous to Lemma 3.3 in [7]) describes the evolution of $W_m(t)$.

Lemma 2. *Given a solution $F_m(t,\sigma)$ to (6), the workload function $W_m(t)$ defined in (14) satisfies*

$$\dot{W}_m = \begin{cases} \rho_m - \varphi_m(t), & Z_m > 0, \\ 0 & Z_m = 0, \end{cases} \tag{15}$$

and therefore the bound $W_m(t) \le W_m(0) + \rho_m t$. In particular, it remains finite for all time.

Proof: Focusing on the case $Z_m > 0$, integrating the PDE with respect to σ and using (5) yields

$$\begin{aligned} \dot{W}(t) &= \int_0^\infty \frac{\partial F_m(t,\sigma)}{\partial t} d\sigma \\ &= \frac{\varphi_m(t)}{Z_m(t)}\Big[F_m(t,\sigma)\Big]_0^\infty + \lambda_m \int_0^\infty G_m(\sigma)d\sigma \\ &= -\varphi_m(t) + \rho_m. \end{aligned}$$

□

A. Lyapunov function

Choose a sufficiently small $\delta > 0$ such that both $\tilde{\rho}_m = (1+\delta)\rho_m$ satisfies $\sum_m R_{lm}\tilde{\rho}_m < c_l$ for all l, and $(1-\delta)(1+\delta)^{\alpha+1} > 1$; recall that $\alpha > 0$ is the fairness parameter used by the congestion control. The second inequality always holds for $0 < \delta < \alpha/(2+\alpha)$; note $\delta \to 0$ as $\alpha \to 0$. Introduce the Lyapunov function

$$\begin{aligned} L(t) &= \sum_m L_m(t) \\ &= \sum_m \frac{1}{\tilde{\rho}_m^\alpha}\int_0^\infty [F_m(t,\sigma)]^{\alpha+1} p_m(\sigma)d\sigma. \end{aligned} \tag{16}$$

Here $p_m(\sigma)$ is a "spatial weight" to be selected shortly; we impose that it is non-negative and bounded in $\sigma \ge 0$, and normalized to $p_m(0) = 1$.

As a first remark, note that since $F_m(t,\sigma)$ is by definition monotonically non-increasing in σ, we have

$$F_m(t,\sigma)^{\alpha+1} \le Z_m(t)^\alpha F_m(t,\sigma),$$

therefore

$$\begin{aligned} L_m(t) &\le \frac{\|p_m\|_\infty Z_m^\alpha(t)}{\tilde{\rho}_m^\alpha}\int_0^\infty F_m(t,\sigma)d\sigma \\ &= \frac{\|p_m\|_\infty Z_m^\alpha(t)}{\tilde{\rho}_m^\alpha} W_m(t), \end{aligned} \tag{17}$$

finite for all time using Lemma 2. Therefore the Lyapunov function is well-defined.

We now compute the time derivative of

$$[\tilde{\rho}_m]^\alpha L_m = \int_0^\infty [F_m(t,\sigma)]^{\alpha+1}p_m(\sigma)d\sigma \tag{18}$$

along the trajectory, for any $m : Z_m > 0$. We have:

$$\begin{aligned} \tilde{\rho}_m^\alpha \dot{L}_m &= \int_0^\infty (\alpha+1)[F_m(t,\sigma)]^\alpha \frac{\partial F_m(t,\sigma)}{\partial t} p_m(\sigma)d\sigma \\ &= \frac{\varphi_m(t)}{Z_m(t)}\int_0^\infty \frac{\partial F_m^{\alpha+1}(t,\sigma)}{\partial\sigma}p_m(\sigma)d\sigma \\ &\quad + \int_0^\infty (\alpha+1)[F_m(t,\sigma)]^\alpha \lambda_m G_m(\sigma)p_m(\sigma)d\sigma. \end{aligned} \tag{19}$$

Integrating by parts in the first term, we have

$$\begin{aligned} &\int_0^\infty \frac{\partial F_m^{\alpha+1}(t,\sigma)}{\partial\sigma}p_m(\sigma)d\sigma \\ &= F_m^{\alpha+1}(t,\sigma)p_m(\sigma)\Big|_{\sigma=0}^\infty - \int_0^\infty F_m(t,\sigma)^{\alpha+1}p_m'(\sigma)d\sigma \\ &= -Z_m(t)^{\alpha+1} - \int_0^\infty F_m(t,\sigma)^{\alpha+1}p_m'(\sigma)d\sigma \end{aligned} \tag{20}$$

Substituting in (19) we obtain

$$\begin{aligned} \tilde{\rho}_m^\alpha \dot{L}_m = &-\varphi_m(t)Z_m(t)^\alpha \\ &+ \int_0^\infty F_m(t,\sigma)^\alpha\Big\{-\frac{\varphi_m(t)}{Z_m(t)}F_m(t,\sigma)p_m'(\sigma) \\ &\qquad + (\alpha+1)\lambda_m G_m(\sigma)p_m(\sigma)\Big\}d\sigma. \end{aligned} \tag{21}$$

B. Choice of weight $p_m(\sigma)$

We now specify that $p_m(\sigma)$ satisfies the following differential equation in σ,

$$p_m'(\sigma) = K\mu_m G_m(\sigma) p_m(\sigma)^{\frac{\alpha+1}{\alpha}} \tag{22}$$

for some $K \in (0, \alpha)$ to be specified later. This equation can be readily solved (for $p_m(0) = 1$) to yield

$$p_m(\sigma) = \left(1 - \frac{K\mu_m}{\alpha}\int_0^\sigma G_m(u)du\right)^{-\alpha}. \tag{23}$$

Note that since $\mu_m \int_0^\infty G_m(u)du = 1$, for $K < \alpha$ the term in brackets is strictly positive, bounded away from zero, so $p_m(\sigma)$ is well-defined, non-negative and bounded. With this choice, (21) becomes

$$\begin{aligned}\tilde{\rho}_m^\alpha \dot{L}_m =& -\varphi_m(t)Z_m(t)^\alpha + \\ &\int_0^\infty F_m(t,\sigma)^\alpha \Big\{ -\frac{\varphi_m(t)}{Z_m(t)} F_m(t,\sigma)p_m(\sigma)^{\frac{1}{\alpha}}K \\ &+ (\alpha+1)\rho_m \Big\} \mu_m G_m(\sigma) p_m(\sigma) d\sigma.\end{aligned} \tag{24}$$

C. Bounding the Lyapunov derivative.

We wish to upper bound the terms involving $F_m(t,\sigma)$ in the above integral. For this we calculate the maximum of the function

$$\gamma(F) := F^\alpha\{(\alpha+1)\rho_m - bF\},$$

over $F \geq 0$. Here we have denoted

$$b = \frac{\varphi_m(t)}{Z_m(t)} p_m(\sigma)^{\frac{1}{\alpha}} K. \tag{25}$$

By differentiation we have

$$\gamma'(F) = F^{\alpha-1}\{\alpha(\alpha+1)\rho_m - (\alpha+1)bF\},$$

which has a root

$$F^* = \frac{\alpha\rho_m}{b} = \frac{\alpha\rho_m Z_m}{K\varphi_m p_m^{1/\alpha}},$$

and yields a maximum

$$\gamma(F^*) = \left(\frac{\alpha\rho_m Z_m}{K\varphi_m p_m^{1/\alpha}}\right)^\alpha \rho_m = \frac{\alpha^\alpha \rho_m^{\alpha+1} Z_m^\alpha}{K^\alpha \varphi_m^\alpha p_m}.$$

Returning to (24), we obtain the bound

$$\begin{aligned}\tilde{\rho}_m^\alpha \dot{L}_m \leq& -\varphi_m(t)Z_m(t)^\alpha + \\ &\int_0^\infty \frac{\alpha^\alpha\rho_m^{\alpha+1}Z_m(t)^\alpha}{K^\alpha \varphi_m(t)^\alpha p_m(\sigma)} \mu_m G_m(\sigma)p_m(\sigma)d\sigma \\ =& -\varphi_m(t)Z_m(t)^\alpha \\ &+ \frac{\alpha^\alpha\rho_m^{\alpha+1}Z_m(t)^\alpha}{K^\alpha\varphi_m(t)^\alpha}\int_0^\infty \mu_m G_m(\sigma)d\sigma \\ =& -\varphi_m(t)Z_m(t)^\alpha + \frac{\alpha^\alpha\rho_m^{\alpha+1}Z_m(t)^\alpha}{K^\alpha\varphi_m(t)^\alpha}.\end{aligned} \tag{26}$$

Since K is a free parameter, restricted only by $K < \alpha$, we can now choose it to satisfy

$$\left(\frac{\alpha}{K}\right)^\alpha = (1-\delta)(1+\delta)^{\alpha+1} > 1. \tag{27}$$

Then (26) becomes

$$\dot{L}_m \leq Z_m(t)^\alpha \left\{ -\frac{\varphi_m(t)}{\tilde{\rho}_m^\alpha} + \frac{\tilde{\rho}_m(1-\delta)}{\varphi_m(t)^\alpha} \right\}. \tag{28}$$

Lemma 3. *For any positive numbers $\tilde{\rho}$, φ,*

$$-\frac{\varphi}{\tilde{\rho}^\alpha} + \frac{\tilde{\rho}}{\varphi^\alpha} \leq (\alpha+1)\frac{(\tilde{\rho}-\varphi)}{\varphi^\alpha}. \tag{29}$$

Proof: Bounding the convex function $h(x) = x^{\alpha+1}$ by its tangent around the point $x = \tilde{\rho}$ gives

$$\varphi^{\alpha+1} \geq \tilde{\rho}^{\alpha+1} + (\alpha+1)\tilde{\rho}^\alpha(\varphi - \tilde{\rho})$$

Dividing by $\tilde{\rho}^\alpha\varphi^\alpha$ and reordering terms yields (29). □

We now state the main result.

Theorem 4. *Suppose $\delta > 0$ is such that $\tilde{\rho}_m = (1+\delta)\rho_m$ satisfy the capacity constraints $\sum_m R_{lm}\tilde{\rho}_m < c_l$ and $(1-\delta)(1+\delta)^{\alpha+1} > 1$. Then, the Lyapunov function L defined in (16) with $p_m(\sigma)$ in (23), where K is chosen as in (27), satisfies*

$$\dot{L} \leq -\delta \sum_{m:Z_m>0} \tilde{\rho}_m \left(\frac{Z_m}{\varphi_m}\right)^\alpha. \tag{30}$$

In particular, $\lim_{t\to\infty} L(t) = 0$.

Proof: We use the bound (29) in (28), and obtain

$$\dot{L}_m \leq (\alpha+1)\left(\frac{Z_m}{\varphi_m}\right)^\alpha (\tilde{\rho}_m - \varphi_m) - \delta\tilde{\rho}_m\left(\frac{Z_m}{\varphi_m}\right)^\alpha \tag{31}$$

for any m where $Z_m > 0$. Note also that $\dot{L}_m = 0$ when $Z_m = 0$, (refer to Remark in Section II-D). Superimposing all terms we get

$$\begin{aligned}\dot{L} = \sum_m \dot{L}_m \leq& (\alpha+1)\sum_{m:Z_m>0}\left(\frac{Z_m}{\varphi_m}\right)^\alpha(\tilde{\rho}_m - \varphi_m) \\ &- \delta\sum_{m:Z_m>0}\tilde{\rho}_m\left(\frac{Z_m}{\varphi_m}\right)^\alpha.\end{aligned} \tag{32}$$

Noting that $\left(\frac{Z_m}{\varphi_m}\right)^\alpha = U_m'\left(\frac{\varphi_m}{Z_m}\right)$, we are in a position to apply (4), with $\psi_m = \tilde{\rho}_m$ that satisfy the capacity constraints. This is the only step that relies on the underlying congestion control resource allocation. We have thus proved (30).

To obtain the asymptotic result, we first note that a bound of the form

$$L_m \leq (A_m + B_m t)Z_m^\alpha \tag{33}$$

holds for appropriately defined constants A_m, B_m. This follows from (17) and Lemma 2. This leads to the inequalities

$$\begin{aligned}\tilde{\rho}_m \left(\frac{Z_m}{\varphi_m}\right)^\alpha &\geq \frac{L_m \tilde{\rho}_m}{(A_m + B_m t)\varphi_m^\alpha} \\ &\geq \frac{L_m \rho_0}{C^\alpha (A_m + B_m t)} \\ &\geq \frac{L_m}{(A + Bt)}.\end{aligned}$$

Here $\rho_0 = \min_m \rho_m$ and $C = \max_l c_l$, and finally A_m, B_m are maximized across m, defining A, B appropriately. Returning to (30), we obtain

$$\dot{L} \leq -\frac{\delta}{(A + Bt)} \sum_m L_m = -\frac{\delta}{(A + Bt)} L.$$

This yields

$$\log L(t) \leq \log L(0) - \int_0^t \frac{\delta}{(A + Bu)} du.$$

Since the right-hand side diverges to $-\infty$, we have $L(t) \to 0$. □

The above result shows *asymptotic* stability, in the Lyapunov sense, of the fluid model (6). Note that the speed of convergence is controlled by δ, and this parameter goes to zero as $\alpha \to 0$. This is consistent with the fact that for $\alpha = 0$, the network need not be stable [1, Example 1].

We also study the possibility of obtaining convergence in finite time, often invoked when connecting fluid and stochastic models. We state the following result.

Proposition 5. *Suppose there exists a constant κ such that the solutions to (6) satisfy the bound*

$$W_m(t) \leq \kappa Z_m(t). \tag{34}$$

Then, under the hypothesis of Theorem 4, $L(t)$ converges to zero in finite time, proportional to

$$L(0)^{\frac{1}{\alpha+1}} = \|F(0, \tau)\|_{\alpha+1},$$

the norm of the initial condition in the function space $\mathcal{L}^{\alpha+1}$.

Proof: Applying (34) to (17) yields

$$L_m \leq \gamma Z_m^{\alpha+1},$$

with $\gamma > 0$ appropriately defined, already maximized over m. This leads to

$$\begin{aligned}\sum_m \tilde{\rho}_m \left(\frac{Z_m}{\varphi_m}\right)^\alpha &\geq \frac{\rho_0}{C^\alpha} \sum_m Z_m^\alpha \\ &\geq \frac{\rho_0}{C^\alpha \gamma^{\frac{\alpha}{\alpha+1}}} \sum_m (L_m)^{\frac{\alpha}{\alpha+1}} \\ &\geq \frac{\rho_0}{C^\alpha \gamma^{\frac{\alpha}{\alpha+1}}} \max_m (L_m)^{\frac{\alpha}{\alpha+1}} \\ &\geq \frac{\rho_0}{C^\alpha \gamma^{\frac{\alpha}{\alpha+1}}} \left(\frac{1}{M} \sum_m L_m\right)^{\frac{\alpha}{\alpha+1}} \\ &= \frac{\rho_0}{C^\alpha (\gamma M)^{\frac{\alpha}{\alpha+1}}} L^{\frac{\alpha}{\alpha+1}}.\end{aligned}$$

Here M is the number of sources. Returning again to (30), we obtain

$$\dot{L} \leq -\epsilon L^{\frac{\alpha}{\alpha+1}}, \qquad \epsilon = \frac{\delta \rho_0}{C^\alpha (\gamma M)^{\frac{\alpha}{\alpha+1}}}.$$

This leads to

$$\begin{aligned}&\frac{d}{dt} L^{\frac{1}{\alpha+1}} \leq -\frac{\epsilon}{\alpha+1} \\ \implies \quad &L(t)^{\frac{1}{\alpha+1}} \leq L(0)^{\frac{1}{\alpha+1}} - \frac{\epsilon}{\alpha+1} t.\end{aligned}$$

Therefore $L(t)$ must reach zero in a finite time bounded by

$$\frac{\alpha+1}{\epsilon} L(0)^{\frac{1}{\alpha+1}} = \frac{(\alpha+1) C^\alpha (\gamma M)^{\frac{\alpha}{\alpha+1}}}{\delta \rho_0} \|F(0, \tau)\|_{\alpha+1}.$$

□

We have thus showed finite time convergence under the additional condition (34). A special case in which this is easily verified is when the file-size distribution is *bounded*, i.e. $G_m(\sigma) = 0$ for $\sigma > \sigma_0$, and the initial condition $F_m(0, \sigma)$ satisfies the same bound. In that case, there are never any residual jobs larger than σ_0,

$$F_m(t, \sigma) = 0 \text{ for } \sigma > \sigma_0, \text{ all } t \geq 0;$$

this can also be verified through (6). From here we immediately have the workload bound

$$W_m(t) = \int_0^{\sigma_0} F_m(t, \sigma) d\sigma \leq \sigma_0 Z_m(t),$$

of the form (34) as required. So finite-time convergence follows in this case.

IV. Stochastic stability discussion

We briefly discuss here the relationship between the fluid-level stability result and the stability of a network with stochastic flow arrivals and departures. A first issue is what is meant by stochastic stability. In the exponential file size case, the stochastic process is a Markov chain where the number of connections $Z = (Z_m)$ per class is the state; in this case stability is usually defined to be positive recurrence of the Markov chain.

The natural generalization of the Markov model to general arrival times and file sizes (G/G) that form a renewal process is the one used by Dai [4]. There, a Markov process is defined where the state keeps track of residual arrival times and service times of currently active jobs, in addition to queue sizes. Stability is defined as the positive Harris recurrence of such Markov process. Dai [4] also obtains a fluid limit model, and defines a notion of stability for fluid models in terms of convergence to zero in finite time, similar to the one obtained in Proposition 5.

Does fluid stability imply stochastic stability? Dai [4] establishes this for service disciplines where the number of residual times in the state remains bounded. This does not cover processor sharing disciplines, where all jobs present in the system receive service, as is the case for our problem. Although [4] claims that extensions to this case "should be evident", we share the view of Gromoll and Williams [6], [7] that such extensions are not straightforward. Indeed, although [6] establishes the fluid limit stability for certain special topologies (linear network, tree networks), the authors stop short of making claims about the stochastic model. To do so would require a theorem that fluid stability implies Harris recurrence in the state-space of measures considered by [7], not currently available to our knowledge.

In the absence of this theorem, an alternative proof route for stochastic stability would be to apply our Lyapunov function directly to the Markov process and invoke generalized versions of Foster's criterion [13]. This possibility remains open for future research.

V. Conclusion

We have considered the conjecture that the natural condition (all mean link loads strictly below capacity) suffices for the stability of a network with randomly arriving files of *general* size distributions, when jobs are served with α-fairness. Building on recent fluid limit studies [7], we formulated a partial differential equation model for the problem, where the state $F_m(t, \sigma)$ represents the residual workload distributions per route.

A Lyapunov function, defined through a suitably weighted spatial $\alpha+1$-norm of F_m, is shown to converge to zero asymptotically along trajectories of the PDE. This gives an affirmative answer to the conjecture, in the sense of fluid models. We have also refined the result to finite-time convergence under an additional condition which is satisfied for the bounded distribution case.

Future work will involve carrying this conclusion through to the stochastic process.

Acknowledgements

This work was supported by NSF under CCF-0835706, by AFOSR-US, and by ANII-Uruguay, grant FCE 2007_265.

References

[1] T. Bonald and L. Massoulié. Impact of fairness on Internet performance. In *Proc. ACM Sigmetrics*, 2001.

[2] M. Chiang, D. Shah and A. Tang. Stochastic stability under fair bandwidth allocation: General file size distribution. In *Proc. Allerton Conference*, 2006.

[3] M. Crovella, A. Bestavros. Self-similarity in World Wide Web Traffic: Evidence and possible causes *IEEE/ACM Trans. Networking*, 5(6):835–846, December 1997.

[4] J.G. Dai, "On positive Harris recurrence of multiclass queueing networks: A unified approach via fluid limit models", *Ann. Appl. Probab.* 5:49–77, 1995.

[5] G. de Veciana, T.-J. Lee, and T. Konstantopoulos. Stability and performance analysis of networks supporting services with rate control — Could the Internet be unstable? In *Proc. IEEE INFOCOM*, volume 2, pp. 802–810, 1999.

[6] H. Gromoll and R. Williams. Fluid model for a data network with α-fair bandwidth sharing and general document size distributions: two examples of stability. In *IMS Collections — Markov Processes and Related Topics: A Festschrift for Thomas G. Kurtz*, vol.4, pp. 253–265, 2008.

[7] H. Gromoll and R. Williams. Fluid limits for networks with bandwidth sharing and general document size distributions. *Ann. Appl. Probab.*, 19:243–280, 2009

[8] F. Kelly, A. Maulloo and D. Tan Rate control in communication networks: shadow prices, proportional fairness and stability, *J. Op. Res. Soc.*, 39:237–252, 1998.

[9] F. P. Kelly and R. J. Williams Fluid model for a network operating under a fair bandwidth-sharing policy. *Ann. Appl. Probab.* 14:1055–1083, 2004.

[10] A. Lakshmikantha, C. Beck and R. Srikant. On the use of SoS methods for analysis of connection-level stability in the Internet. In *Proc. American Control Conference*, Portland, OR, 2005.

[11] X. Lin, N. Shroff and R. Srikant. On the connection-level stability of congestion-controlled communication networks. *IEEE Trans. Information Theory*, 54(5):2317–2338, May 2008

[12] L. Massoulié. Structural properties of proportional fairness: stability and insensitivity. *Ann. Appl. Probab.*, 17(3):809–839, 2007.

[13] S. P. Meyn and R. L. Tweedie. Stability of Markovian Processes III: Foster-Lyapunov criteria for continuous time processes. *Adv. in Appl. Probab.* 25:518–548, 1993.

[14] J. Mo and J. Walrand Fair end-to-end window based congestion control. *IEEE/ACM Trans. Networking*, 8(5):556–567, Oct. 2000.

[15] J. Roberts, L. Massoulié. Bandwidth Sharing and Admission Control for Elastic Traffic, In *Proc. ITC Specialist Seminar*, Yokohama, October 1998.