



UvA-DARE (Digital Academic Repository)

A stability conjecture on bandwidth sharing networks

Walton, N.S.; Mandjes, M.R.H.

DOI

[10.1007/s11134-011-9233-2](https://doi.org/10.1007/s11134-011-9233-2)

Publication date

2011

Document Version

Final published version

Published in

Queueing Systems

[Link to publication](#)

Citation for published version (APA):

Walton, N. S., & Mandjes, M. R. H. (2011). A stability conjecture on bandwidth sharing networks. *Queueing Systems*, 68(3-4), 237-250. <https://doi.org/10.1007/s11134-011-9233-2>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, P.O. Box 19185, 1000 GD Amsterdam, The Netherlands. You will be contacted as soon as possible.

A stability conjecture on bandwidth sharing networks

N.S. Walton · M.R.H. Mandjes

Received: 9 May 2011 / Revised: 9 May 2011 / Published online: 9 July 2011
© The Author(s) 2011. This article is published with open access at Springerlink.com

Abstract We consider a queueing model where documents are simultaneously transferred over a communication network. The bandwidth allocated to each document transfer is assumed to be the solution of a utility optimization problem. Under a natural stability condition and under the assumption that document arrivals are Poisson and that document sizes are independent exponential distributions, such queueing models have been proven to be positive recurrent. It has been conjectured for a decade that the assumption of exponentially distributed documents can be removed. There exist numerous generalizations without this exponential assumption, but a general proof remains elusive.

Keywords Utility optimization · Bandwidth sharing · Internet congestion control · Stability · Open problem

Mathematics Subject Classification (2000) Primary 90B20 · Secondary 60K25 · 60K30 · 90B18

1 Introduction

In the Internet, documents are divided into a number of packets, which are then sent sequentially across the network from sender to receiver. The Internet could be unstable in two senses: the number of packets within the network could be diverge, or the number of documents in transfer could be diverge.

N.S. Walton (✉) · M.R.H. Mandjes
Korteweg-de Vries Institute for Mathematics, University of Amsterdam, P.O. Box 94248,
1090 GE Amsterdam, The Netherlands
e-mail: n.s.walton@uva.nl

M.R.H. Mandjes
e-mail: m.r.h.mandjes@uva.nl

Internet routers have a fixed buffer size, and so, the number of packets in flight, however large, is always finite. In this sense, the Internet is stable, and the question of how to send packets across the Internet is more concerned with providing desirable transfer rates or *bandwidth*. This task of finding good transfer rates is implemented by a sender and receiver with the Transmission Control Protocol (TCP). At one conceptual level, TCP can be thought of as optimizing the aggregate utility of the bandwidth received by different document transfers on the Internet; see [20]. Given such a utility optimizing bandwidth allocation policy, can a network process all the document transfers in progress? More precisely, we will ask if a specific stochastic model of document transfer is positive Harris recurrent. Throughout this paper, the term *stable* refers to the positive Harris recurrence of a Markov process, and *unstable* refers to transience of a Markov process.

Suppose that documents arrive on each route r at rate ν_r and each of these documents has expected size μ_r^{-1} . The amount of work arriving per unit time is then $\rho_r = \frac{\nu_r}{\mu_r}$. If the rate of work arriving at a network resource j is strictly above the resource capacity to process this work, C_j , then certainly instability will arise. So, the condition that the rate of work arriving is less than the capacity provides a natural necessary condition for stability,

$$\sum_{r:j \in r} \rho_r < C_j, \quad j \in \mathcal{J}. \quad (1.1)$$

Even so, with this condition in place, one could imagine a traffic regime where a network is unable to process a document transfer before a new transfer arrives. Such instability, for example, is known to occur in models of radio packets switch networks; see [1].

Given that users receive a bandwidth allocation that optimizes a utility function, is condition (1.1) sufficient to guarantee the stability of the Markov process describing the document transfers in progress? When documents arrive from independent Poisson process and documents sizes are independent exponentially distributed, then it has been shown that condition (1.1) is sufficient for stability; see [3]. Document sizes for Internet transfers are known to follow a heavy-tailed distribution, and thus the assumption that document transfers are exponentially distributed is not entirely satisfactory. In this article we discuss the conjecture that condition (1.1) is sufficient for stability for general document size distributions.

If such a conjecture holds, then this would suggest that congestion controllers provide a viable mechanism to provide the maximum level of stability. As we will review shortly, the resolution of this conjecture has spurred a great deal of interesting research literature; nonetheless, the conjecture remains open.

We now concentrate on describing the Markov chain model which we wish to study and reviewing the progress made by different authors. The conjecture which we discuss is presented in Sect. 2.4.

2 Model, conjecture and literature review

We now build the various components of our model of document transfer over a communication network. First, we define the structure of a communication network. Sec-

only, we define the utility optimizing bandwidth allocations that might be achieved by a congestion controller. We then define a stochastic model of document transfer when documents are exponentially distributed and then finally when documents are of a general size. Throughout this section we introduce and discuss literature relevant to our conjecture. The conjecture itself is presented in Sect. 2.4.

2.1 Network structure

We let \mathcal{J} index the set of resources of a communication network. With each resource $j \in \mathcal{J}$, we associate a capacity $C_j > 0$. A route through the network, r , is a (nonempty) set of resources. We let $\mathcal{R} \subset 2^{\mathcal{J}}$ index the set of routes. We let n_r give the number of document transfers in progress on route r . As a function of $n = (n_r : r \in \mathcal{R})$, route r document transfers are allocated an amount of bandwidth $\Lambda_r(n) \in \mathbb{R}_+$. This bandwidth is then shared equally amongst all the transfers present on route r . In other words, within each route, a processor sharing discipline is used. This important feature of our model will be discussed later. The bandwidth available to each route is constrained by the available capacity at each resource. So,

$$\sum_{r:j \in r} \Lambda_r(n) \leq C_j, \quad j \in \mathcal{J}. \tag{2.1}$$

We call any such function $\Lambda(\cdot)$ a *bandwidth allocation policy*. In a communication network, each document transfer receives a rate which may vary depending on the number of transfers present on each route. Here a bandwidth allocation policy is used in order to define this transfer rate.

2.2 Network utility maximization and fairness

Congestion control within the Internet is chiefly governed by the Transmission Control Protocol (TCP). It has been suggested that TCP implicitly attempts to maximize the utility of the transfer rate received by its users; see [20]. With this in mind, with each route $r \in \mathcal{R}$, we associate a utility function $U_r(\cdot)$, a real-valued increasing, strictly concave function of the bandwidth allocated to each route r user. A congestion controller then allocates bandwidth in order to optimize the following objective:

$$\text{maximize} \quad \sum_{r \in \mathcal{R}} n_r U_r \left(\frac{\Lambda_r}{n_r} \right) \tag{2.2a}$$

$$\text{subject to} \quad \sum_{r:j \in r} \Lambda_r \leq C_j, \quad j \in \mathcal{J}, \tag{2.2b}$$

$$\text{over} \quad \Lambda_r \geq 0, \quad r \in \mathcal{R}. \tag{2.2c}$$

In addition, if $n_r = 0$, we assume that $\Lambda_r(n) = 0$. The solution to the above optimization represents the equilibrium rate of packet transfer in a communications network given the number of documents in transfer on each route. Over different values of

$n = (n_r : r \in \mathcal{R}) \in \mathbb{Z}_+^{\mathcal{R}}$, the solutions to this optimization problem, $\Lambda(n)$, define a bandwidth allocation policy achieved by a congestion control mechanism.

We think of a utility function $U_r(x_r)$ as summarizing an Internet connections benefit from receiving bandwidth x_r . We then think of the optimization problem (2.2) as allocating the available bandwidth in order to maximize the total benefit its users. Because a bandwidth allocation policy may attempt to allocate some amount of bandwidth to all transfers, one often thinks of a bandwidth allocation being *fair*.

2.2.1 Literature on network utility maximization

In general, a bandwidth allocation policy need not be particularly fair. For example, it could attempt to maximize the total bandwidth within the network. This maximization leads to some transfers receiving zero bandwidth. As we shall see, this *maximum throughput* allocation can lead to bad stability properties. A bandwidth allocation can be very fair. For example, subject to network capacity constraints, one could attempt to maximize the transfer rate of the document with the smallest bandwidth and then maximize the second smallest and so forth. . . . Such a policy would achieve *max-min fairness*. In the context of communication networks, this was the first fairness criterion to be considered; see [2]. The utility optimization framework considered here was first introduced by Kelly [18], and in particular, *proportional fairness* where $U_r(x_r) = \log(x_r)$ has proved a particularly tractable criterion for analysis.

A popular class of utility functions introduced by Mo and Walrand [31] is the set of *weighted α -fair utility functions*. For this class of utility functions, U_r takes the form

$$U_r(x_r) = \begin{cases} w_r \frac{x_r^{1-\alpha}}{1-\alpha} & \text{for } \alpha > 0, \\ w_r \log(x_r) & \text{for } \alpha = 1. \end{cases}$$

Here the class of utility functions is parameterized by weights $w = (w_r : r \in \mathcal{R}) \in \mathbb{R}_+$ and fairness parameter $\alpha > 0$.

The weighted α -fair class has proved popular as it contains a number of popular fairness criteria: proportional fairness ($\alpha = w_r = 1$); TCP fairness ($\alpha = 2, w_r = \frac{1}{T_r}$), which also converges to maximum throughput ($\alpha \rightarrow 0, w_r = 1$), and max-min fairness ($\alpha \rightarrow \infty, w_r = 1$). Although certainly other fairness criteria exist, the weighted α -fair family parameterizes the set of fairness criteria of interest. For this reason, weighted α -fairness often acts as a convenient starting point for results on fair bandwidth sharing.

Although it is not the central objective of this article, good introductions to the relationship between utility optimization and congestion control are the short article of Kelly [17] and the book by Srikant [33].

2.3 A stochastic model with exponentially distributed document sizes

Until now, we have discussed an equilibrium that might be achieved by packets in transfer across a communication network. We, now, define a model that incorporates the stochastic arrival and departure of documents and also allocates service with an allocation policy. We ask if this stochastic model has an equilibrium. This model can

be thought of as a model of document transfer across the Internet. We introduce our stochastic model for a bandwidth allocation policy $\Lambda(\cdot)$.

Documents to be transferred arrive as a Poisson process. Route $r \in \mathcal{R}$ documents arrive as an independent Poisson process of rate $\nu_r > 0$. Each document has a size that is independent exponentially distributed with mean μ_r^{-1} . We, thus, define the traffic intensity of work arriving at route r by $\rho_r = \frac{\nu_r}{\mu_r}$. When there are $n = (n_r : r \in \mathcal{R})$ documents in transfer on each route, each route r document is served at rate $\frac{\Lambda_r(n)}{n_r}$. Documents are then processed at this rate until the number of documents in transfer changes either by a document departure or arrival.

Here $n = (n_r : r \in \mathcal{R})$, the number of documents in transfer, gives the state of a Markov chain with nonzero transition rates

$$q(n, n + e_r) = \nu_r, \tag{2.3a}$$

$$q(n, n - e_r) = \mu_r \Lambda_r(n) \quad \text{if } n_r > 0. \tag{2.3b}$$

This model Markov chain model was introduced by Massoulié and Roberts [25]. We now discuss the stability of this stochastic model.

2.3.1 An example of instability

We note here that seemingly sensible bandwidth allocation policies need not be stable. The following example is taken from [3].

Consider the network that consists of two resources $\mathcal{J} = \{A, B\}$, each with capacity 1, and three routes: $r_0 = \{A, B\}$, $r_1 = \{A\}$, and $r_2 = \{B\}$. This network is called a two-node linear network. This is because our network forms a line, r_0 , with two types of cross traffic, r_1 and r_2 .

Suppose that we use a policy that maximizes the total bandwidth allocated. In this case, it is preferred to allocate bandwidth to routes r_1 and r_2 over r_0 because route r_0 uses capacity at both resources A and B , whilst r_1 and r_2 only use one resource. So we use a bandwidth allocation policy that gives all available bandwidths to routes r_1 and r_2 whenever there is a document in transfer on either route. Only when there are no documents in transfer on route r_1, r_2 , we allocate capacity to route r_0 , i.e.,

$$(\Lambda_{r_0}(n), \Lambda_{r_1}(n), \Lambda_{r_2}(n)) = \begin{cases} (0, 1, 1) & \text{if } n_{r_1} > 0 \text{ and } n_{r_2} > 0, \\ (0, 1, 0) & \text{if } n_{r_1} > 0 \text{ and } n_{r_2} = 0, \\ (0, 0, 1) & \text{if } n_{r_1} = 0 \text{ and } n_{r_2} > 0, \\ (1, 0, 0) & \text{if } n_{r_0} > 0, n_{r_1} = 0, \text{ and } n_{r_2} = 0, \\ (0, 0, 0) & \text{if } n_{r_0} = 0, n_{r_1} = 0, \text{ and } n_{r_2} = 0. \end{cases}$$

This policy maximizes the total amount of bandwidth allocated and therefore places the maximum amount of effort toward processing documents. But is this allocation policy stable under condition (1.1)? Unfortunately, it is not stable.

Suppose that condition (1.1) holds. We note that route r_1 , and similarly r_2 , has priority over route r_0 and so route r_1 behaves as a single-server queue with server capacity 1. Also, as they do not interact, one can see that the number of documents

transfers on route r_2 is independent of route r_1 . So for routes 1 and 2, the stability condition (1.1) is sufficient for routes 1 and 2 to be positive recurrent. If we look at route r_0 , we notice that it can only process documents when both routes r_1 and r_2 have no documents in transfer. The long-run proportion of time that this occurs is given by the stationary probability that no documents are in transfer on routes r_1 and r_2 , which is $(1 - \rho_1)(1 - \rho_2)$. For the stability of class r_0 , the long-run proportion of capacity devoted to r_0 documents must be strictly larger than the rate at which work arrives, and thus we derive the stability condition

$$\rho_0 < (1 - \rho_1)(1 - \rho_2).$$

This condition is more restrictive than condition (1.1), and so, condition (1.1) is not sufficient for stability in this example.

2.3.2 Literature on stability with exponentially distributed documents

As we have mentioned, the stochastic model considered above with exponentially distributed documents sizes is first considered by Massoulié and Roberts [25]. Massoulié and Roberts [25] prove that condition (1.1) is sufficient for stability for the proportionally fair allocation policy on a linear network. Massoulié and Roberts [25] do this by explicitly calculating the stationary distribution of the proportionally fair model on this network topology.

Subsequently, De Veciana and Konstantopoulos [12] prove that (1.1) is sufficient for the stability for both proportionally fair and max-min fair allocation policies under a general network topology. De Veciana and Konstantopoulos [12] proved the stability by the construction of an appropriate Lyapunov function and then by applying Foster's lemma.

The generalization of these results to the rest of the weighted α -fair family was proven by Bonald and Massoulié [3]. Once again, their arguments consisted of constructing an appropriate Lyapunov function. Instead of applying their Lyapunov function directly to their Markov processes, the Lyapunov function is applied to the “fluid model” associated with the Markov process (2.3). A *fluid model* is a deterministic processes associated with the Markov processes (2.3). In the case of a weighted α -fair bandwidth allocation policy $\Lambda(\cdot)$, the appropriate *fluid model* is a positive solution to the differential equation

$$\frac{dn_r(t)}{dt} = \nu_r - \mu_r \Lambda_r(n), \quad n_r > 0, r \in \mathcal{R}. \quad (2.4)$$

Intuitively, this differential equation makes sense: it relates the rate that documents arrive and depart to changes in the number of documents present.

To derive (2.4), one must prove a “fluid limit,” a formal law of large numbers argument that proves the convergence of the Markov chain model (2.3) to a solution of the differential equation (2.4). For weighted α -fair allocation policies, this fluid limit result is proven in [19]. For general discussions on proving fluid limits, see [10] and [11].

Given that the fluid model (2.4) is related to the Markov process (2.3), the stability of the differential equation (2.4) may perhaps relate to the positive recurrence

of the Markov process (2.3). A result proved by Dai [8] shows that if any fluid limit model eventually reaches zero and stays there, regardless of the initial system configuration, then the underlying queueing process is positive Harris recurrent. Bonald and Massoulié prove the stability of the associated fluid model with the Lyapunov function

$$F(n) := \sum_{r \in \mathcal{R}} w_r \mu_r^{-1} \rho_r^{-\alpha} \frac{n_r^{1+\alpha}}{\alpha + 1}.$$

To give some feeling as to how such Lyapunov functions are applied, we present and discuss the proof of Bonald and Massoulié [3] in the [Appendix](#).

The Lyapunov function of Bonald and Massoulié [3] was next generalized to apply to a general utility function by Ye [34]. Aside from a few technical conditions bounding derivatives, this result demonstrated that condition (1.1) was sufficient for the network stability when utility functions were general and when documents had exponentially distributed sizes.

Finally, we note that there are some Markov chain generalizations beyond the model presented. Liu et al. [23] present and prove the stability of a model where the capacity region (2.1) may be nonconvex and may vary in time. Also, in practice congestion controllers may take some time to converge to a utility optimizing bandwidth allocation. The paper [22] discusses and proves the stability when a utility optimizing bandwidth allocation is not achieved instantaneously.

2.3.3 A further note on fluid limits and positive recurrence

A key component of the proof of Bonald and Massoulié [3] was the stability result presented in [8] and [9]. The work of Dai applies to queueing network where customers have general service requirements, but the result of Dai comes with the caveat that if a queueing model processes jobs in different classes/routes, the service discipline must be *head-of-the-line*, meaning that within each class documents are processed with a first-in-first-out service discipline. As we described in Sect. 2.1, service within a route is processor sharing not head-of-the-line. Due to the memoryless property, processor sharing and head-of-the-line service disciplines are equivalent when documents have an exponentially distributed size. So the results of Dai do apply to the case where documents are exponentially distributed, but unfortunately they do not directly apply to document sizes that are not exponentially distributed. To deal with this issue, new theory must be developed.

2.4 A stochastic model with generally distributed document sizes

The assumption of exponentially distributed document sizes is somewhat idealistic, and so it is desirable to model the transfer of documents of a generally distributed size. The description of such a Markov process is a natural extension of our previous Markov chain model.

Documents to be transferred arrive as a Poisson process. Route $r \in \mathcal{R}$ documents arrive as an independent Poisson processes of rate $\nu_r > 0$. Each document on route r has a size that is independent identically distributed according to some random

variable X_r . A route r document’s distribution has support on $(0, \infty)$ and has mean μ_r^{-1} . Once again, we define the traffic insensitivity of route r by $\rho_r = \nu_r/\mu_r$. Between document arrival and departure events, on each route r , a document is processed linearly at rate $\frac{A_r(n)}{n_r}$.

In order to give an Markov process description of our communication network, we must record the residual sizes of documents, that is, the amount of each document which remains to be transferred. Given that there are n_r documents to be transfer on route r for $k = 1, \dots, n_r$, we let $y_{rk} > 0$ be the residual size of the k th document in transfer on route r . We index elements y_{rk} so that $y_{r1} \leq y_{r2} \leq \dots \leq y_{rn_r}$. The state of the documents in transfer on route r is then given by the vector

$$y_r = (y_{rk} : k = 1, \dots, n_r),$$

and the state of the documents in transfer is given by

$$y = (y_r : r \in \mathcal{R}).$$

The variable y_{rk} decreases linearly at rate $\frac{A_r(n)}{n_r}$ until a document departure or arrival occurs. That is, given that the Markov process takes state y^0 at time t_0 ,

$$y_r(t) = y_r^0 - (t - t_0) \frac{A_r(n^0)}{n_r^0} \quad \text{for } t > t_0,$$

while $y_{r1}(t) > 0$ (before a document departure) and while $(t - t_0) < \min_r \{S_r\}$ (before a document arrival). Here each S_r is an independent exponential random variable with parameter ν_r . We let $y(t^-)$ give the state of our Markov process instantaneously before time t .

If a document departure occurs at time t , that is, $y_{r1}(t^-) = 0$ and $(t - t_0) < \min_r \{S_r\}$, then $n_r(t)$ is updated to equal $n_r(t^-) - 1$, and $y_r(t)$ is updated so that

$$y(t) = (y_{rk}(t^-) : k = 2, \dots, n_r).$$

The random variables S_r represent the Poisson arrival of documents. If a document arrival occurs on route r , i.e., $t - t_0 = S_r = \min_{\tilde{r}} S_{\tilde{r}}$ and $y_{\tilde{r}1}(t) > 0$ for all $\tilde{r} \in \mathcal{R}$, then an independent random variable X_r is drawn according to the distribution of route r documents. This gives the size of the arriving document. Accordingly, $n_r(t)$ is updated to equal $n_r(t^-) + 1$, and $y_r(t)$ is updated to give

$$y_r(t) = \begin{cases} (X_r, y_{r1}(t^-), \dots, y_{rn_r}(t^-)) & \text{for } X_r \leq y_{r1}(t^-), \\ (y_{r1}(t^-), \dots, y_{rk}(t^-), X_r, y_{rk+1}(t^-), \dots, y_{rn_r}(t^-)) \\ \quad \text{for } y_{rk}(t^-) \leq X_r \leq y_{rk+1}(t^-), \\ (y_{r1}(t^-), \dots, y_{rn_r}(t^-), X_r) & \text{for } y_{rn_r}(t^-) \leq X_r. \end{cases}$$

A new exponential parameter ν_r random variable S'_r is then drawn to replace S_r .

The process we describe here is a piecewise linear cádlág Markov processes. We could also describe the state of each route y_r as a measure with a Dirac mass at y_{rk} ; this *measure-valued* description is used by Gromoll and Williams [15] and Bramson [6].

Remark 1 We remark that this model could be generalized so that interarrival times are independent identically distributed rather than exponentially distributed. Residual interarrival time could then be incorporated to give a Markov state description. We do not include this extension, firstly, because a Poisson process is justifiable arrival process, it is the aggregation of a large number of rare events, and secondly, because allowing general document sizes is our primary modeling objective.

We are interested in knowing when our Markov process has a stationary distribution. We cannot use the theory of countable-state-space Markov chains, and so, in this setting, a slightly different notion of recurrence and positive recurrence is used. In particular, the corresponding notions of recurrence and positive recurrence are *Harris recurrence* and *positive Harris recurrence*. We will shortly give pointers to precise definitions and literature on this topic, but first we state our conjecture.

The principle conjecture of interest is the following:

Conjecture 1 *The Markov process described above, a bandwidth sharing network operating under a bandwidth allocation which optimizes (2.2) and with general document sizes, is positive Harris recurrent under the condition*

$$\sum_{r:j \in r} \rho_r < C_j, \quad j \in \mathcal{J}.$$

This conjecture is widely expected to hold for weighted α -fair utility functions. Like the result of Ye [34], for general utility functions, some additional restrictions may be required.

2.4.1 Discussion on positive Harris recurrence and queueing

The theory of Harris recurrence was firstly developed by its namesake [16]. A processes is said to be *Harris recurrent* if there exists a σ -finite measure ϕ , defined on the state space of our Markov process, S , such that

$$\phi(A) > 0 \quad \text{implies} \quad \mathbb{P}_y(\tau_A < \infty), \quad x \in S,$$

where \mathbb{P}_y is the law of our Markov process started at y , A is a Borel-measurable subset of our Markov process' state space, and τ_A is the hitting time of the Markov process on set A .

As in the countable state space theory, it has been shown that Harris recurrence is sufficient for a Markov process to have a stationary measure, see [13]. If this measure is a finite measure, then a stationary distribution exists, and our process is said to be *positive Harris recurrent*.

A crucial role in proving positive Harris recurrence is the existence of a petite set. In words, a petite set is a set of states from which there is some measure determined lower bound on the probability of reaching any other set of states. The existence of a petite set that has finite expected hitting time after leaving itself is equivalent to a Markov process being positive Harris recurrent. Thus, to prove the positive Harris recurrence, one must find a petite set and prove that its expected return time is finite.

An excellent account of this approach for discrete-time Markov processes is given by Meyn and Tweedie [28]. A continuous-time treatment is given by the same authors in the papers [27, 29, 30].

So, in the context of our queueing applications and in the approach of Dai [8, 9], one needs to characterize a petite set and then show that it has a finite expected return time. To characterize a petite set, some technical assumptions must be made on document interarrival times. These assumptions are presented in [26] and further in [5, Sect. 4.2]. These assumptions are satisfied for an independent Poisson arrival process. To prove that expected return times are finite, one can apply Foster's lemma. Proposition 4.6 of [5] presents a form of Foster's lemma, the *multiplicative Foster's criterion*. This version of Foster's lemma is used to execute the stability proof in [8, 9]. The book of Bramson [5] provides a superb account of the theory relevant to proving the stability of queueing networks and should be the first place to start on this topic.

2.4.2 Literature on stability with generally distributed documents

For general document sizes, initial results on the stability of utility optimizing bandwidth allocations were proven by Bonald and Proutière [4]. Bonald and Proutière [4] studied the behavior of *insensitive bandwidth allocations*. Bonald and Proutière [4] showed that the proportionally fair allocation policy coincided with an insensitive policy for line, grid, and hypercube network topologies. They thus proved that (1.1) gave the stability for the proportionally fair policy with phase-type document sizes for these specific networks. For lines and the special case of a 2×2 grid, this result is also observed by Lakshminantha et al. [21].

By further developing the link between proportional fairness and insensitivity, Massoulié [24] proves stability of proportionally fair bandwidth allocation policies with the assumption that documents sizes are from phase-type. Massoulié [24] proves results for any network topology.

Under a relatively mild second-moment condition, Bramson [6] proves the stability for networks operating under a max-min fair allocations policy with generally distributed document sizes.

With the aim of establishing a fluid limit proof of stability, as given by Dai [8] for head-of-the-line disciplines, a number of authors have established fluid limits for bandwidth sharing networks with general document sizes. In particular, Chiang et al. [7] prove fluid limit results for weighted α -fair bandwidth allocations with α sufficiently close to zero. Chiang et al. [7] use this result to place probabilistic bounds on their queueing network's state. Similarly, Gromoll and Williams [15] provide a formal fluid limit argument for bandwidth networks with general document sizes. Gromoll and Williams [14] show the fluid limit stability for tree network and line network topologies. Further, a fluid model of weighted α -fair bandwidth sharing is stated in the paper [32]. Here the authors prove the stability of their partial differential equation model. As it stands, there is no proof that the fluid model stability implies the positive Harris recurrence, and so theory must be developed so that these papers imply the stability of the underlying Markov process.

Open Access This article is distributed under the terms of the Creative Commons Attribution Noncommercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

Appendix: Fluid limit and stability

The paper [3] considers the Markov chain (2.3) for a weight α -fair bandwidth allocations. The paper proves that a necessary and sufficient condition for positive recurrence of the Markov chain is that $\sum_{r:j \in \mathcal{R}} \rho_r < C_j$ for each $j \in \mathcal{J}$. We will state one result needed for the proof, introduce some notation, and then give their proof.

Theorem 1 ([8]) *If any fluid model n satisfying (2.4) is such that $\exists T > 0$ s.t. $\forall t > T, n(t) = 0$, then the Markov chain (2.3) is positive recurrent.*

Consider the objective function of the weighted α -fair optimization problem

$$G_n(\Lambda) = \begin{cases} \sum_{r \in \mathcal{R}} w_r n_r \frac{1}{1-\alpha} \left(\frac{\Lambda_r(n)}{n_r}\right)^{1-\alpha} & \text{for } \alpha \neq 1, \\ \sum_{r \in \mathcal{R}} w_r n_r \log \Lambda_r & \text{for } \alpha = 1. \end{cases}$$

From $G_n(\Lambda)$ we can heuristically construct a Lyapunov function. Assuming the condition (1.1), the weighted α -fair bandwidth allocation $\Lambda(n)$ and traffic intensity ρ both satisfy the networks capacity constraints. Because $\Lambda(n)$ is optimal, $G_n(\Lambda(n)) > G_n(\rho)$. In this sense, the processing rate $\Lambda(n)$ is out performing the work arrival rate ρ . We now ask what is the long-run effect of this. By applying a Taylor expansion, we have

$$\begin{aligned} & \int_0^T [G_{n(t)}(\Lambda(n(t))) - G_{n(t)}(\rho)] dt \\ &= \int_0^T (\Lambda(n(t)) - \rho) \cdot \nabla G_{n(t)}(\rho) + o(|\Lambda(n(t)) - \rho|) dt \\ &= - \int_0^T \sum_r (\rho_r - \Lambda_r(n(t))) w_r \left(\frac{n_r(t)}{\rho_r}\right)^\alpha dt + o\left(\int_0^T |\Lambda(n(t)) - \rho| dt\right) \\ &= - \int_0^T \frac{1}{\mu_r} \frac{dn_r(t)}{dt} w_r \left(\frac{n_r(t)}{\rho_r}\right)^\alpha dt + o\left(\int_0^T |\Lambda(n(t)) - \rho| dt\right) \\ &= F(n(0)) - F(n(t)) + o\left(\int_0^T |\Lambda(n(t)) - \rho| dt\right), \end{aligned}$$

where we define the nonnegative function

$$F(n) := \sum_{r \in \mathcal{I}} w_r \mu_r^{-1} \rho_r^{-\alpha} \frac{n_r^{1+\alpha}}{\alpha + 1}.$$

Thus,

$$\begin{aligned} F(n(t)) &= F(n(0)) - \int_0^T [G_{n(t)}(\Lambda(n(t))) - G_{n(t)}(\rho)] dt \\ &\quad + o\left(\int_0^T |\Lambda(n(t)) - \rho| dt\right). \end{aligned}$$

$F(n)$ determines the cumulative benefit of $\Lambda(n)$ out performing ρ . If the network is stable, in the long run, $\Lambda(n)$ is approximately ρ . So, the little o term above should have a small effect on the evolution of $F(n(t))$.

Now we can give the theorem:

Theorem 2 ([3]) *The Markov chain given by (2.3) is positive recurrent iff $\sum_{r:j \in \mathcal{R}} \rho_r < C_j$ for each $j \in \mathcal{J}$*

Proof We first prove the if direction of the result. As we observed above,

$$\frac{dF(n(t))}{dt} = \nabla G_n(\rho) \cdot (\rho - \Lambda(n(t))),$$

and $G_n(\Lambda(n)) \geq G_n(u)$ for any $u \in \mathbb{R}_+^{\mathcal{R}}$ satisfying the networks capacity constraints $\sum_{r:j \in \mathcal{R}} u_r \leq C_j$. Thus, the partial derivate of the convex function $G_n(\cdot)$ in the direction $\Lambda(n) - u$ is positive, i.e.,

$$\nabla G_n(u) \cdot (\Lambda(n(t)) - u) \geq 0.$$

In particular, we take $u = (\rho_r(1 + \epsilon) : r \in \mathcal{R})$, where $\epsilon > 0$ is sufficiently small so that u satisfies the network capacity constraints. So,

$$0 \geq \nabla G_n(u) \cdot (\rho(1 + \epsilon) - \Lambda(n)) \times (1 + \epsilon)^\alpha = \sum_{r \in \mathcal{R}} w_r \left(\frac{n_r}{\rho_r}\right)^\alpha (\rho_r(1 + \epsilon) - \Lambda_r).$$

Thus, for $n(t) \neq 0$,

$$\begin{aligned} -\epsilon \sum_{r \in \mathcal{R}} w_r \rho_r \left(\frac{n_r}{\rho_r}\right)^\alpha &\geq \sum_{r \in \mathcal{R}} w_r \left(\frac{n_r}{\rho_r}\right)^\alpha (\rho_r - \Lambda_r) = \nabla G_n(\rho) \cdot (\Lambda(n(t)) - \rho) \\ &= \frac{dF(n(t))}{dt}. \end{aligned}$$

Notice that $\|n\|_1 := (\sum_r w_r \rho_r^{1-\alpha} n_r^\alpha)^{1/\alpha}$ and $\|n\|_2 := (F(n))^{1/\alpha+1}$ both define norms on $\mathbb{R}^{\mathcal{R}}$. Since all finite-dimensional norms are equivalent, there is a constant β such that $\|n\|_1 \geq \beta \|n\|_2$, and so we have

$$\frac{dF(n)}{dt} \leq -\epsilon \beta^\alpha F(n(t))^{\alpha/(\alpha+1)}. \tag{A.1}$$

If $F(n(t)) = 0$, then (A.1) implies that $\frac{dF(n(t))}{dt} = 0$, and so $F(n(t')) = 0 \forall t' > t$, and hence, $n(t') = 0 \forall t' > t$. Hence, all we need now is to check that $F(n(T)) = 0$ for some $T \geq 0$. Integrating (A.1) yields

$$\begin{aligned} F(n(T))^{\frac{1}{\alpha+1}} - F(n(0))^{\frac{1}{\alpha+1}} &= \int_0^T F(n(s))^{\frac{-\alpha}{1+\alpha}} dF(n(s)) \leq -\epsilon \int_0^T \beta^\alpha ds = -\epsilon \beta^\alpha T \\ \Rightarrow F(n(T)) &\leq \left(F(n(0))^{1/(1+\alpha)} - \frac{\epsilon \beta^\alpha}{1 + \alpha} T \right)_+^{\frac{1}{\alpha+1}}. \end{aligned}$$

We can find $T > 0$ such that the left-hand side of the above equation is 0, hence making $F(n(T)) = 0$.

We now show the only if direction of the theorem. Suppose that, on the contrary, there exists a link j such that $\sum_{r:j \in r} \rho_r \geq C_j$. Set \tilde{W} such that, for our Markov chain (2.3), every time a document arrives on a route r with $j \in r$, an amount of work equal to the size of that document is added to \tilde{W} . Let the work at \tilde{W} be processed at constant rate $\tilde{C} = C_j$. Hence, $\tilde{C} \geq \sum_{r:j \in r} \Lambda_r(N_r(t))$ for all t , and so in this coupling the amount of work \tilde{W} is greater than the amount of work on all routes associated with queue j . \tilde{W} is the workload process of a multiclass single-server queue with different input flows for each r such that $j \in r$. It is known that such a multiclass single-server queue is positive recurrent iff $\sum_{r:j \in r} \frac{\rho_r}{C_j} < 1$. But by assumption it is not, and so the expected time for \tilde{W} to reach zero is infinite. As the expected time for this workload process to reach zero is infinite, the same must be true for our bandwidth model. Hence, the Markov chain (2.3) is not positive recurrent, and so we have a contradiction. \square

References

1. Aldous, D.: Ultimate instability of exponential back-off protocol for acknowledgment-based transmission control of random access communication channels. *IEEE Trans. Inf. Theory* **33**, 219–223 (1987)
2. Bertsekas, D., Gallager, R.: *Data Networks*. Prentice Hall, New York (1987)
3. Bonald, T., Massoulié, L.: Impact of fairness on Internet performance. *Proc. ACM Sigmetrics* **29**, 82–91 (2001)
4. Bonald, T., Proutière, A.: Insensitive bandwidth sharing in data networks. *Queueing Syst.* **44**, 69–100 (2003)
5. Bramson, M.: *Stability of Queueing Networks*. Lecture Notes in Mathematics, vol. 1950. Springer, Berlin (2008). Lectures from the 36th Probability Summer School held in Saint-Flour, July 2–15, 2006
6. Bramson, M.: Network stability under max-min fair bandwidth sharing. *Ann. Appl. Probab.* **20**, 1126–1176 (2010). doi:10.1214/09-AAP649
7. Chiang, M., Shah, D., Tang, A.: Stochastic stability under network utility maximization: General file size distribution. In: *Proceedings of Allerton*, vol. 5, pp. 49–77 (2006)
8. Dai, J.G.: On positive Harris recurrence of multiclass queueing networks: A unified approach via fluid limit models. *Ann. Appl. Probab.* **5**(1), 49–77 (1995)
9. Dai, J.G.: Stability of open multiclass queueing networks via fluid models. In: *Stochastic Networks. The IMA Volumes in Mathematics and Its Applications*, vol. 71, pp. 71–90. Springer, Berlin (1995)
10. Darling, R.W.R.: Fluid limits of pure jump Markov processes: A practical guide. *CoRR arXiv:math/0210109v3* (2002)
11. Darling, R.W.R., Norris, J.R.: Differential equation approximations for Markov chains. *Probab. Surv.* **5**, 37 (2008). doi:10.1214/07-PS121
12. De Veciana, G.L., Konstantopoulos, T.: Stability and performance analysis of networks supporting services with rate control—could the Internet be unstable? *IEEE/ACM Trans. Netw.* **9**(1), 2–14 (1999)
13. Gettoor, R.K.: Transience and recurrence of Markov processes. In: *Séminaire de Probabilités, XIV (Paris, 1978/1979)* (French). *Lecture Notes in Math.*, vol. 784, pp. 397–409. Springer, Berlin (1980)
14. Gromoll, H., Williams, R.: Fluid model for a data network with alpha-fair bandwidth sharing and general document size distributions: two examples of stability. *Markov Process. Relat. Top.* **4**, 253–265 (2008). A Festschrift for Thomas G. Kurtz
15. Gromoll, H., Williams, R.: Fluid limits for networks with bandwidth sharing and general document size distributions. *Ann. Appl. Probab.* **19**(1), 243–280 (2009)
16. Harris, T.E.: The existence of stationary measures for certain Markov processes. In: *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability, 1954–1955*, vol. II, pp. 113–124. University of California Press, Berkeley (1956)

17. Kelly, F.: The mathematics of traffic in networks. The Princeton Companion to Mathematics (2008). www.statslab.cam.ac.uk/~frank/PAPERS/princeton.html
18. Kelly, F.P.: Charging and rate control for elastic traffic. *Eur. Trans. Telecommun.* **8**, 33–37 (1997)
19. Kelly, F.P., Williams, R.J.: Fluid model for a network operating under a fair bandwidth-sharing policy. *Ann. Appl. Probab.* **14**, 1055–1083 (2004)
20. Kelly, F.P., Maulloo, A.K., Tan, D.K.H.: Rate control in communication networks: shadow prices, proportional fairness and stability. *J. Oper. Res. Soc.* **49**, 237–252 (1998)
21. Lakshminantha, A., Beck, C., Srikant, R.: On the use of sos methods for analysis of connection-level stability in the Internet. In: American Control Conference, 2005. Proceedings of the 2005, vol. 4, pp. 2705–2708 (2005). doi:[10.1109/ACC.2005.1470377](https://doi.org/10.1109/ACC.2005.1470377)
22. Lin, X., Shroff, N.B.: On the stability region of congestion control (2004)
23. Liu, J., Proutiere, A., Yi, Y., Chiang, M., Poor, H.V.: Flow-level stability of data networks with non-convex and time-varying rate regions. *SIGMETRICS Perform. Eval. Rev.* **35**, 239–250 (2007)
24. Massoulié, L.: Structural properties of proportional fairness: Stability and insensitivity. *Ann. Appl. Probab.* **17**, 809–839 (2007)
25. Massoulié, L., Roberts, J.: Bandwidth sharing and admission control for elastic traffic. *Telecommun. Syst.* **15**, 185–201 (1998)
26. Meyn, S.P., Down, D.: Stability of generalized Jackson networks. *Ann. Appl. Probab.* **4**(1), 124–148 (1994)
27. Meyn, S.P., Tweedie, R.L.: Generalized resolvents and Harris recurrence of Markov processes. In: Doeblin and Modern Probability, Blaubeuren, 1991. *Contemp. Math.*, vol. 149, pp. 227–250. Amer. Math. Soc., Providence (1993)
28. Meyn, S.P., Tweedie, R.L.: *Markov Chains and Stochastic Stability*. Springer, London (1993)
29. Meyn, S.P., Tweedie, R.L.: Stability of Markovian processes. II. Continuous-time processes and sampled chains. *Adv. Appl. Probab.* **25**(3), 487–517 (1993). doi:[10.2307/1427521](https://doi.org/10.2307/1427521)
30. Meyn, S.P., Tweedie, R.L.: Stability of Markovian processes. III. Foster–Lyapunov criteria for continuous-time processes. *Adv. Appl. Probab.* **25**(3), 518–548 (1993). doi:[10.2307/1427522](https://doi.org/10.2307/1427522)
31. Mo, J., Walrand, J.: Fair end-to-end window-based congestion control. *IEEE/ACM Trans. Netw.* **8**, 556–567 (2000)
32. Paganini, F., Tang, A., Ferragut, A., Andrew, L.L.H.: Stability of networks under general file size distribution with alpha fair rate allocation. In: Proceedings of the 47th Annual Allerton Conference on Communication, Control, and Computing, Allerton'09, pp. 136–142. IEEE Press, Piscataway (2009). <http://portal.acm.org/citation.cfm?id=1793974.1793997>
33. Srikant, R.: *The Mathematics of Internet Congestion Control*. Birkhäuser, Basel (2004)
34. Ye, H.Q.: Stability of data networks under an optimization-based bandwidth allocation. *IEEE Trans. Autom. Control* **48**(7), 1238–1242 (2003). doi:[10.1007/s11134-011-9233-2](https://doi.org/10.1007/s11134-011-9233-2)