

Fixed-Point Approximations of Bandwidth-Sharing Networks with Rate Constraints

Maria Frolkova
CWI
P.O. Box 94079
NL-1090GB Amsterdam
m.frolkova@cwi.nl

Josh Reed
NYU Stern
Kaufman Management Center
44 West 4th Street
New York, NY 10012
jreed@stern.nyu.edu

Bert Zwart
CWI
EURANDOM
VU University Amsterdam
Georgia Tech
Bert.Zwart@cwi.nl

ABSTRACT

Bandwidth sharing networks are important flow level models of communication networks. We focus on the fact that it takes a significant number of users to saturate a link, necessitating the inclusion of individual rate constraints. In particular we extend work of Reed & Zwart on fluid models of bandwidth sharing with rate constraints under Markovian assumptions: we consider a bandwidth sharing network with rate constraints, where job sizes and deadlines have a general joint distribution. We introduce a fluid model and investigate several of its properties. In particular we show that its invariant point approximates the invariant distribution of the bandwidth sharing network if capacities are large.

Keywords

Bandwidth sharing, flow level Internet model, impatience, measure-valued processes, fluid limits, invariant laws.

1. INTRODUCTION

Bandwidth-sharing networks as considered by Massoulié & Roberts [10, 14] provide a natural framework for modeling the dynamic interaction among competing elastic flows that traverse several links along their source-destination paths, and offer insight in the behavior of communication networks. An extensive effort of contemporary research is devoted to bandwidth-sharing networks. A variety of results related to stability can be found in [15, 2, 11, 9]. Fluid and diffusion limits are considered in [8, 6, 3, 4, 7, 16]. The latter works are important, but ignore the impact of individual peak rate limitations, as has been pointed out by Roberts [13]. Ayesta & Mandjes [1] allow peak rate limitations and construct fluid and diffusion approximations in two specific settings. Reed & Zwart [12] consider a general bandwidth-sharing network with arbitrary topology, where the rate constraints are endogenous part of the network utility maximization procedure.

This paper builds upon [12] by relaxing the stochastic assumptions. In [12], the size B and lead time D of a flow are both exponential and independent, while in the present work, (B, D) can have an arbitrary distribution. Note that it is realistic to assume some dependence between B and D . In particular, it may not make sense for a user to abandon if it is always served at maximum rate. This work is also related to previous work on a single class model with impatience [5] and on bandwidth-sharing networks in overload [4, 3].

We study the model in terms of measure-valued processes.

We propose a fluid model which is shown to arise as the limit of a scaled sequence of stochastic processes. The scaling involves letting the system capacities and input rates grow large. We also show that the fluid model has a unique invariant point in many cases, and construct an example with multiple invariant points, which is a feature that is distinctive from earlier cited works. The ideas behind the proofs are similar to those used in [5, 3, 12]. A new type of result is that we show that the invariant distribution of the scaled stochastic model converges to the invariant point of the fluid model, if the latter is unique. The invariant point of the fluid model can be found by solving a concave programming problem with a polyhedral capacity set. Thus, we establish an approximation of the bandwidth-sharing network that is not only computable in polynomial time, but also valid for arbitrary network topologies, and holds under non-Markovian assumptions.

The paper is organized as follows. A model description is provided in Section 2. In Section 3 we define a fluid model and discuss uniqueness of its solution for a non-zero initial state as well as existence and uniqueness of its invariant solution. The main results on convergence to the fluid model and convergence of invariant laws are presented in Section 4. In Section 5 we discuss performance of a single link and, in particular, give an example of multiple invariant points.

2. MODEL DESCRIPTION

Network parameters Consider a network which consists of a finite number of links labeled by $j \leq J$. Link j has capacity $C_j \in (0, \infty)$. There is a finite number of entities called *routes*, labeled by $i \leq I$. Associated with each is a non-empty subset of links along which traffic offered to the network can be transferred. Let A be a $J \times I$ incidence matrix such that $A_{ji} = 1$ if route j contains link i , and $A_{ji} = 0$ otherwise. Traffic is represented by I classes of flows, class i flows are transferred along route i . If there are multiple flows on a route, then all of them are transferred at the same rate, which is at most $m_i \in (0, \infty)$ for route i . We also assume that, while being transferred, a flow takes simultaneous possession of all links on its route.

Suppose there are z_i flows on route i . Let Λ_i denote the *bandwidth allocated to route i* which is the sum of rates allocated to flows on route i . Then $A\Lambda \leq C$ and $\Lambda \leq m \cdot z$, where Λ , C and m are the vectors of bandwidth allocations, link capacities and rate constraints, and $z = (z_1, \dots, z_I)$ and $m \cdot z = (m_1 z_1, \dots, m_I z_I)$.

Bandwidth allocation policy To each flow on route i

we assign a utility \mathcal{U}_i which is a function of the rate allocated to that flow. We assume that $\mathcal{U}_i(\cdot)$ is strictly increasing and strictly concave in $[0, \infty)$, and twice differentiable in $(0, \infty)$ with $\mathcal{U}'_i(0) = \infty$. Then, given a population z of flows in the network, the bandwidth allocation vector $\Lambda(z)$ is determined as the unique solution to the following optimization problem:

$$\max_{\substack{\Lambda: \mathbb{A}\Lambda \leq C, \\ \Lambda \leq m \cdot z}} \sum_{i=1}^I z_i \mathcal{U}_i(\Lambda_i/z_i), \quad (1)$$

where $\Lambda_i/0 := 0$. By $\mathcal{U}'_i(0) = \infty$, we have $\Lambda_i(z) > 0$ if $z_i > 0$. Also the function $\Lambda(\cdot)$ is Lipschitz continuous in any compact set $[\delta, \Delta]^I \subset (0, \infty)^I$ (see [12]) and continuous in \mathbb{R}_+^I .

Dynamic assumptions Suppose at time 0 there are Z_i^0 initial flows on route i , $Z_i^0 < \infty$ a.s. New flows arrive on route i according to a Poisson process $E_i(\cdot)$ with rate η_i , and U_{ik} denotes the arrival time of flow k on route i . R.v.'s (B_{il}^0, D_{il}^0) and (B_{ik}, D_{ik}) represent the initial size and the initial lead time of initial flow l and of new flow k on route i respectively. A flow abandons the network as soon as it has been transferred or its lead time has run out, depending on what happens earlier. For each i , $\{(B_{il}^0, D_{il}^0)\}_{l=1}^{\infty}$ are $(0, \infty)^2$ -valued r.v.'s, and $\{(B_{ik}, D_{ik})\}_{k=1}^{\infty}$ are i.i.d. copies of a $(0, \infty)^2$ -valued r.v. (B_i, D_i) with a finite mean value. We allow dependence between B_i and D_i , and assume that $D_i \geq B_i/m_i$ a.s., which in particular implies that a flow will not abandon when transferred at maximum rate. The stochastic primitives $E_i(\cdot)$, $\{(B_{ik}, D_{ik})\}_{k=1}^{\infty}$, $i \leq I$, are mutually independent.

Time evolution of the network For each $t \geq 0$, let $Z(t) = (Z_1(t), \dots, Z_I(t))$ denote the population of flows in the network at time t . We call $Z(\cdot)$ the *queue-length process*. For all $t \geq s$, let $S_i(s, t)$ denote the cumulative bandwidth allocated per flow on route i during time interval $[s, t]$. Since the bandwidth allocated to a route is shared equally by all flows on that route, we have $S_i(s, t) = \int_s^t \Lambda_i(Z(u))/Z_i(u) du$, where the integrand is defined to be 0 when $Z_i(u) = 0$. For $x \in \mathbb{R}$, let $x^+ = \max\{x, 0\}$. The *residual size* and the *residual lead time* at time t of initial flow $l \leq Z_i^0$ on route i , and those of flow $k \leq E_i(t)$ on route i are given by: $B_{il}^0(t) = (B_{il}^0 - S_i(0, t))^+$, $D_{il}^0(t) = (D_{il}^0 - t)^+$, and $B_{ik}(t) = (B_{ik} - S_i(U_{ik}, t))^+$, $D_{ik}(t) = (D_{ik} - (t - U_{ik}))^+$.

We now introduce a measure-valued process that we call the *state descriptor*, and that keeps track of the residual sizes and the residual lead times of all flows in the network. For $(x, y) \in \mathbb{R}_+^2$, let $\delta_{(x, y)}^+$ denote the Dirac point measure at (x, y) if $\min\{x, y\} > 0$, otherwise $\delta_{(x, y)}^+$ is zero measure. Then the state of the network at time t is represented by the vector of random measures $\mathcal{Z}(t) = (\mathcal{Z}_1(t), \dots, \mathcal{Z}_I(t))$, where

$$\mathcal{Z}_i(t) = \sum_{l=1}^{Z_i^0} \delta_{(B_{il}^0(t), D_{il}^0(t))}^+ + \sum_{k=1}^{E_i(t)} \delta_{(B_{ik}(t), D_{ik}(t))}^+. \quad (2)$$

The total mass of the state descriptor coincides with the queue length, $Z(t) = \mathcal{Z}(t)(\mathbb{R}_+^2) = (\mathcal{Z}_1(t)(\mathbb{R}_+^2), \dots, \mathcal{Z}_I(t)(\mathbb{R}_+^2))$.

3. FLUID MODEL

Existence and uniqueness In this section we define a deterministic fluid model that later will be shown to arise as the limit of the stochastic model described in the previous section under a proper scaling (and that will imply existence of the fluid model). Let \mathbf{M} denote the set of finite non-negative Borel measures on \mathbb{R}_+^2 , endowed with the weak

topology. Let $\zeta^0 = (\zeta_1^0, \dots, \zeta_I^0) \in \mathbf{M}^I$ be such that, for all i , the projections $\zeta_i^0(\cdot \times \mathbb{R}_+)$ and $\zeta_i^0(\mathbb{R}_+ \times \cdot)$ are free of atoms in \mathbb{R}_+ . Put $z^0 = \zeta^0(\mathbb{R}_+^2)$, and let (B_i^0, D_i^0) be a r.v. with distribution ζ_i^0/z_i^0 if $z_i^0 > 0$, otherwise let $(B_i^0, D_i^0) = (0, 0)$ a.s. For all i , take r.v.'s (B_i, D_i) as defined in Section 2. Then a continuous function $\zeta(\cdot): [0, \infty) \rightarrow \mathbf{M}^I$ is called a *measure-valued fluid model solution (m.v.f.m.s.) with initial state ζ^0* if, for $i \leq I$, all $t \geq 0$ and all $(x, y) \in \mathbb{R}_+^2$,

$$\begin{aligned} \zeta_i(t)((x, \infty) \times [y, \infty)) &= z_i^0 \mathbb{P}\{B_i^0 \geq x + s_i(0, t), D_i^0 \geq y + t\} \\ &+ \eta_i \int_0^t \mathbb{P}\{B_i \geq x + s_i(s, t), D_i \geq y + t - s\} ds, \end{aligned} \quad (3)$$

where $s_i(s, t) = \int_s^t \Lambda_i(z(u))/z_i(u) du$ and $z(\cdot)$ is the total-mass function, $z(t) = \zeta(t)(\mathbb{R}_+^2)$. The function $z(\cdot)$ is called simply a *fluid model solution (f.m.s.) with initial state ζ^0* .

Note that uniqueness of a f.m.s. is equivalent to uniqueness of a m.v.f.m.s. because they are uniquely defined by each other. In the following theorem we show uniqueness of a f.m.s. with a non-zero and Lipschitz continuous initial state.

THEOREM 1. *Suppose that $\zeta_i^0(\mathbb{R}_+^2) > 0$, $i \leq I$. Suppose further that there exists a constant $L \in (0, \infty)$ such that, for all i , all $x < x'$ and all y , $\zeta_i^0([x, x'] \times [y, \infty)) \leq L|x - x'|$. Then a (measure-valued) f.m.s. with initial state ζ^0 is unique.*

The first assumption implies that f.m.s.'s are bounded away from 0 in all coordinates. They are also bounded from above. Then Lipschitz continuity of the rate allocation function $(\Lambda_1(z)/z_1, \dots, \Lambda_I(z)/z_I)$ in any compact set $[\delta, \Delta]^I \subset (0, \infty)^I$, and Lipschitz continuity of the initial state ζ^0 imply that a f.m.s. must be unique.

Invariant point To study invariant points, we consider only f.m.s.'s because m.v.f.m.s.'s are uniquely defined by them. Invariant f.m.s.'s are given by the equations $\Lambda_i = g_i(\Lambda_i/z_i)$, $i \leq I$, where $g_i(x) = \eta_i \mathbb{E} \min\{B_i, xD_i\}$.

THEOREM 2. *Suppose that, for $i = 1, \dots, I$, the left most point of the support of distribution of D_i/B_i is $1/m_i$. Then an invariant (measure-valued) f.m.s. exists and is unique.*

Let $\rho_i = \eta_i \mathbb{E} B_i$, $\rho = (\rho_1, \dots, \rho_I)$. Under the conditions of the theorem, the functions $g_i^{-1}(\cdot)$ are strictly increasing in $[0, \rho_i]$. For any invariant f.m.s. z , $\Lambda(z)$ is a solution to the following optimization problem:

$$\max_{\substack{\Lambda: \mathbb{A}\Lambda \leq C, \\ \Lambda \leq \rho}} \sum_{i=1}^I G_i(\Lambda_i), \quad (4)$$

where functions $G_i(\cdot)$ are such that $G'_i(x) = \mathcal{U}'_i(g_i^{-1}(x))$. By strict concavity of (4), $\Lambda(z) = \Lambda^*$ is unique, and the unique invariant f.m.s. is given by $z_i^* = \Lambda_i^*/g_i^{-1}(\Lambda_i^*)$, $i \leq I$.

4. LARGE CAPACITY SCALING

Convergence to the fluid model With large capacity scaling, we let global parameters of the network, link capacities and arrival rates, grow to ∞ , while characteristics of an individual flow remain of a fixed order. More precisely, consider a sequence of stochastic models as defined in Section 2, indexed by $n \in \mathbb{N}$. Let capacities and arrival rates grow linearly in n , and rate constraints, generic flow sizes and lead times be the same in all models: $C^n = nC$, $\eta^n = n\eta$, $m^n = m$, $(B_i^n, D_i^n) = (B_i, D_i)$ for all i . For all n , introduce the fluid scaled versions of the queue-length process $Z^n(\cdot)$ and the state descriptor $\mathcal{Z}^n(\cdot)$ of the n -th model:

$\bar{Z}^n(\cdot) = Z^n(\cdot)/n$ and $\bar{Z}^n(\cdot) = Z^n(\cdot)/n$. The processes $\bar{Z}^n(\cdot)$ and $\bar{Z}^n(\cdot)$ take values in the Skorokhod spaces $D([0, \infty), \mathbb{R}_+^I)$ and $D([0, \infty), \mathbf{M}^I)$ of right-continuous functions with left limits. Finally, we make an assumption about the initial conditions $Z^n(0)$. For $\zeta \in \mathbf{M}^I$, put $\langle \chi, \zeta \rangle = (\int x_1 d\zeta_1, \int x_2 d\zeta_1, \dots, \int x_1 d\zeta_I, \int x_2 d\zeta_I)$. Let $(\bar{Z}^n(0), \langle \chi, \bar{Z}^n(0) \rangle) \Rightarrow (\zeta^0, \langle \chi, \zeta^0 \rangle)$ as $n \rightarrow \infty$, where $\zeta^0 \in \mathbf{M}^I$ and $\langle \chi, \zeta^0 \rangle$ is finite, and the projections $\zeta_i^0(\cdot \times \mathbb{R}_+)$ and $\zeta_i^0(\mathbb{R}_+ \times \cdot)$ are free of atoms in \mathbb{R}_+ for all i . We have the following fluid limit result.

THEOREM 3. *Suppose that the conditions of Theorem 1 hold. Then the sequence $\{\bar{Z}^n(\cdot)\}_{n=1}^\infty$ as defined above converges in distribution to the unique measure-valued f.m.s. with initial state ζ^0 .*

By means of a technique developed in [6], we show that the sequence $\{\bar{Z}^n(\cdot)\}_{n=1}^\infty$ is tight. Due to rate constraints, for any weak limit point $\tilde{Z}(\cdot)$, its total mass function is bounded away from 0 outside $t = 0$, which allows to show that $\tilde{Z}(\cdot)$ satisfies the fluid model equation (3) a.s.

Convergence of invariant laws Since the total population of the network (as defined in Section 2) is bounded from above by the length of an M/G/ ∞ queue with generic service time $D = \max D_i$, the state descriptor is regenerative and has a unique stationary distribution. Consider a sequence of stochastic models with parameters $C^n = nC$, $\eta^n = n\eta$, $m^n = m$, $(B_i^n, D_i^n) = (B_i, D_i)$ for all i . Let \mathcal{Y}^n have the stationary distribution of the state descriptor of the n -th model, and put $\bar{\mathcal{Y}}^n = \mathcal{Y}^n/n$.

THEOREM 4. *Suppose that, for all i , B_i has a bounded density. Suppose also that the condition of Theorem 2 holds. Then the sequence $\{\bar{\mathcal{Y}}^n\}_{n=1}^\infty$ converges in distribution to the unique invariant measure-valued f.m.s.*

First we check that a criteria of tightness holds for the sequence $\{\bar{\mathcal{Y}}^n\}_{n=1}^\infty$. Then consider a weak limit point $\tilde{\mathcal{Y}}$. By boundedness of flow size densities and by M/G/ ∞ upper bounds, $\tilde{\mathcal{Y}}$ is Lipschitz continuous in both coordinates. For each n , run the n -th network starting from $\bar{Z}^n(0) = \bar{\mathcal{Y}}^n$. Then Theorem 3 implies that $\bar{Z}^n(\cdot)$ converges to the f.m.s. with initial state $\tilde{\mathcal{Y}}$. Since all $\bar{Z}^n(\cdot)$ are stationary processes, the limit is stationary, too. Hence $\tilde{\mathcal{Y}}$ is an invariant f.m.s.

5. EXAMPLES

A single link in overload Let $I = C_1 = 1$ and the utilities $\mathcal{U}_i(x) = \kappa_i \log x$. Assume that the network is overloaded ($\sum_{i=1}^K \rho_i > 1$), and that the condition of Theorem 2 holds. Given the network is in equilibrium, which classes are served at the full rate ($\Lambda_i = \rho_i$) and which are not ($\Lambda_i < \rho_i$)? Suppose $\kappa_1/m_1 \geq \dots \geq \kappa_I/m_I$. Let x^* be the unique solution to $\sum_{i=1}^K g_i(\kappa_i x^*) = 1$, and let $i^* = \min\{i: \kappa_i x^* < m_i\}$. By the KKT conditions for (1), the unique invariant point is given by: $z_i = \rho_i / \min\{m_i, \kappa_i x^*\}$ for $i < i^*$, $z_i = g_i(\kappa_i x^*) / (\kappa_i x^*)$ for $i \geq i^*$; and $\Lambda_i = \rho_i$ for $i < i^*$, $\Lambda_i < \rho_i$ for $i \geq i^*$.

Multiple invariant points In the previous example assume critical load ($\sum_{i=1}^I \rho_i = 1$) and that, for all i , $D_i = a_i B_i$ with $a_i \geq 1/m_i$. Then in the invariant point equation $\Lambda_i = \rho_i$ are unique. If $a_i = 1/m_i$ for all i , then the invariant point $z = (\rho_1/m_1, \dots, \rho_I/m_I)$ is unique, too. Let $a_1 > 1/m_1$, which violates the condition of Theorem 2. If the set $S_1 = [\kappa_1/m_1, \kappa_1 a_1] \cap_{i>1} (0, \kappa_i/m_i]$ is non-empty, then, for any $p \in S_1$, $z = (p\rho_1/\kappa_1, \rho_2/m_2, \dots, \rho_I/m_I)$ is an invariant point.

6. REFERENCES

- [1] Ayesta, U., Mandjes, M. Bandwidth-sharing networks under a diffusion scaling. *Ann. Oper. Res.* **170**, 41–58, 2009.
- [2] Bonald, T., Massoulié, L. Impact of fairness on Internet performance. In: *Proc. ACM Sigmetrics & Performance 2001 Conf.*, Boston MA, USA, 82–91.
- [3] Borst, S.C., Egorova, R., Zwart, B. Fluid limits of bandwidth-sharing networks in overload. Submitted, 2009.
- [4] Egorova, R., Borst, S.C., Zwart, B. Bandwidth-sharing networks in overload. *Perform. Evaluation* **64**, 978–993, 2007.
- [5] Gromoll, H.C., Robert, Ph., Zwart, B. Fluid limits for processor sharing queues with impatience. *Math. Oper. Res.* **33**, 375–402, 2008.
- [6] Gromoll, H.C., Williams, R.J. Fluid limits for networks with bandwidth sharing and general document size distributions. *Ann. Appl. Probab.* **19**, 243–280, 2009.
- [7] Kang, W.N., Kelly, F.P., Lee, N.H., Williams, R.J. State space collapse and diffusion approximation for a network operating under a fair bandwidth-sharing policy. *Ann. Appl. Probab.* **19**, 1719–1780, 2009.
- [8] 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.
- [9] Massoulié, L. Structural properties of proportional fairness: stability and insensitivity. *Ann. Appl. Probab.* **17**, 809–839, 2007.
- [10] Massoulié, L., Roberts, J.W. Bandwidth sharing: objectives & algorithms. In: *Proc. IEEE Infocom '99*, New York NY, USA, 1395–1403.
- [11] Mo, J., Walrand, J. Fair end-to-end window-based congestion control. *IEEE/ACM Trans. Netw.* **8**, 556–567, 2000.
- [12] Reed, J., Zwart, B. Limit theorems for bandwidth sharing networks with rate constraints. Submitted, 2010.
- [13] Roberts, J.W. A survey on statistical bandwidth sharing. *Comput. Netw.* **45**, 319–332, 2004.
- [14] Roberts, J.W., Massoulié, L. Bandwidth sharing and admission control for elastic traffic. In: *Proc. ITC Specialist Seminar*, Yokohama, Japan, 1998.
- [15] De Veciana, G., Lee, T.-L., Konstantopoulos, T. Stability and performance analysis of networks supporting elastic services. *IEEE/ACM Trans. Netw.* **9**, 2–14, 2001.
- [16] Ye, H., Yao, D.D. Heavy-traffic optimality of a stochastic network under utility-maximizing resource control. *Oper. Res.* **56**, 453–470, 2008.