Concave switching in single-hop and multihop networks

Neil Walton¹

Received: 21 July 2014 / Revised: 17 April 2015 / Published online: 23 May 2015

© The Author(s) 2015. This article is published with open access at Springerlink.com

Abstract Switched queueing networks model wireless networks, input-queued switches, and numerous other networked communications systems. We consider an (α, g) -switch policy; these policies provide a generalization of the MaxWeight policies of Tassiulas and Ephremides (IEEE Trans Autom Control 37(12):4936–1948, 1992) and the weighted α -fair with allocations of Mo and Walrand (IEEE/ACM Trans Netw 8(5):556–567, 2000) which are typically applied to Bandwidth Sharing Networks (Massoulié and Roberts in IEEE/ACM Trans Netw 10(3):320–328, 2002). For single-hop switch networks, we prove the maximum stability property for this class of randomized policies. Thus these policies have the same first-order behavior as the MaxWeight policies. However, for multihop networks some of these generalized policies address a number of critical weakness of the MaxWeight/BackPressure policies. For multihop networks with fixed routing, we consider a policy called the Proportional Scheduler (or $(1, \log)$ -policy). In this setting, the BackPressure policy is maximum stable, but must maintain a queue at each node for every route destination, which typically grows rapidly with a network's size. However, the Proportional Scheduler only needs to maintain a queue for each outgoing link, which is typically bounded in number. As is common with Internet routing, by maintaining per-link queueing, each node only needs to know the next hop for each packet and not its entire route. Further, in contrast to BackPressure, the Proportional Scheduler does not compare downstream queue lengths to determine weights; only local link information is required. This leads to greater potential for decomposed implementations of the policy. Through a reduction argument and an entropy argument, we demonstrate that,

---

✉ Neil Walton
n.s.walton@uva.nl

¹ University of Amsterdam, Science Park 904, 1098 XH Amsterdam, The Netherlands

while maintaining substantially less queueing overhead, the Proportional Scheduler achieves maximum throughput stability.

Keywords Proportional fairness · MaxWeight · BackPressure · Switch network · Utility maximization · Maximum stable

Mathematics Subject Classification 60K25

1 Introduction

The MaxWeight/BackPressure scheduling policies were first introduced by Tassiulas and Ephremides as a model of wireless communication [36]. Their policy was applicable to the class of switched queueing networks. Here there are constraints on which queues can be served simultaneously. Subsequently, to analyze Internet Protocol routers, McKeown et al. [25] applied this policy and paradigm to the example of input-queued switches. Extensions of these policies can be found in Stolyar [34]. The MaxWeight and BackPressure policies have proved popular and, accordingly, have been generalized [1, 12, 15, 31–34]. The defining feature of each generalization of MaxWeight and BackPressure is that it minimizes the drift of a certain Lyapunov function. This Lyapunov function is then used to analyze the policy's stability properties given the queueing network's arrival process.

Switched networks can be classed as either single-hop or multihop. In single-hop switched networks, packets depart the network after service at a queue, while in multihop networks, packets depart after receiving service at one or more queues. When applied to single-hop networks, the policy of Tassiulas and Ephremides is called *MaxWeight* and in multihop networks it is referred to as *BackPressure*.

Both MaxWeight and BackPressure have the key property of having a maximal stability region while not requiring explicit estimation of traffic arrival rates. Roughly stated, a queue scheduling policy is maximum stable if it is positive recurrent for every arrival rate for which there exists a stabilizing policy. If the vector of arrival rates is known, then maximum stability is trivial: at each time, one can choose a random schedule whose average service rate dominates each arrival rate. However, in practice, explicit knowledge of arrival rates is not achievable, particularly when these rates may vary over time. A striking feature of the MaxWeight/BackPressure models is that, in addition to being maximum stable, it is *myopic*—meaning that only current state information is required to make a scheduling decision.

A further class of maximum stable Internet models were first introduced by Massoulié and Roberts [24]. When specifically applied as flow level Internet models, these systems have been called Bandwidth Sharing Networks, though their applicability is potentially broader than internet Bandwidth Allocation. Similar to MaxWeight and BackPressure, policies on these networks are often defined by an optimization and are maximum stable; however, unlike BackPressure, they are not constructed from a Lyapunov function. Stability proofs for these systems can be found in Bonald and Massoulié [3] and Massoulié [23]. Within this class of networks, we will be particularly interested in a *proportional fair* allocation; see Kelly [19]. Stability results on

proportional fairness can be found in [10, 23]. Further progress on the large deviations and heavy traffic behavior of proportional fairness can be found in Jonckheere and López [18] and Vlasiov, Zhang, Zwart [37]. The paper generalizes the MaxWeight policies by allowing for resource sharing similar to that found in Bandwidth Sharing Networks. To the knowledge of this author, discussions applying these models to switch networks were first made by Shah and Wischik [32] and Zhong [38] for the weighted α -fair bandwidth allocation policies of Mo and Walrand [26].

This paper generalizes the MaxWeight to a new class of policies called the (α, g) -switch policies. We prove stability of these policies for single-hop switch networks. We show that the generalized policies can be applicable to multihop networks. In particular, for a proportionally fair case, we consider a policy called the Proportional Scheduler. We prove that the Proportional Scheduler is maximum stable for multihop networks with fixed routing. In essence, the paper leverages and generalizes results to switch networks that originate from Bonald and Massoulié [3] and Massoulié [23]. We can then advocate the potential applicability of this paradigm by comparison with MaxWeight/BackPressure policies. In particular, because of their substantially reduced queueing complexity, there are a number of significant structural advantages which have not previously been observed in switch systems.

1.1 Policies and results

We now briefly describe the switch policies considered in this paper. A more formal description is given in Sect. 3.

A single-hop switch network consists of a (finite) set of queues, $\mathcal{J}$. In a single-hop network, a packet departs the network after service. In this discrete-time model, there are constraints on which queues can be served simultaneously. The set of permissible schedules are given by the finite set $\mathcal{S} \subset \mathbb{Z}_+^{\mathcal{J}}$. Further, we allow for randomization among these discrete choices. We let $\langle \mathcal{S} \rangle \subset \mathbb{R}_+^{\mathcal{J}}$ be the convex hull of the set $\mathcal{S}$.

We define a switch policy as follows: given a parameter $\alpha > 0$ and concave functions $g = (g_j : j \in \mathcal{J})$, the (α, g) -switch policy chooses a random schedule which solves

$$\text{maximize} \quad \sum_{j \in \mathcal{J}} g_j(s_j) Q_j^\alpha \quad \text{over} \quad s \in \langle \mathcal{S} \rangle,$$

where $Q = (Q_j : j \in \mathcal{J})$ is the vector of queue sizes and where, allowing for randomization, the maximization is taken over a set of admissible schedules $\langle \mathcal{S} \rangle$. The interior of the set $\langle \mathcal{S} \rangle$ also gives the network's *stability region*—the largest set of arrival rates for which a policy can stabilize the network. A policy which is stable for all arrival rates in $\langle \mathcal{S} \rangle$ is maximum stable. In order to place these policies on a par with the traditional MaxWeight policies, we show that these policies have the same maximum stability property as MaxWeight. In Theorem 1, we provide a proof which is analogous to the proof of Bonald and Massoulié for weighted α -fair policies [3]. This is the first contribution of the paper.

Next we consider multihop networks with fixed routing. In a multihop network, after service a packet can rejoin the network at a queue as specified by the packet's

route. Further, one can think of a multihop network as a set of nodes $\mathcal{N}$, and the set $\mathcal{J}$ as a series of links between these nodes, $j = (n, n')$ for $n, n' \in \mathcal{N}$. In this setting, BackPressure is the canonical maximal stable policy. The BackPressure policy essentially consists of two stages: the first stage, where weights are determined for each link j by comparing current and downstream stream queues sizes for each commodity/route, and a second stage, where we optimize these weights over the set of schedules. Notice, for BackPressure to make a scheduling decision, the policy must know the routes of all packets present at each node. This is not practical when the number of routes processed by a link is large in number, or when the routes processed by the link are not completely known to that link.

Given these structural issues, we are interested in policies that employ *per-link queueing*; that is, policies only require information about the number of packets wishing to traverse a link j , rather than *per-route queueing* policies—like BackPressure—which require information on the number of packets from each route at the link.

In this context, we analyze a Proportional Scheduler. This consists of solving the following proportional fair optimization:

$$\text{maximize} \quad \sum_{j \in \mathcal{J}} Q_j \log s_j \quad \text{over} \quad s \in \langle S \rangle.$$

Here $j = (n, n')$ represents a directed link between two nodes of a communication network and Q_j is the number of packets present at n which are currently waiting to pass through link j . Then,

a schedule is then chosen whose mean is the solution to this optimization and a job is chosen at random from each scheduled queue.

Notice this is a special case of the (α, g) -switch policies above: specifically, the $(1, \log)$ -switch policy. The Proportional Scheduler applies per-link queueing. So, we could use the terms *link* and *queue* interchangeably. However, since this equivalence does not hold for BackPressure, we use the term *link* when referring to index $j \in \mathcal{J}$.

We consider the fluid model for this fixed route multihop network. We show that the Proportional Scheduler has a fluid model which is maximum stable. We give two proofs of this result. Each emphasizes different features of the policy:

1. The first exploits the structure of our random service discipline and use the properties of the logarithm to reduce our multiclass queueing system with fixed routing to a single-class system with Markovian routing. The reduction shows that per-link and per-route queueing are equivalent for the Proportional Scheduler. From this reduction, we can apply the fluid analysis of Massoulié [23] to prove fluid stability.
2. We provide a second and direct proof that emphasizes the policy's implicit ability to balance the entropy between separate networked components. The fluid stability argument given is compared with the proof of Bramson [4] for head-of-the-line processor sharing networks.

These two proofs originate from two different queueing networks: first, bandwidth sharing networks and, second, classical queueing networks. However, the importance of the result is its consequence for the third setting of switched networks.

1.2 Contributions

Let us discuss in more detail a few of the advantages of this Proportional Scheduler for multihop switched networks. Firstly, as we consider increasingly large communication networks, we expect the degree of nodes (routers) to stay bounded, but the number of routes or destinations processed by a router to grow substantially. Thus the Proportional Scheduler, which maintains a queue per outgoing link, requires a number of queues that is far smaller than BackPressure, which maintains a queue for each route destination at each node. Second, to make a scheduling decision with BackPressure, we need to know the entire route, or at least the destination, of each packet, while for the Proportional Scheduler, knowing the next hop of the packet is sufficient to implement the policy. Third, if we add new nodes to our network then, to implement BackPressure, the newly added node must be aware of the entire topology and route/class structure of the network, while the Proportional Scheduler only requires information about neighboring links. So, routing decisions can be determined by packets. Fourth, when different schedules within the network can be implemented independently—here we imagine separate non-interfering network components—for BackPressure, messages must be shared between the separate components in order to determine the differences in queue sizes. However, for the Proportional Scheduler, the optimization relevant to the policy will decompose into an independent optimization for each component which can then be solved separately. Finally, as a general rule, there are almost no single-hop communication systems. Communication systems consist of a large number of interconnected components. Single-hop switch networks are often considered in performance analysis for reasons of tractability. This statement is particularly true for wireless systems. However, stability in a single-hop system does not in general imply stability in a corresponding multihop system. See Lu and Kumar [22], Rybko and Stolar [28] and Bramson [5] for instability examples which can easily be adapted to switch systems. Notable, recent work of Dieker and Shin [11] considers schemes which extend local/single-hop stability results to global stability results. These allow for multiclass routing. Further recent work of Ji, Joo, and Shroff [16] considers stability results for switched networks where packets are queued per-link rather than per-route, as is typically applied for BackPressure. Here stability is achieved by running an appropriate queueing system in the memory of the algorithm. By estimating queue sizes and loads in this way, stability is achievable for switched networks as long as routes do not form a loop. However, these policies require route information from each packet to make scheduling decisions (like BackPressure) or require estimation of long-run arrival rates (which may in practice vary over time). To the best of the knowledge of this author, this is the first general proof of maximum stability for a multihop switch network where neither routing information nor estimation of arrival rates are required in order to make a scheduling decision.

We now briefly discuss extensions and future consequences of this analysis. The Proportional Scheduler which we consider applies a processor-sharing discipline within each queue, meaning that the scheduled packet from a queue is chosen at random. However, it would be desirable to consider first-in-first-out queueing. In this way, packets will be served by the network in sequence and jitter effects will be reduced.

Surprisingly, the Proportional Scheduler is provably maximum stable for FIFO queueing disciplines. The proof of this result is substantially more involved than the short proof that we can provide for the processor-sharing case. The FIFO case follows due to underlying connections between proportional fairness and quasi-reversible queueing systems—see [23,30] for some related discussion. This analysis leads to further structural advantages, namely product-form resource-pooling effects and better delay scaling complexity compared to BackPressure. We refer the interested reader to the preprint (Bramson M, D’Auria B, Walton N, Proportional switching in fifo routed networks “in preparation”).

In summary, the contributions of this work are as follows:

- We generalize the set of MaxWeight policies and weighted α -fair Bandwidth allocation policies to provide a new class of switch policies, and we prove maximum stability for these policies.
- Using a random discipline within queues, we consider the Proportional Scheduler in multihop networks with fixed routing. We give two proofs of fluid stability for this system.
- We emphasize several important consequences of this last result for switch networks:
  - the policy only needs to maintain a queue for each outgoing link;
  - packets can be routed by only knowing each packet’s next hop;
  - network nodes do not need to know their network topology; and
  - messages do not need to be sent to calculate weights and so scheduling decisions can be completely decomposed between independently functioning components.
- We provide the first general proof of maximum stability for a multihop switch network where routing information is not required in order to make a scheduling decision.

Informally speaking, because the Proportional Scheduler does not need to know the route structure of the network to make a scheduling decision, while BackPressure does, the policy’s implementation scales better with network size.

1.3 Organization

The remaining sections of the paper are organized as follows. In Sect. 2, we describe the network and queueing process for both single-hop and multihop networks. In Sect. 3, we define the (α, g) -switch policies, proportional fairness, and we discuss how these relate to existing policies, namely MaxWeight, BackPressure, utility optimizing bandwidth allocation policies, and weighted α -fairness. We also define the fluid model associated with each of these systems. In Sect. 4, we present and discuss the main results of this paper: Theorem 1, which proves maximum stability for the (α, g) -switch policies in single-hop switched networks; and Theorem 2, which proves maximum stability for the Proportional Scheduler in multihop switched networks. In Sect. 5, we provide a proof of Theorem 1. This consists of characterizing the fluid limit of the (α, g) -policy, proving stability of that fluid system and then using this to prove positive recurrence of the prelimit process. In Sect. 6, we give two proofs of Theorem 2. The

first involves a reduction to a proof of Massoulié [23]. The second proof can be argued from the fluid analysis of Bramson [4].

2 Switch network notation

We define single-hop and multihop switch networks. These discrete-time queueing networks have restrictions on which queues can be served simultaneously.

The following notation is used both for single-hop and multihop networks. We assume that time is slotted; that is, each time, t , belongs to the positive integers, $\mathbb{Z}_+$. We let the finite set $\mathcal{J}$ index the set of queues or links of a network. We let the finite set $\mathcal{S}$ be the set of schedules. Each schedule $\sigma = (\sigma_j : j \in \mathcal{J}) \in \mathcal{S}$ is a vector in $\mathbb{Z}_+^{\mathcal{J}}$, where σ_j gives the number of jobs that will be served from queue j under schedule σ . We assume the zero vector, $\mathbf{0}$, belongs to $\mathcal{S}$. We let $\sigma_{\max}$ give the maximum value of σ_j for $\sigma \in \mathcal{S}$. We let $\langle \mathcal{S} \rangle$ be the convex combination of schedules in $\mathcal{S}$. We assume $\langle \mathcal{S} \rangle$ has a non-empty interior. For real numbers $q, s \in \mathbb{R}$, we define $q \wedge s = \min\{q, s\}$. For vectors $q, s \in \mathbb{R}^{\mathcal{J}}$, we define $q \wedge s = (q_j \wedge s_j : j \in \mathcal{J})$. For a vector of queue sizes $Q \in \mathbb{Z}_+^{\mathcal{J}}$, we define

$$\mathcal{S}_Q = \{\sigma \wedge Q : \sigma \in \mathcal{S}\}.$$

Given that not all jobs can be served from a queue j when $Q_j < \sigma_j$, the set $\mathcal{S}_Q$ gives the set of schedules available given the vector of queue sizes Q . We let $\langle \mathcal{S}_Q \rangle$ be the convex combination of points in $\mathcal{S}_Q$.

2.1 Single-hop switched network

For *single-hop networks* once a packet has been served at its queue, it departs the network.

We let $a(t) = (a_j(t) : j \in \mathcal{J}) \in \mathbb{Z}_+^{\mathcal{J}}$ be the number of arrivals occurring at each queue at time $t \in \mathbb{Z}_+$. We assume $\{a(t)\}_{t=0}^\infty$ is a sequence of independent identically distributed random vectors with finite mean $\bar{a} \in (0, \infty)^{\mathcal{J}}$ and finite second moment $\mathbb{E}[a_j(t)^2] \leq K$ for all $j \in \mathcal{J}$.¹

We let $Q(0) = (Q_j(0) : j \in \mathcal{J})$ be the number of jobs in each queue at time $t = 0$. From a sequence of schedules $\{\sigma(t)\}_{t=1}^\infty$, we can define the queue size vector $Q(t) = (Q_j(t) : j \in \mathcal{J})$ by

$$Q_j(t+1) = Q_j(t) - \sigma_j(t+1) + a_j(t+1), \quad (1)$$

for $j \in \mathcal{J}$, and $t \in \mathbb{N}$. We assume that service occurs before arrivals. So, for positive queue sizes, it is required that $\sigma(t+1) \in \mathcal{S}_{Q(t)}$, for all $t \in \mathbb{Z}_+$.²

¹ The assumption of finite variance is not essential; however, it allows for bounds more convenient for our proofs.

² Observe, this choice of notation is equivalent to defining the queueing process by $Q_j(t+1) = [Q_j(t) - \sigma_j(t+1)]_+ + a_j(t+1)$.

Given $\{Q(t)\}_{t=0}^\infty$ defines a Markov chain, we say the queue size process is *maximum stable* if it is positive recurrent whenever the vector of arrival rates, $\bar{a}$, belongs to the interior of $\langle \mathcal{S} \rangle$.

2.2 Multihop switched network

In a multihop network, we allow packets to visit multiple queues within the network before departing. We will develop our previous notation in a similar way to Tassiulas and Ephremides [36].

As in the previous section, we let the finite set $\mathcal{J}$ index the set of queues or links of a network. We let $\mathcal{N}$ define the set of nodes of a network. We consider each link $j \in \mathcal{J}$ to be a directed edge in this network, $j = (n, n') \in \mathcal{N} \times \mathcal{N}$. We let $\mathcal{R}$ be a set of routes through the network. Here each $r = (n_1^r, \dots, n_{k_r}^r) \in \mathcal{R}$ is an ordered set of nodes such that $j_k^r := (n_k^r, n_{k+1}^r) \in \mathcal{J}$ for $k = 1, \dots, k_r - 1$. A route- r packet served at node n_k must next visit node n_{k+1} using link j_k^r . We assume that the final destination node, $n_{k_r}^r$, on a route r is used uniquely by that route. We apply the notation $j \in r$ if link j is part of route r . Further, for each $j \in r$, we let j_-^r and j_+^r denote the previous (upstream) and the next (downstream) link on route r .

We let $a(t) = (a_{jr}(t) : j \in \mathcal{J}, r \in \mathcal{R}) \in \mathbb{Z}_+^{\mathcal{J} \times \mathcal{R}}$ be the number of external arrivals occurring at the ingress $j = j_1^r$ of each route r at each time $t \in \mathbb{Z}_+$. Note $a_{jr}(t) = 0$ unless the queue is the first on its route, $j = j_1^r$. We assume $\{a(t)\}_{t=0}^\infty$ is a sequence of independent identically distributed random vectors with finite mean $\bar{a} = (\bar{a}_{jr} : r \in \mathcal{R}, j \in \mathcal{J}) \in \mathbb{R}_+^{\mathcal{J} \times \mathcal{R}}$ and uniformly bounded second moment $\mathbb{E}[a_{jr}(t)^2] \leq K$ for all $r \in \mathcal{R}, j \in \mathcal{J}$. We define the load on queue $j \in \mathcal{J}$ by

$$\bar{a}_j = \sum_{r: j \in r} \bar{a}_{jr}. \quad (2)$$

Also we let $\bar{a}_r := \bar{a}_{j_1^r r}$ be the arrival rate of route r .

We let $Q_j(t)$ denote the number of packets whose next service is at link j . We let $X_{jr}(t)$ be the number of route r packets who will next be served over link j at time t . Thus, it holds that

$$Q_j(t) = \sum_{r: j \in r} X_{jr}(t). \quad (3)$$

We must describe the queueing dynamics for the queue size process $X(t) = (X_{jr}(t) : j \in \mathcal{J}, r \in \mathcal{R})$. In particular, because we are considering a multiclass queueing network, once the links have been scheduled, we must decide which route will be served. For the BackPressure policy, this would involve comparing the queue size for each route at the queue and at the next downstream node. Shortly, we will explain how this decision is made as part of a Proportional Scheduler. In any case, given a schedule $\sigma(t) \in \mathcal{S}_{Q(t)}$ is chosen at time t , we let $\xi(t) = (\xi_{jr}(t) : j \in \mathcal{J}, r \in \mathcal{R})$ denote the number of route r packets served from queue j under schedule $\sigma(t)$. For $\xi(t)$, the following constraints must hold:

$$\sum_{r:j \in r} \xi_{jr}(t) = \sigma_j(t), \quad \xi_{jr}(t) \leq X_{jr}(t);$$

that is, the number of jobs served from each class is the number scheduled and we cannot schedule more jobs from each class than there are at the queue. Given this, the queueing dynamics of a multihop switch network are as follows:

$$X_{jr}(t+1) = X_{jr}(t) + a_{jr}(t+1) - \xi_{jr}(t+1) + \xi_{j'r}(t+1), \quad (4)$$

where for each $r \in \mathcal{R}$ and each link is $j \in r$ its upstream link is j' .³ Notice that once a packet has been served at its queue it joins the next queue on its route. We call this a *multihop network with fixed routing*.

Given the process $\{X(t)\}_{t=0}^{\infty}$ is a Markov chain, we say the queue size process is *maximum stable* if it is positive recurrent whenever the vector of arrival rates, $(\bar{a}_j : j \in \mathcal{J})$, belongs to the interior of $\langle S \rangle$.

3 Switch policies

We now define the main switching policies considered in this paper. For single-hop networks, we introduce the (α, g) -switch policies and, for multihop networks, we consider the Proportional Scheduler.

3.1 The (α, g) -switch policy

We let α be a positive real number and, for each $j \in \mathcal{J}$, we let $g_j : \mathbb{R}_+ \rightarrow \mathbb{R}$ be a strictly increasing, differentiable, strictly concave function. Given a queue size vector $Q(t) = (Q_j(t) : j \in \mathcal{J})$ at time t , we define $\bar{\sigma}(t+1) = (\bar{\sigma}_j(t+1) : j \in \mathcal{J})$ to be a solution to the optimization

$$\text{maximize} \quad \sum_{j \in \mathcal{J}} g_j(s_j) Q_j(t)^\alpha \quad (5a)$$

$$\text{over} \quad s \in \langle S_{Q(t)} \rangle. \quad (5b)$$

In general, $\bar{\sigma}(t+1)$ need not belong to the set of schedules $S_{Q(t)}$. However, $\bar{\sigma}(t+1)$ is a convex combination of points in $\langle S_{Q(t)} \rangle$. Thus we let $\sigma(t+1)$ be a random variable with support on $S_{Q(t)}$ and mean $\bar{\sigma}(t+1)$. The (α, g) -switch policy scheduling policy is the policy that chooses schedule $\sigma(t)$ at each time $t \in \mathbb{N}$.

To be concrete, we assume the random variables $\sigma(t)$ are, respectively, a function of $Q(t-1)$ and an independent (uniform) random variable. This ensures that the queue size process $\{Q(t)\}_{t=0}^{\infty}$ associated with the (α, g) -switch policy scheduling policy is a discrete-time Markov chain. Further, the constraints (5b) ensure that a schedule never exceeds the queue it serves.

³ If no upstream link exists for j , then we set $\xi_{j'r}(t+1) = 0$.

We use the (slightly non-standard) notation $\mathcal{S}_Q$ because it allows for a convenient representation of the set of randomized schedules which allow for positive queue size, $\langle \mathcal{S}_Q \rangle$. However, when we consider the fluid limit of this system, we will see that this notation is no longer required and we need only consider the sets $\mathcal{S}$ and $\langle \mathcal{S} \rangle$.

3.1.1 MaxWeight, utility optimization, and α -fairness

The motivation for the (α, g) -policies is to generalize MaxWeight and network utility maximizing allocations. The (α, g) -switch policies generalize MaxWeight. For reasons discussed below, the policies do not completely generalize network utility maximizing allocations. However, the policies allow for a general form of utility maximization and an important instance, weight α -fairness, is a special case of our new class.

We now briefly compare the class of policies above with MaxWeight- α and network utility optimizing allocations. MaxWeight- α is the special case of the optimization problem (5) when $g_j(s_j) = s_j$:

$$\text{maximize } \sum_{j \in \mathcal{J}} s_j Q_j(t)^\alpha \quad \text{over } s \in \mathcal{S}. \quad (6)$$

A network utility maximizing allocation in the sense of [19] corresponds to an optimization of the form

$$\text{maximize } \sum_{j \in \mathcal{J}} g_j \left(\frac{s_j}{Q_j(t)} \right) Q_j(t) \quad \text{over } s \in \mathcal{C}, \quad (7)$$

where $\mathcal{C}$ is some convex set. In comparison with the (α, g) -switch policies, notice that, for a network utility maximizing allocation, the terms $s_j/Q_j(t)$ are applied to concave g_j , while the (α, g) -policy just applies s_j . As a model of bandwidth sharing, the $s_j/Q_j(t)$ term is included for Bandwidth Sharing Networks because it applies a processor-sharing system to each transfer. However, for the (α, g) -switched policies, we do not divide by Q_j since a switch network applies its discipline to each component of the switch, which can affect stability analysis.

A further interesting case is the α -fair family of allocation policies, as first introduced by Mo and Walrand [26]. For $\alpha > 0$ and $\alpha \neq 1$, these are defined as follows:

$$\text{maximize } \sum_{j \in \mathcal{J}} \frac{s_j^{1-\alpha}}{1-\alpha} Q_j(t)^\alpha \quad \text{over } s \in \langle \mathcal{S} \rangle.$$

We remark that the (α, g) -switch policy coincides with the weighted α -fair policy when $g_j(s_j) = s_j^{1-\alpha}/(1-\alpha)$. For switch networks, the α -fair policy was first considered by Shah and Wischik [32].

3.1.2 Fluid model

The fluid model associated with the (α, g) -switch policy is a positive, absolutely continuous process $q(t) = (q_j(t) : j \in \mathcal{J}), t \in \mathbb{R}_+$, where, for $j \in \mathcal{J}$ and for almost every⁴ time $t \in \mathbb{R}_+$, if $q_j(t) > 0$ then

$$\frac{dq_j(t)}{dt} = \bar{a}_j - \sigma_j^*(q(t)), \quad (8)$$

where $\sigma^*(q)$ solves the optimization

$$\text{maximize} \quad \sum_{j \in \mathcal{J}} g_j(s_j) q_j^\alpha \quad \text{over} \quad s \in \langle S \rangle. \quad (9)$$

Note that we optimize over $\langle S \rangle$ rather than $\langle S_q \rangle$. Also note that if $q(t) = 0$ then, since $q(t)$ is positive, the derivative must be zero at a point of differentiability.

We say that a fluid model is *stable* if there exists a time $T > 0$ such that, for every $\{q(t)\}_{t \in \mathbb{R}_+}$ satisfying (8) and with $\|q(0)\|_1 = 1$,

$$q_j(t) = 0, \quad j \in \mathcal{J}, \quad (10)$$

for all $t \geq T$. Here and hereafter, $\|\cdot\|_1$ is the L_1 -norm.

3.2 Proportional scheduler

Given a vector of link queue sizes $(Q_j(t) : j \in \mathcal{J})$, the Proportional Scheduler for multihop networks is defined as follows:

PS1. Over set of schedules $\langle S_{Q(t)} \rangle$ solve the optimization

$$\text{maximize} \quad \sum_{j \in \mathcal{J}} Q_j(t) \log(s_j) \quad (11a)$$

$$\text{over} \quad s \in \langle S_{Q(t)} \rangle. \quad (11b)$$

As previously, let $\sigma(t+1)$ be a random variable on $S_{Q(t)}$ whose mean solves this optimization.

PS2. From each queue $j \in \mathcal{J}$, serve $\sigma_j(t+1)$ packets selected uniformly at random from $Q_j(t)$ packets at the queue. These then join their next downstream node as determined by their route class.

Notice we do not need to know the routes of packets in order to define the above policy. We assume that, after service at a queue, packets go to the next hop on their respective routes. The optimization (11) is due to Kelly [19] and is called the proportional fair optimization. We refer to the policy as the *Proportional Scheduler*.⁵

⁴ By *almost every*, we mean on all points except for a set of Lebesgue measure zero.

⁵ Since fairness is not, as such, an objective in our scheduling discipline, we omit the use of the word fair.

3.2.1 The BackPressure policy

We now briefly compare the Proportional Scheduler with the BackPressure policy. For multihop networks, the principle maximal stable policy has been the BackPressure policy. Given the vector of queue sizes $(X_{jr}(t) : j \in \mathcal{J}, r \in \mathcal{R})$, this policy is defined by the following three steps:

BP1. For each link $j = (n, n') \in \mathcal{J}$ directed out of each node $n \in \mathcal{N}$, calculate weights by comparing with upstream queue lengths,

$$w_j(X(t)) = \max_{r: j \in r} \left\{ X_{jr}(t) - X_{j_+^r}^r(t), 0 \right\}, \quad (12)$$

and let $r_j^*(X(t))$ be the solution to this maximization.⁶

BP2. Over set of schedules $\mathcal{S}$, solve the optimization

$$\max_{\sigma \in \mathcal{S}} \sum_{j \in \mathcal{J}} \sigma_j w_j(X(t)), \quad (13)$$

and let $\sigma^*(X(t))$ be the solution to this optimization.

BP3. If $w_j(X(t)) > 0$, at the next time instance schedule $\sigma_j^*(X(t))$ packets from route $r_j^*(X(t))$ from each link $j \in \mathcal{J}$, else, do not schedule any packets on link j .

MaxWeight is the special case of BackPressure when we consider a single-hop network. Notice, in the first step above, information must be exchanged along links to make a queue size comparison (12).

In contrast to BackPressure, the Proportional Scheduler does not maintain a queue for each route. It maintains a queue for each link. Further, the Proportional Scheduler does not compare queue sizes with downstream queues. Thus the policy is, in a certain sense, more decentralized. In the Proportional Scheduler, packets are selected at random from the queue. In most communication systems, packets are served according to a first-in-first-out discipline. We remark that it is possible to prove maximum stability results when the queueing discipline PS2 is replaced by a FIFO queueing discipline. However, the proof is significantly more technical in nature. We refer the reader to the forthcoming article for a proof (Bramson M, D'Auria B, Walton N, Proportional switching in fifo routed networks “in preparation”).

3.2.2 Tree network example

To emphasize the advantages of the queueing structure of our policy, we consider a tree network of degree d and diameter D . Here each non-leaf node of degree d , the diameter of the graph, is of degree D and there is a central node such that each leaf node is a distance $D/2$ from the central node. See Fig. 1 for a tree of degree $d = 3$ and diameter $D = 6$. BackPressure, as described above, requires each node to maintain a

⁶ Here if there is no next link after j on route r then we set $X_{j_+^r}^r = 0$.

Fig. 1 A tree network of depth $d = 3$ and diameter $D = 6$. There are 12 leaf nodes and thus $132 = 12 \cdot 11$ directed routes between leaf nodes

queue for each route passing through it. Thus, one can verify that for this network the central node of a tree must maintain a number of queues given by

$$d(d-1)^{(D-1)}.$$

For this example, all routes with the same destination can be merged once they intersect. For this reason, BackPressure can be implemented slightly more efficiently in each node by storing a queue for each destination; however, even in this case the number of queues required for the central node grows as

$$d(d-1)^{(D/2-1)}.$$

In other words to implement BackPressure, the nodes of this network must maintain a number of queues that grow exponentially with the depth of the network. In general, one expects the number of queues required to implement BackPressure to grow with the network's size.

However, we observe that the Proportional Scheduler maintains a queue for each outgoing link from a node; that is, for each node n the policy maintains a queue for each $j = (n, n') \in \mathcal{J}$. Thus, in our example, a Proportional Scheduler would require d queues at the central node. This is significantly smaller than the number required by BackPressure. In general, for the Proportional Scheduler to be implemented, each node requires a queue for each neighboring node, which is smaller than the network size and is often bounded as we increase the size of a network.

3.2.3 Fluid model

The fluid model associated with the multihop network operating under the Proportional Scheduler is a positive, absolutely continuous process $x(t) = (x_{jr}(t) : j \in \mathcal{J}, r \in \mathcal{R})$

and $q(t) = (q_j(t) : j \in \mathcal{J}), t \in \mathbb{R}_+$, where for almost every⁷ time $t \in \mathbb{R}_+$,

$$\frac{dx_{jr}(t)}{dt} = \frac{x_{j_r^-(t)}}{q_{j_r^-(t)}} \sigma_{j_r^-}^*(q(t)) - \frac{x_{jr}(t)}{q_j(t)} \sigma_j^*(q(t)), \quad (14)$$

where $\sigma^*(q)$ solves the optimization

$$\text{maximize} \quad \sum_{j \in \mathcal{J}} q_j \log s_j \quad \text{over} \quad s \in \langle S \rangle. \quad (15)$$

Above, we apply the convention that

$$\sum_{r: j \in r} x_{jr}(t) = q_j(t). \quad (16)$$

Further, in Eq. (14), if the queue j is the first queue on route r , i.e., $j = j_1^r$, then we apply the convention that⁸

$$\frac{x_{j_r^-(t)}}{q_{j_r^-(t)}} \sigma_{j_r^-}^*(q(t)) = \bar{a}_{jr}.$$

In this way, we account for external arrivals.

4 Main results

The main objectives of this article are to prove maximum stability for the (α, g) -switch policies in single-hop networks and to prove maximum stability for the Proportional Scheduler in multihop networks with fixed routing. These results are stated as follows.

Theorem 1 *For the (α, g) -switch policy applied to a single-hop switched network, if the vector of average arrival rates $\bar{a}$ belongs to the interior of $\langle S \rangle$ then the queue size process $\{Q(t)\}_{t=0}^\infty$ is positive recurrent.*

Theorem 2 *For the Proportional Scheduler applied to a multihop switched network, if the vector of average arrival rates $\bar{a}$ belongs to the interior of $\langle S \rangle$ then the fluid model of this system is stable.*

The first result is proven by appropriately modifying the Lyapunov argument of Bonald and Massoulié [3] for weighted α -fair networks. For the second result, we give two proofs. The first proof shows how the result can be demonstrated within the framework of Massoulié [23]. Essentially, Massoulié's result proves maximum

⁷ By *almost every*, we mean on all points except for a set of Lebesgue measure zero.

⁸ This is similar to the proof of Bramson [4] for Head-of-the-line processor-sharing networks. Bramson introduces an additional queue for each route as a proof device, while here, the introduction of j_0^r is purely for notational and conceptual convenience.

stability of proportional fairness in single-class networks with probabilistic routing (i.e., Jackson type [14]), while we wish to prove maximum stability of proportional fairness for multiclass networks with fixed routing (i.e., Kelly type [2, 20]). We note that, due to the structure of the logarithm, the latter problem can be reduced to the former. The second proof observes that the result can be proven within the framework of Bramson [4]. Bramson shows the stability of fluid models of networks of processor-sharing queues. We show by the differentiability properties of the proportionally fair objective that the fluid system for the proportionally fair system has essentially the same behavior as a network of processor-sharing queues. Both these proofs and a number of consequences and conjectures rely on deeper connections with reversible queueing systems. These are described in the forthcoming paper (Bramson M, D’Auria B, Walton N, Proportional switching in fifo routed networks “in preparation”).

Theorem 1 shows that the class of (α, g) -switch policies has the same stability properties as the MaxWeight policies. This, like MaxWeight, is proven for single-hop networks; however, for multiclass queueing networks, results that hold for single-hop systems do not naturally extend when directly applied multihop systems. See Chapter 3 of Bramson [5] for various counterexamples of this type. This point is particularly poignant for switch networks where there has been a great deal of recent progress on the stability of single-hop switch networks, [17, 29]. Theorem 2 gives the first result where stability of a single-hop switch network implies stability of its multihop counterpart. The result relies on the specific properties of proportional fairness. Nonetheless, the hope is this initial result will lead to general methods for extending stability results from single-hop to multihop networks.

There are a number of immediate advantages of the proportional fair optimization in comparison to BackPressure. We only need to maintain a queue for each outgoing link, while BackPressure requires a vastly more complex data structure. BackPressure must know the route class of each packet within the network to make a routing decision while for our policy, packets can be routed by only knowing each packet’s next hop. Messages do not need to be sent to calculate weights and so scheduling decisions can be completely decomposed between independently functioning components; that is, if $\mathcal{S}$ and $\mathcal{J}$ can be split into independent components $\prod_k \mathcal{S}^{(k)}$ and $\cup_k \mathcal{J}^{(k)}$, then

$$\max_{s \in \mathcal{S}} \sum_{j \in \mathcal{J}} Q_j(t) \log s_j = \sum_k \max_{s \in \mathcal{S}^{(k)}} \sum_{j \in \mathcal{J}^{(k)}} Q_j(t) \log s_j.$$

Here we imagine each $\mathcal{S}^{(k)}$ corresponds to the constraints of an independent network component, for example, an input-queued switch for each k . The above statement shows that, without messaging, we can solve this optimization separately for each component, implement the decision, and still guarantee stability.

Informally speaking, the Proportional Scheduler does not need to know the route structure of the network to make a routing decision, while BackPressure does. For this reason, the Proportional Scheduler scales far better with network size.

5 Proof of Theorem 1

The proof of Theorem 1 follows the fluid limit approach of Dai [8]. We prove that a certain fluid model is satisfied by the limit of a sequence of Markov chains using the (α, g) -switch policy. We prove that the associated fluid model with the (α, g) -policy is stable. We then give a proof of positive recurrence of the original Markov chain.

5.1 Fluid model

In this section, we state fluid model equations and the fluid limit associated with the (α, g) -switch policy. The proof of the fluid limit is given in the appendix.

Let $\{Q^{(c)}\}_{c \in \mathbb{N}}$ be a sequence of versions of our queueing process for the (α, g) -switch policy, where $\|Q(0)\|_1 = c$. We define

$$\bar{Q}^{(c)}(t) = \frac{Q^{(c)}(\lfloor ct \rfloor)}{c}, \quad (17)$$

for $t \in \mathbb{R}_+$. The following result formalizes a fluid model q , (8), as the limit of $\{\bar{Q}^{(c)}\}_{c \in \mathbb{N}}$. In informal terms, it states that the only possible limit of the sequence $\{\bar{Q}^{(c)}\}_{c \in \mathbb{N}}$ as $c \rightarrow \infty$ is a process q satisfying (8).

Proposition 1 (Fluid Limit) *The sequence of stochastic processes $\{\bar{Q}^{(c)}\}_{c \in \mathbb{N}}$ is tight⁹ with respect to the topology of uniform convergence on compact time intervals. Moreover, any weakly convergent subsequence of $\{\bar{Q}^{(c)}\}_{c \in \mathbb{N}}$ converges to a Lipschitz continuous process satisfying fluid equations (8).*

The proof—and indeed statement—of Proposition 1 is somewhat technical. However, the main point is that we can compare the queueing process to a tractable fluid model, q , satisfying equations (8). Proposition 1 is proven in the appendix. We now analyze the stability of the fluid models.

5.2 Fluid stability

We consider a process $\{q(t)\}_{t \in \mathbb{R}_+}$, which satisfies the fluid limit equations (8). In the following theorem, we show that these fluid solutions are stable in the sense that they hit the zero state in finite time. This result will be sufficient to prove positive recurrence of the (α, g) -switch policy.

Proposition 2 (Fluid Stability) *Given $(\bar{a}_j : j \in \mathcal{J})$ belonging to the interior of $\langle S \rangle$, there exists a time $T > 0$ such that, for every fluid model $\{q(t)\}_{t \in \mathbb{R}_+}$ satisfying (8) and with $\|q(0)\|_1 = 1$,*

$$q_j(t) = 0, \quad j \in \mathcal{J}, \quad (18)$$

⁹ Recall that a sequence of random processes $\{\bar{Q}^{(c)}\}_c$ is tight if every subsequence of $\{\bar{Q}^{(c)}\}_c$ has a weakly convergent subsequence.

for all $t \geq T$.

The main idea is to consider the gradient of the tangent line of the (α, g) -switch policy objective between two points: the arrival rate and the optimal solution, see (20) below. Integrating this obtains a Lyapunov function (22). This idea is used by Bonald and Massoulié [3] in their analysis of weighted α -fair networks. For switch networks, Proposition 2 follows analogously.

Proof We define $G_q(s)$ to be the objective of the (α, g) -switch policy optimization,

$$G_q(s) = \sum_{j \in \mathcal{J}} g_j(s_j) q_j^\alpha. \quad (19)$$

Recall that in our fluid equations

$$\sigma^*(q(t)) \in \operatorname{argmax}_{s \in \langle S \rangle} G_{q(t)}(s).$$

Any vector ρ belonging to the interior of $\langle S \rangle$ is not optimal. Thus $G_{q(t)}(\rho) < G_{q(t)}(\sigma^*(q(t)))$. As $G_{q(t)}(\cdot)$ is strictly concave, $G_{q(t)}(\cdot)$ must be increasing along the line connecting ρ to $\sigma^*(q(t))$. In other words, for all ρ in the interior of $\langle S \rangle$ and for $q(t) \neq 0$,

$$\left(\sigma^*(q(t)) - \rho \right) \cdot \nabla G_{q(t)}(\rho) > 0. \quad (20)$$

Here $\nabla G_q(\rho) = (g'_j(\rho_j) q_j(t)^\alpha : j \in \mathcal{J})$. Since $\bar{a}$ belongs to the interior of $\langle S \rangle$, there exists $\epsilon > 0$ such that $(1 + \epsilon)\bar{a} \in \langle S \rangle$. We define $\rho = (1 + \epsilon)\bar{a}$. In this case, we can re-express the inequality (20) as follows:

$$\sum_{j \in \mathcal{J}} \left(\bar{a}_j - \sigma_j^*(q(t)) \right) g'_j(\rho_j) q_j(t)^\alpha \leq -\epsilon \sum_{j \in \mathcal{J}} g'_j(\rho_j) q_j(t)^\alpha. \quad (21)$$

As $\frac{dq_j}{dt} = \bar{a}_j - \sigma_j^*(q(t))$, we define the Lyapunov function

$$L(q) = \sum_{j \in \mathcal{J}} g'_j(\rho_j) \frac{q_j^{1+\alpha}}{1+\alpha}, \quad (22)$$

$q \in \mathbb{R}_+^{\mathcal{J}}$. The function $L(q)$ is positive and $L(q) = 0$ iff $q_j = 0$ for all $j \in \mathcal{J}$. We now observe

$$\frac{dL(q(t))}{dt} = \sum_{j \in \mathcal{J}} \left(\bar{a}_j - \sigma_j^*(q(t)) \right) g'_j(\rho_j) q_j(t)^\alpha \leq -\epsilon \sum_{j \in \mathcal{J}} g'_j(\rho_j) q_j(t)^\alpha. \quad (23)$$

The equality holds by the chain rule and the inequality holds by (21).

We define two norms on $\mathbb{R}_+^{\mathcal{J}}$:

$$\|q\|_{1+\alpha} = (L(q))^{\frac{1}{1+\alpha}}, \quad \|q\|_{\alpha} = \left(\sum_{j \in \mathcal{J}} g'_j(\rho_j) q_j(t)^{\alpha} \right)^{\frac{1}{\alpha}}. \quad (24)$$

By the Lipschitz equivalence of norms, there is a constant $\gamma > 0$ such that

$$\gamma \|q\|_{1+\alpha} \leq \|q\|_{\alpha},$$

for all $q \in \mathbb{R}_+^{\mathcal{J}}$.¹⁰ Applying this observation to the inequality (23), we see that

$$\frac{dL(q(t))}{dt} \leq -\epsilon \gamma^{\alpha} L(q(t))^{\frac{\alpha}{1+\alpha}}. \quad (25)$$

Observe, by the above inequality, if $L(q(T)) = 0$ for any differentiable point T then $L(q(t)) = 0$ for all $t \geq T$. Now, while $L(q(t)) > 0$, we have from (25) that

$$\begin{aligned} L(q(t))^{\frac{1}{1+\alpha}} - L(q(0))^{\frac{1}{1+\alpha}} &= \int_0^t (1+\alpha)^{-1} L(q(t))^{\frac{-\alpha}{1+\alpha}} \frac{dL(q(t))}{dt} dt \\ &\leq -\epsilon (1+\alpha)^{-1} \gamma^{\alpha} t. \end{aligned}$$

Rearranging this expression, we see that for all times t

$$L(q(t)) \leq \left(L(q(0))^{\frac{1}{1+\alpha}} - \epsilon (1+\alpha)^{-1} \gamma^{\alpha} t \right)_+^{1+\alpha}. \quad (26)$$

The function $L(q)$ is continuous and therefore bounded above by a constant, K , for all values of q with $\|q\|_1 = 1$. Hence, if $\|q(0)\|_1 = 1$, $L(q(t)) = 0$ for all $t \geq T$ where

$$T = \frac{(1+\alpha)K^{\frac{1}{1+\alpha}}}{\epsilon \gamma^{\alpha}},$$

and thus, as required, $q_j(t) = 0$, $j \in \mathcal{J}$, for all $t \geq T$. $\square$

5.3 Positive recurrence

We are now in a position to combine Propositions 1 and 2 to prove Theorem 1. We can at this point apply the general stability results of Dai [8,9] and Bramson [5]. Nonetheless, we can also provide a self-contained proof with Proposition 3 proven in the Appendix.

¹⁰ Note, $\|q\|_{1+\alpha} \leq (1+\alpha)^{-\frac{1}{1+\alpha}} |\mathcal{J}| \max_j g'_j(\rho_j) q_j$ and also note that $\max_j g'_j(\rho_j) q_j \leq \|q\|_{\alpha}$. So, for instance, we can take $\gamma = (1+\alpha)^{\frac{1}{1+\alpha}} |\mathcal{J}|^{-1}$.

Proof For every $t \geq 0$, the sequence of queue sizes $\{\bar{Q}^{(c)}(t)\}_{c \in \mathbb{N}}$ is uniformly integrable. This is proven in Lemma 8 in the appendix. By Proposition 1, for any unbounded sequence in $\mathbb{N}$, there is a subsequence $\{c_k\}_{k \in \mathbb{N}}$ for which $\bar{Q}^{(c_k)}$ converges in distribution to the fluid solution q . Let T be the time given in Proposition 2, where $q_j(T) = 0$ for $j \in \mathcal{J}$. Since $\{|\bar{Q}^{(c_k)}(T)|\}_{c_k}$ is uniformly integrable and converges in distribution to $q(T)$, we also have L_1 convergence:

$$\lim_{c_k \rightarrow \infty} \mathbb{E} \|\bar{Q}^{(c_k)}(T)\|_1 = \mathbb{E} \|q(T)\|_1 = 0. \quad (27)$$

This implies that there exists a κ such that for all $c > \kappa$

$$\mathbb{E} \|\bar{Q}^{(c)}(T)\|_1 < (1 - \epsilon). \quad (28)$$

Note that if (28) did not hold, then we could find a subsequence for which (27) did not hold; thus, we would have a contradiction. Expanding this inequality (28), we have the following *multiplicative Foster's condition*: for $\|Q(0)\|_1 > \kappa$

$$\mathbb{E} \left[\|Q(T\|Q(0)\|_1)\|_1 - \|Q(0)\|_1 \mid Q(0) \right] < -\epsilon \|Q(0)\|_1. \quad (29)$$

Foster's criterion in the form (69) was introduced, and first used to prove network stability, in [27], and is further elaborated in Bramson [5]. This then implies that our process is positive recurrent; see [5, Proposition 4.6] or, alternatively, a short proof of this can be found in Proposition 3 of the appendix. $\square$

6 Proof of Theorem 2

We give two proofs of Theorem 2. The first proof reduces our multiclass fixed routing structure (Kelly type) to a single-class routing structure (Jackson type). From here we can apply the fluid stability proof of Massoulié [23]. The second proof observes that the Lyapunov function arguments of Bramson [4], which apply for head-of-the-line processor-sharing queueing networks, can be extended when services rates are proportionally fair instead of fixed.

6.1 First proof

We state the following result on the fluid model for a single-class Proportional Scheduler with probabilistic routing. Massoulié [23] considered a process $(q_j(t) : j \in \mathcal{J})$ satisfying, for $q_j(t) > 0$, the following system of fluid equations:

$$\frac{dq_j(t)}{dt} = a_j + \sum_{l \in \mathcal{J}} \sigma_l^*(q) p_{lj} - \sigma_j^*(q), \quad (30)$$

where $\sigma_j^*(q)$ gives the solution to the proportionally fair optimization (15). The matrix $P = (p_{lj})_{l,j}$ gives the probability that a link l job next joins link j . So all packets depart

this system, it is assumed that this matrix is sub-stochastic. From this one can calculate from $a = (a_j : j \in \mathcal{J})$ the load induced on each link as

$$\bar{a} = (1 - P^\top)^{-1}a. \quad (31)$$

Massoulié proves the following fluid stability result.

Theorem 3 ([23]) *If $\bar{a}$ belongs to the interior of the set $\langle \mathcal{S} \rangle$ then the fluid equations (30) are stable.*

Here stable is meant in the same sense as in Proposition 2.

Since in Massoulié's model each packet leaving a link must have the same routing behavior, it is not immediately clear that the result extends to multiclass queueing networks where routing depends on a packet's route class. However, the following lemma makes this possible.

Lemma 1 *For $j \in \mathcal{J}$, take a positive vector $x^j = (x_{jr} : r \in \mathcal{R}, j \in r)$ and positive constant σ_j . Let q_j be the sum of the components of x^j . The following holds:*

$$q_j \log \sigma_j + \sum_{r:j \in r} x_{jr} \log \frac{x_{jr}}{q_j} \quad (32)$$

$$= \max_{\gamma^j \geq 0} \sum_{r:j \in r} x_{jr} \log \gamma_{jr} \text{ s.t. } \sum_{j \in r} \gamma_{jr} = \sigma_j \quad (33)$$

and the above optimization is solved by $\gamma_{jr} = \frac{x_{jr}\sigma_j}{q_j}$.

The lemma is left as an exercise. What the lemma says is that if packets $x^j = (x_{jr} : r \in \mathcal{R}, j \in r)$ are sharing the resource j then under a proportional fair optimization they share it equally and, perhaps more importantly, the equality (32) shows that the form of the logarithm in the optimization remains unchanged after this optimization step. Thus, to solve a proportionally fair optimization given x_{jr} , we can solve the proportionally fair optimization given q_j , i.e., optimizing to $q_j \log \sigma_j$ terms in (32), and then assign rates according to a processor-sharing discipline: $\gamma_{jr} = \frac{x_{jr}\sigma_j}{q_j}$. This describes the per-link queueing mechanism described in Sect. 3.2, which possesses significant structural simplification when compared to other policies known to be maximum stable for multihop-routed switch networks.

Proof As a consequence of Lemma 1, we have that the optimization

$$\text{maximize } \sum_{j \in \mathcal{J}} \sum_{r \in \mathcal{R}} x_{jr} \log \gamma_{jr} \quad (34a)$$

$$\text{subject to } \sum_{r:j \in r} \gamma_{jr} = \sigma_j, \quad j \in \mathcal{J}, \quad (34b)$$

$$\text{over } \sigma \in \langle \mathcal{S} \rangle, \quad (34c)$$

is equivalent to solving the optimization

$$\text{maximize} \quad \sum_{j \in \mathcal{J}} q_j \log \gamma_{jr} \quad (35a)$$

$$\text{over} \quad \sigma \in \langle \mathcal{S} \rangle, \quad (35b)$$

and setting $\gamma_{jr} = x_{jr} \sigma_j(q) / q_j$.

Consider the process $x(t) = (x_{jr}(t) : r \in \mathcal{R}, j \in r)$, a bandwidth sharing network of the type (30) where there is a class is now indexed by (j, r) ; that is, the stations of the network are no longer indexed by j , but by pairs (j, r) . After completing service at (j, r) , packets next join (j_r^+, r) with fixed probability. This system obeys the fluid equations¹¹

$$\frac{dx_{jr}(t)}{dt} = \gamma_{j_r^-}^*(x(t)) - \gamma_{jr}^*(x(t)). \quad (36)$$

These equations are exactly of the form (30). Thus, by Massoulié [23], this system is fluid stable. However, the equivalence of the optimizations (34) and (35) shows that this fluid system is the same as our Proportional Scheduler's fluid model:

$$\frac{dx_{jr}(t)}{dt} = \frac{x_{j_r^-}(t)}{q_{j_r^-}(t)} \sigma_{j_r^-}^*(q(t)) - \frac{x_{jr}(t)}{q_j(t)} \sigma_j^*(q(t)).$$

Thus the stability of our fluid model now follows from Massoulié [23].

The observation of Lemma 1 and its implications for multihop switch systems are critical. After observing Lemma 1, the reduction to Massoulié's proof is straightforward. However, to this author at least, it was surprising that the multiclass system can be reduced to a single-class system in this way. For instance, it is not immediately clear that one can reduce results on the stability of networks with BCMP-Kelly (multiclass) routing as a direct reduction from stability results networks with Jackson (single-class) routing. Nonetheless, the fluid stability proof of Massoulié is not so straightforward. A second, direct proof can argued from Bramson [4]. As a brief aside, with Lemma 1 and Lemma 2 (below), the fluid stability result of Bramson for processor-sharing networks [4] and the fluid result of Massoulié for proportional fair networks [23] have some equivalence. However, such equivalences are, perhaps, less striking than their resulting implications for switch systems. We now see that, for the proportionally fair case, stability of a single-hop queueing network also implies stability of its multihop-routed counterpart. Moreover, the queueing structure of this system is far simpler than BackPressure, the canonical class of multihop throughput optimal switch policies.

¹¹ Here $\gamma_{j_r^-}^*(x(t)) = \bar{a}_r$ when j is the first link on route r .

6.2 Second proof

Bramson [4] considers the fluid model of a multiclass processor-sharing queueing network. Here the service rate at queues $\sigma \in \mathbb{R}_+^J$ is not a function of queue sizes. So, for $x_{jr}(t) > 0$, the fluid model of Bramson is as follows:

$$\frac{dx_{jr}(t)}{dt} = \frac{x_{j'-r}(t)}{q_{j'-}(t)} \sigma_{j'-} - \frac{x_{jr}(t)}{q_j(t)} \sigma_j.$$

Bramson considers a Lyapunov function

$$\hat{H}(x) = \sum_{r \in \mathcal{R}} \sum_{j \in r} x_{jr} \log \left(\frac{x_{jr} \sigma_j}{q_j \bar{a}_r} \right). \quad (37)$$

This is the large deviations rate function associated with a network of processor-sharing queues. Informally speaking, this proof relies on the fact that the partial derivatives of this function are, for $x_{jr} > 0$,

$$\frac{\partial \hat{H}(x)}{\partial x_{jr}} = \log \left(\frac{x_{jr} \sigma_j}{q_j \bar{a}_r} \right). \quad (38)$$

If we now replace σ with the proportionally fair solution $\sigma^*(q)$ and consider the following Lyapunov function

$$H(x) = \sum_{r \in \mathcal{R}} \sum_{j \in r} x_{jr} \log \left(\frac{x_{jr} \sigma_j^*(q)}{q_j \bar{a}_r} \right), \quad (39)$$

due to the properties of the proportionally fair optimization, the partial derivatives of this optimization remain the same. This is demonstrated in the following lemma.

Lemma 2 For $x_{jr} > 0$

$$\frac{\partial H(x)}{\partial x_{jr}} = \log \left(\frac{x_{jr} \sigma_j^*(q)}{q_j \bar{a}_r} \right). \quad (40)$$

This lemma is proven in the Appendix. With the addition of a few technical lemmas, we can now give a direct proof of Theorem 2. These technical lemmas are stated and proven after the following proof.

Proof It is proven in Lemma 4 that $H(x)$ is positive and is minimized when $x = 0$. Applying Lemma 2, we can see that the following sequence of equalities hold:

$$\frac{dH(x(t))}{dt} = \sum_{r,j \in \mathcal{R}} \left(\frac{x_{jr}^r}{q_{jr}^r} \sigma_{jr}^*(q) - \frac{x_{jr}}{q_j} \sigma_j^*(q) \right) \log \left(\frac{x_{jr} \sigma_j^*(q)}{q_j \bar{a}_r} \right) \quad (41)$$

$$= - \sum_{r \in \mathcal{R}} \bar{a}_r \sum_{j \in r} \frac{x_{jr} \sigma_j^*(q)}{q_j \bar{a}_r} \log \left(\frac{\left[\frac{x_{jr} \sigma_j^*(q)}{q_j \bar{a}_r} \right]}{\left[\frac{x_{j_+^r}^r \sigma_{j_+^r}^*(q)}{q_{j_+^r}^r \bar{a}_r} \right]} \right). \quad (42)$$

To derive (42), we increment the j -index terms of summation (41) so that the coefficients of the logarithm are equal. After multiplying and dividing by $\bar{a}_r$ we gain expression (42).

Notice, in the first equality, we include the term $j = j_0^r$ since in this case $\log((x_{jr} \sigma_j^*(q_j))/(q_j \bar{a}_r)) = \log 1 = 0$. The following sequence of inequalities hold, for $q \neq 0$:

$$\begin{aligned} \frac{dH(x(t))}{dt} &\leq - \sum_{r \in \mathcal{R}} \left[\bar{a}_r \left(\sum_{j \in r} \frac{x_{jr} \sigma_j^*(q)}{q_j \bar{a}_r} \right)^{-1} \right. \\ &\quad \times \sum_{j \in r} \left(\frac{x_{jr} \sigma_j^*(q)}{q_{jr} \bar{a}_r} - \frac{x_{j_+^r}^r \sigma_{j_+^r}^*(q)}{q_{j_+^r}^r \bar{a}_r} \right)^2 \Big] \\ &\leq - \sum_{r \in \mathcal{R}} \left[\bar{a}_r \left(\frac{\sigma_{\max}}{\bar{a}_r} \right)^{-1} \times \sum_{j \in r} \left(\frac{x_{jr} \sigma_j^*(q)}{q_j \bar{a}_r} - \frac{x_{j_+^r}^r \sigma_{j_+^r}^*(q)}{q_{j_+^r}^r \bar{a}_r} \right)^2 \right] \\ &\leq -\epsilon. \end{aligned}$$

where ϵ is defined in Lemma 6. In the first inequality, we apply Lemma 5 to (42). In the second inequality, we apply the bound

$$\sum_{j \in r} \frac{x_{jr} \sigma_j^*(q)}{q_j \bar{a}_r} \leq \sum_{j \in r} \frac{x_{jr} \sigma_{\max}}{q_j \bar{a}_r} = \frac{\sigma_{\max}}{\bar{a}_r}, \quad (43)$$

where ϵ is defined in Lemma 6. In the third inequality, we apply Lemma 6. This then shows that our Lyapunov function has strictly negative drift.

By Lemma 4, $H(x) = 0$ only if $x = 0$. Thus, for all $x(0)$ such that $H(x(0)) \leq h$ for some constant $h > 0$, we have for all $t \geq h\epsilon^{-1}$

$$q(t) = 0. \quad (44)$$

Thus, fluid stability holds.

We now prove the lemmas required in the above proof. For two probability distributions p and q defined on the same finite set $\mathcal{X}$, the relative entropy between p and

q is defined to be

$$D(p||q) = \sum_{x \in \mathcal{X}} p_x \log \frac{p_x}{q_x}.$$

The following bound on $D(p||q)$ holds.

Lemma 3 (Pinsker's Inequality)

$$\sqrt{D(p||q)} \geq \sum_{x \in \mathcal{X}} |p_x - q_x|.$$

A proof of this bound is found in Cover and Thomas [7]. Relative entropy is positive, continuous in p and minimized when $p = q$. This gives the following.

Lemma 4 *The function $H(x)$ is positive, is bounded when $\|x\|_1 = 1$, and is minimized when $x = 0$.*

Proof of Theorem 2 Observe that $H(x)$ can be expressed as linear combination of relative entropy terms as follows:

$$\begin{aligned} H(x) &= \sum_{j \in \mathcal{J}} q_j D\left(\left(\frac{x_{jr}}{q_j}\right)_{r \ni j} \left\| \left(\frac{\bar{a}_r}{a_j}\right)_{r \ni j}\right.\right) + \sum_{j \in \mathcal{J}} q_j \log \frac{\sigma_j^*(q)}{a_j} \\ &\geq \sum_{j \in \mathcal{J}} q_j \log \frac{\sigma_j^*(q)}{a_j} \geq 0. \end{aligned} \quad (45)$$

The first inequality follows by the positivity of the relative entropy. The second inequality follows by the optimality of $\sigma_j^*(q)$ and the fact the vector $(a_j : j \in \mathcal{J})$ belongs to the interior of $\langle S \rangle$. From the form of the entropy equality above, we see that $H(x)$ is continuous for $\|x\|_1 = 1$ and so is bounded. Further, we note that the inequalities above hold with equality iff $x = 0$. $\square$

If we now do not assume that p and q are probability distributions but instead we assume that they are positive and sum to the same constant, then $D(p||q)$ is still well-defined. The following is a consequence of the above result.

Lemma 5 *If p and q are two positive vectors components indexed by $\mathcal{X}$ and with*

$$\sum_{x \in \mathcal{X}} p_x = \sum_{x \in \mathcal{X}} q_x \quad (46)$$

then

$$\sum_{x \in \mathcal{X}} p_x \log \frac{p_x}{q_x} \geq \frac{1}{\sum_{x \in \mathcal{X}} p_x} \cdot \sum_{x \in \mathcal{X}} (p_x - q_x)^2. \quad (47)$$

Proof Define

$$\tilde{p}_x = \frac{p_x}{\sum_{y \in \mathcal{X}} p_y} \quad \text{and} \quad \tilde{q}_x = \frac{q_x}{\sum_{y \in \mathcal{X}} q_y}. \quad (48)$$

Then

$$\begin{aligned}
 \sum_{x \in \mathcal{X}} p_x \log \frac{p_x}{q_x} &= D(\tilde{p} \parallel \tilde{q}) \sum_{x \in \mathcal{X}} p_x \\
 &\geq \left(\sum_{x \in \mathcal{X}} |\tilde{p}_x - \tilde{q}_x| \right)^2 \sum_{x \in \mathcal{X}} p_x \\
 &= \frac{1}{\sum_{x \in \mathcal{X}} p_x} \left(\sum_{x \in \mathcal{X}} |p_x - q_x| \right)^2 \\
 &\geq \frac{1}{\sum_{x \in \mathcal{X}} p_x} \sum_{x \in \mathcal{X}} (p_x - q_x)^2.
 \end{aligned}$$

In the first equality above, we apply the definition of $D(\tilde{p} \parallel \tilde{q})$; in the second, we apply Pinsker's Inequality; we then rearrange this term and bound to get the result. $\square$

Lemma 6 *There exists $\epsilon > 0$ such that for $x \in \mathbb{R}_+^{\mathcal{R} \times \mathcal{J}}$*

$$\sum_{r \in \mathcal{R}} \bar{a}_r \left(\frac{\sigma_{\max}}{\bar{a}_r} \right)^{-1} \sum_{j \in r} \left(\frac{x_{jr} \sigma_j^*(q)}{q_j \bar{a}_r} - \frac{x_{j_+^r} \sigma_{j_+^r}^*(q)}{q_{j_+^r} \bar{a}_r} \right)^2 > \epsilon. \quad (49)$$

Proof First let us look what is required for the term on the left-hand side of (49) to equal zero—which, by the statement of the lemma, should give a contradiction. So, we would require that for each route $r \in \mathcal{R}$, for all $j \in r$

$$\frac{x_{jr} \sigma_j^*(q)}{q_j \bar{a}_r} - \frac{x_{j_+^r} \sigma_{j_+^r}^*(q)}{q_{j_+^r} \bar{a}_r} = 0. \quad (50)$$

So the terms $\frac{x_{jr} \sigma_j^*(q)}{q_j \bar{a}_r}$ are constant over $j \in r$, including j_0 . But $\frac{x_{jr} \sigma_j^*(q)}{q_j \bar{a}_r} = 1$. So we would require that, for all r and $j \in r$

$$\bar{a}_r = \frac{x_{jr} \sigma_j^*(q)}{q_j}.$$

Now summing over $r \in j$, we have that

$$\bar{a}_j = \sum_{r: j \in r} \bar{a}_r = \sigma_j^*(q). \quad (51)$$

However, we now have a contradiction: by assumption, $\bar{a}_j$ does not belong to the boundary of the capacity set $\langle S \rangle$, while $\sigma_j^*(q)$ does. So the above equality cannot hold. So the left-hand side of (49) cannot equal zero.

Now let us use this last contradiction to extrapolate back and prove the lemma. Since $(\bar{a}_j : j \in \mathcal{J})$ belongs to the interior of $\langle \mathcal{S} \rangle$ and $\sigma_j^*(q) \in \mathcal{A}(\mathcal{S})$, there exists a $\delta > 0$ such that for each vector $q \neq 0$ there is some $j \in \mathcal{J}$ such that

$$\rho_j + \delta < \sigma_j^*(q). \quad (52)$$

Thus, for the above inequality to hold, there exists some r with $j \in r$ such that

$$\bar{a}_r + \frac{\delta}{|r|} < \frac{x_{jr}}{q_j} \sigma_j^*(q) \quad (53)$$

or expressed differently,

$$\frac{x_{j_0^r} \sigma_{j_0^r}^*(q)}{q_{j_0^r} \bar{a}_r} + \frac{\delta}{|r| \bar{a}_r} < \frac{x_{jr} \sigma_j^*(q)}{q_j \bar{a}_r}. \quad (54)$$

Thus, for this to hold from j_0^r to j on route r , there must be two links in the sequence between j_0^r and j such that

$$\frac{\delta}{|r|^2 \bar{a}_r} < \frac{x_{j_+^r} \sigma_{j_+^r}^*(q)}{q_{j_+^r} \bar{a}_r} - \frac{x_{jr} \sigma_j^*(q)}{q_j \bar{a}_r}. \quad (55)$$

Setting $\epsilon = (\frac{\delta}{|r|^2})^2 \sigma_{\max}$, we see that (49) must hold. $\square$

7 Conclusions and future work

In this paper, we have considered a generalized class of switch policies. We have shown that these share the same maximum stability properties as the MaxWeight class of policies. For the Proportional Scheduler, we have extended these stability results to multihop networks with fixed routing. This is the first general proof where a single-hop switch system is proven to be stable when applied to a multihop network. This is also significant given the greatly simplified queueing structure required to define the policy.

Future work might consider the implementation of our policies, which require the solution of a concave optimization problem. Thus, using, for instance, an interior point method, one can approximate a solution to this optimization in computational time that is polynomial in the constraints of the set of feasible schedules $\langle \mathcal{S} \rangle$. After this one can decompose the mean vector onto the set of schedules $\mathcal{S}$. For instance, in the case of an Input-Queued switch, one could use a Birkhoff von Neumann decomposition. Approaches of this type have been considered previously [30]. Compared to MaxWeight, an advantage of the strict convexity of our optimization is that the solutions of our optimization move continuously with continuous changes in queue size. Thus one could consider an alternative, online convex optimization schema where one continuously adapts the current schedule in order to track the optimum (α, g) -policy.

As queue sizes increase, the solution of our optimization will change more slowly and thus a tracking policy would expect to converge to the correct scheduling decision [13, 39].

We have emphasized the structural benefits of implementing a Proportional Scheduler but have not thus far discussed the statistical benefits of the policy. As we mentioned in the introduction, there is a close relationship between proportional fairness and the maximum stability of reversible systems. With this, one can begin to show that the randomized queueing discipline considered in the Proportional Scheduler can be replaced by a FIFO service discipline and still maintain its stability properties. This is important from a practical perspective as most communications systems implement FIFO queueing. From a theoretical perspective, the result is far more involved and technical than the relatively straightforward proofs for the randomized service which fit within this paper. However, with this analysis, further statistical benefits can be observed. For instance, delay is known to increase quadratically with route length for the BackPressure policy; see [6, 35]. However, one can argue that for the Proportional Scheduler, delay will grow linearly with route length.

As mentioned above, we have given the first example where a single-hop maximum stability policy implies stability of its routed multihop counterpart. It is reasonable to conjecture that this holds more generally than for just the Proportional Scheduler. The hope is that the arguments started here could give a general proof technique. One could then hope to extend recent, notable developments on single-hop networks [17, 25, 29, 34] to multihop networks. Thus we begin to provide performance analysis of switched communication networks where packets are communicated.

Acknowledgments The work of Neil Walton is funded by the VENI research programme, which is financed by the Netherlands Organisation for Scientific Research (NWO).

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

Appendix: Proof of fluid limit

For $q \in \mathbb{R}_+^{\mathcal{J}}$, we let $\hat{\sigma}(q)$ be the solution to the optimization problem

$$\text{maximize} \quad \sum_{j \in \mathcal{J}} g_j(s_j) q_j^\alpha \quad \text{over } s \in \langle \mathcal{S}_q \rangle, \quad (56)$$

and we recall the following proof.

Proof To prove the tightness of a sequence of process $\{X^{(c)}\}_c$, from Robert [27, Theorem C.9], we see that we must prove, for each $\eta > 0$ and $t > 0$, that

$$\lim_{\delta \rightarrow 0} \mathbb{P} \left(\sup_{\substack{u, v: u, v < t \\ |v-u| < \delta}} \|X^{(c)}(v) - X^{(c)}(u)\|_1 \geq \eta \right) = 0. \quad (57)$$

We wish to demonstrate this for $\{\bar{Q}^{(c)}\}_c$. For $j \in \mathcal{J}$, we can express $\bar{Q}_j^{(c)}(t)$ as follows:

$$\bar{Q}_j^{(c)}(t) = \bar{Q}_j^{(c)}(0) + \bar{a}_j \frac{\lfloor ct \rfloor}{c} - \bar{S}_j^{(c)}(t) + \bar{M}_j^{(c)}(t) + \bar{N}_j^{(c)}(t), \quad (58)$$

where, from $a(t)$, $\bar{a}$, $\sigma(t)$, $\bar{\sigma}(t)$ and for each $j \in \mathcal{J}$, we define

$$\begin{aligned} \bar{S}_j^{(c)}(t) &= \frac{1}{c} \sum_{s=1}^{\lfloor ct \rfloor} \bar{\sigma}_j(s) = \int_0^{\lfloor ct \rfloor / c} \hat{\sigma}_j(Q_j^{(c)}(cs)) ds, \\ \bar{M}_j^{(c)}(t) &= \frac{1}{c} \sum_{s=1}^{\lfloor ct \rfloor} (a_j(s) - \bar{a}_j), \\ \bar{N}_j^{(c)}(t) &= \frac{1}{c} \sum_{s=1}^{\lfloor ct \rfloor} (\bar{\sigma}_j(s) - \sigma_j(s)). \end{aligned}$$

It is clear that $Q_j(0) + \bar{a}_j \lfloor ct \rfloor / c$ satisfies (57). By the triangle inequality, it is sufficient to show (57) holds for each of the above terms.

The term $\bar{\sigma}(s)$ is bounded by some constant K_1 for all s . So, we have the Lipschitz condition

$$|\bar{S}_j(t) - \bar{S}_j(s)| < K_1 |t - s| + 2K_1 c^{-1}$$

which satisfies (57) with $\delta < \eta K_1^{-1}$.

The process $\bar{M}_j(t)$ is a martingale. Thus, by Doob's L_2 inequality

$$\mathbb{P} \left(\sup_{\substack{u, v: u, v < t \\ |v - u| < \delta}} |\bar{M}_j^{(c)}(v) - \bar{M}_j^{(c)}(u)| \geq \eta \right) \quad (59)$$

$$\leq \mathbb{P} \left(\sup_{u: u < t} |\bar{M}_j^{(c)}(u)| \geq \frac{\eta}{2} \right) \quad (60)$$

$$\leq \frac{4}{\eta^2} \mathbb{E} [\bar{M}_j^{(c)}(t)^2] = \frac{4t}{\eta^2} \left[\frac{1}{c} \mathbb{E}(a_j(1) - \bar{a}_j)^2 \right] \xrightarrow{c \rightarrow \infty} 0. \quad (61)$$

The process $\bar{N}_j^{(c)}(t)$ is a martingale with summands bounded by some constant K_2/c , so by the Azuma–Hoeffding Inequality

$$\mathbb{P}(|\bar{N}_j^{(c)}(t)| \geq \eta) \leq 2e^{-\frac{c\eta^2}{2tK_2^2}}. \quad (62)$$

Again applying Doob's L_2 inequality,

$$\begin{aligned} \mathbb{P} \left(\sup_{\substack{u,v: u,v < t \\ |v-u| < \delta}} |\bar{N}_j^{(c)}(v) - \bar{N}_j^{(c)}(u)| \geq \eta \right) &\leq \mathbb{P} \left(\sup_{u: u < t} |\bar{N}_j^{(c)}(u)| \geq \frac{\eta}{2} \right) \\ &\leq \frac{4}{\eta^2} \mathbb{E} [\bar{N}_j^{(c)}(t)^2] \leq \left(\frac{4K_2 t}{\eta c} \right)^2 \xrightarrow{c \rightarrow \infty} 0. \end{aligned} \quad (63)$$

$$(64)$$

Thus our sequence of processes $\{\bar{Q}^{(c)}\}_c$ is tight.

Note the inequalities above, (59) and (63), also prove that the sequences $\bar{N}^{(c)}$ and $\bar{M}^{(c)}$ converge in distribution to zero.

It remains to show that any fluid limit process satisfies the fluid equations (8). By the Skorohod Representation Theorem, see [27, Theorem C.8], and since $\{\bar{Q}^{(c)}\}_c$ is tight, for any convergent subsequence, we can choose a subsequence where convergence occurs almost surely, along an appropriately chosen probability space. Thus, using this and the definition of $\bar{S}_j$

$$\begin{aligned} \lim_{c \rightarrow \infty} \bar{Q}_j^{(c)}(t) &= \lim_{c \rightarrow \infty} \left\{ \bar{Q}_j^{(c)}(0) + \bar{a}_j \frac{\lfloor ct \rfloor}{c} - \bar{S}_j^{(c)}(t) + \bar{M}_j^{(c)}(t) + \bar{N}_j^{(c)}(t) \right\} \\ &= q_j(0) + \bar{a}_j t - \lim_{c \rightarrow \infty} \bar{S}_j^{(c)}(t) \\ &= q_j(0) + \bar{a}_j t - \lim_{c \rightarrow \infty} \int_0^{\lfloor ct \rfloor / c} \hat{\sigma}_j(Q_j^{(c)}(cs)) ds = q_j(t). \end{aligned}$$

As σ^* is bounded, it is clear that $q_j(t)$ is Lipschitz continuous. If $q_j(t) > 0$ then, by continuity, $q_j(s) > 0$ for an open region $s \in (t - \delta, t + \delta)$. We have for this choice of t and s

$$\begin{aligned} q_j(t) - q_j(s) &= \bar{a}_j(t - s) - \lim_{c \rightarrow \infty} \int_{\lfloor cs \rfloor / c}^{\lfloor ct \rfloor / c} \hat{\sigma}_j(Q_j^{(c)}(cu)) du \\ &= \bar{a}_j(t - s) - \int_{\lfloor cs \rfloor / c}^{\lfloor ct \rfloor / c} \lim_{c \rightarrow \infty} \hat{\sigma}_j(Q_j^{(c)}(cu)) du \\ &= \bar{a}_j(t - s) - \int_{\lfloor cs \rfloor / c}^{\lfloor ct \rfloor / c} \sigma_j^*(q_j(u)) du. \end{aligned}$$

As Lemma 7 states $\hat{\sigma}_j(q_j)$ is continuous for $q_j > 0$, we apply the bounded convergence theorem in the second equality above. Now, in addition to being tight, we have proven that the fluid equations, (8), hold. $\square$

Other Lemmas

The following lemma assists our fluid limit proof, Proposition 1. Recall our definitions of $\hat{\sigma}(q)$, (56), and $\sigma^*(q)$, (15). We let $\{q^{(c)}\}_c$ be any unbounded sequence in $\mathbb{Z}_+^{\mathcal{J}}$ with $\|q^{(c)}\|_1 = c$. From this, we define $\bar{q}^{(c)} = q^{(c)}/c$.

Lemma 7 *If $\bar{q}^{(c)} \rightarrow q$ as $c \rightarrow \infty$ then, for each $j \in \mathcal{J}$ with $q_j > 0$,*

$$\hat{\sigma}_j(q^{(c)}) \rightarrow \sigma_j^*(q)$$

as $c \rightarrow \infty$.

Proof We can take an appropriate subsequence and without loss of generality, we may assume that the bounded sequence $\{\hat{\sigma}(q^{(c)})\}_c$ converges to some value $\hat{\sigma}$. In addition, we take M such that $M \geq \sigma_j$ for all $\sigma \in \mathcal{S}$ and $j \in \mathcal{J}$ and define $\kappa \in \mathbb{R}^{\mathcal{J}}$ by

$$\kappa_j = \begin{cases} M, & \text{if } q_j > 0, \\ 0, & \text{if } q_j = 0. \end{cases}$$

Notice, since $\bar{q}^{(c)} \rightarrow q$,

$$q^{(c)} \geq \kappa \quad (65)$$

eventually as $c \rightarrow \infty$.

By definition

$$\sum_{j \in \mathcal{J}} g_j(\hat{\sigma}_j) q_j^\alpha \leq \sum_{j \in \mathcal{J}} g_j(\sigma_j^*(q)) q_j^\alpha. \quad (66)$$

By strict concavity, $g_j(p s_j^1 + (1-p)s_j^2) q_j > p g_j(s_j^1) q_j + (1-p) g_j(s_j^2) q_j$, the optimal solution $\sigma^*(q)$ is unique for each component j with $q_j > 0$. We define G^* by

$$G^* = \max_{s \in \langle \mathcal{S} \rangle} \sum_{j \in \mathcal{J}} g_j(s_j) q_j^\alpha. \quad (67)$$

By definition $\sigma^*(q)$ is optimal for the above optimization. Also, notice the above optimum remains equal to G^* if we maximize over $\langle \mathcal{S}_\kappa \rangle$ instead of $\langle \mathcal{S} \rangle$. Thus

$$\begin{aligned} G^* &= \max_{\sigma \in \langle \mathcal{S}_\kappa \rangle} \sum_{j \in \mathcal{J}} g_j(\sigma) q_j^\alpha = \lim_{c \rightarrow \infty} \max_{\sigma \in \langle \mathcal{S}_\kappa \rangle} \sum_{j \in \mathcal{J}} g_j(\sigma_j) (\bar{q}_j^{(c)})^\alpha \\ &\leq \lim_{c \rightarrow \infty} \max_{\sigma \in \langle \mathcal{S}_{q^{(c)}} \rangle} \sum_{j \in \mathcal{J}} g_j(\sigma_j) (\bar{q}_j^{(c)})^\alpha \\ &= \lim_{c \rightarrow \infty} \sum_{j \in \mathcal{J}} g_j(\hat{\sigma}_j(q^{(c)})) (\bar{q}_j^{(c)})^\alpha = \sum_{j \in \mathcal{J}} g_j(\hat{\sigma}_j) q_j^\alpha \end{aligned}$$

The first equality hold since we maximize over smaller set $\langle \mathcal{S}_\kappa \rangle$ rather than of $\langle \mathcal{S} \rangle$; the second equality holds by continuity of the objective function and compactness of

$\langle \mathcal{S}_\kappa \rangle$; the next inequality holds by monotonicity (65); and the remaining equalities hold by definition and continuity.

Thus $\hat{\sigma}$ is optimal and thus $\hat{\sigma}_j = \sigma_j^*(q)$ for all j with $q_j > 0$. So, as this limit $\hat{\sigma}_j$ is unique for all subsequences, the lemma must hold. $\square$

Lemma 8 *For every t , the random variables $\{\bar{Q}^{(c)}(t)\}_{c \in \mathbb{N}}$ are uniformly integrable.*

Proof It is sufficient to prove that these random variables are bounded above by a sequence of L_2 -bounded random variables. The queue size is bounded above by the number of arrivals, so, for all $t \geq 0$

$$\sum_{j \in \mathcal{J}} \bar{Q}_j^{(c)}(t) \leq 1 + \sum_{j \in \mathcal{J}} \sum_{s=1}^{\lceil ct \rceil} \frac{a_j(s)}{c}. \quad (68)$$

By assumption, $a_j(s)$ has finite variance; thus, the right-hand side of inequality (68) is L_2 -bounded; and so, $\sum_{j \in \mathcal{J}} \bar{Q}_j^{(c)}(t)$ is uniformly integrable. $\square$

Proposition 3 *If for $\|Q(0)\|_1 > \kappa$*

$$\mathbb{E} \left[\left| \|Q(T\|Q(0)\|_1)\|_1 - \|Q(0)\|_1 \right| Q(0) \right] < -\epsilon \|Q(0)\|_1 \quad (69)$$

then our process is positive recurrent.

Proof We now use Proposition 3 to prove positive recurrence of the event $\{\|Q(t)\|_1 \leq \kappa\}$. We consider our Markov chain, Q , at specific stopping times: $\tau_0 := 0$ and for $n \in \mathbb{N}$,

$$\begin{aligned} \tau_n &:= \tau_{n-1} + T\|Q(\tau_{n-1})\|_1, & \text{if } \|Q(\tau_{n-1})\|_1 > \kappa, \\ \tau_n &:= \tau_{n-1} + 1, & \text{if } \|Q(\tau_{n-1})\|_1 \leq \kappa, \end{aligned}$$

We define the stopping time $N = \min\{n : \tau_n \geq t\}$. Now,

$$\begin{aligned} 0 &\leq \mathbb{E}\|Q(\tau_N)\|_1 \\ &= \mathbb{E}\|Q(0)\|_1 + \mathbb{E}\left[\sum_{n=1}^N \|Q(\tau_n)\|_1 - \|Q(\tau_{n-1})\|_1\right] \\ &\leq \mathbb{E}\|Q(0)\|_1 - \epsilon \mathbb{E}\left[\sum_{n=1}^N \|Q(\tau_{n-1})\|_1 \mathbb{I}[\|Q(\tau_{n-1})\|_1 > \kappa]\right] \\ &\quad + \|\bar{a}\|_1 \mathbb{E}\left[\sum_{n=1}^N \mathbb{I}[\|Q(\tau_{n-1})\|_1 \leq \kappa]\right] \end{aligned} \quad (70)$$

$$\begin{aligned} &= \mathbb{E}\|Q(0)\|_1 - \frac{\epsilon}{T} \mathbb{E}\left[\sum_{n=1}^N (\tau_n - \tau_{n-1}) \mathbb{I}[\|Q(\tau_{n-1})\|_1 > \kappa]\right] \\ &\quad + \|\bar{a}\|_1 \mathbb{E}\left[\sum_{n=1}^N \mathbb{I}[\|Q(\tau_{n-1})\|_1 \leq \kappa]\right] \end{aligned} \quad (71)$$

$$= \mathbb{E}\|Q(0)\|_1 - \frac{\epsilon}{T} \mathbb{E}[\tau_N] + \left(\|\bar{a}\|_1 + \frac{\epsilon}{T}\right) \mathbb{E}\left[\sum_{n=1}^N \mathbb{I}[\|Q(\tau_{n-1})\|_1 \leq \kappa]\right] \quad (72)$$

$$\leq \mathbb{E}\|Q(0)\|_1 - \frac{\epsilon}{T} t + \left(\|\bar{a}\|_1 + \frac{\epsilon}{T}\right) \mathbb{E}\left[\sum_{s=1}^t \mathbb{I}[\|Q(s)\|_1 \leq \kappa]\right]. \quad (73)$$

In inequality (70), we condition on the event $\{\|Q(\tau_{n-1})\|_1 > \kappa\}$ and apply the multiplicative Foster's condition (69); for equality (71), we observe that, by definition, $\tau_n - \tau_{n-1} = T\|Q(\tau_{n-1})\|_1$ and then adding appropriate terms, we get an interpolating sum for Eq. (72); and finally, for (73), we observe the sequence $\|Q(\tau_n)\|_1$, for $\tau_n < t$, must have hit below κ fewer times than the sequence $\|Q(s)\|_1$, $s \leq t$.

Rearranging the positive expression (73), dividing by t , and taking limits, we gain

$$\liminf_{t \rightarrow \infty} \frac{1}{t} \mathbb{E}\left[\sum_{s=1}^t \mathbb{I}[\|Q(s)\|_1 \leq \kappa]\right] \geq \frac{\epsilon}{T\|\bar{a}\|_1 + \epsilon} > 0.$$

Thus, Q is positive recurrent because it is positive recurrent in the finite set of states $\{q \in \mathbb{Z}_+^J : \|q\|_1 \leq \kappa\}$. $\square$

Proof of Lemma 2

Proof We take the partial derivative of $H(x)$ with respect to x_{jr} from first principles. For $h > 0$, define $x_{jr}^h = x_{jr} + h$ and $x_{j'r'}^h = x_{j'r'}$ for $r' \neq r$ and $j' \neq j$. Then

$$\begin{aligned}
 & \frac{H(x^h) - H(x)}{h} \\
 &= \frac{1}{h} \left[\sum_{r \in \mathcal{R}} \sum_{j \in r} x_{jr}^h \log \left(\frac{x_{jr}^h \sigma_j^*(q^h)}{q_j^h \bar{a}_r} \right) - \sum_{r \in \mathcal{R}} \sum_{j \in r} x_{jr} \log \left(\frac{x_{jr} \sigma_j^*(q)}{q_j \bar{a}_r} \right) \right] \\
 &\geq \frac{1}{h} \left[\sum_{r \in \mathcal{R}} \sum_{j \in r} x_{jr}^h \log \left(\frac{x_{jr}^h \sigma_j^*(q)}{q_j^h \bar{a}_r} \right) - \sum_{r \in \mathcal{R}} \sum_{j \in r} x_{jr} \log \left(\frac{x_{jr} \sigma_j^*(q)}{q_j \bar{a}_r} \right) \right] \\
 &= \log \left(\frac{\sigma_j^*(q)}{\bar{a}_r} \right) + \frac{1}{h} (x_{jr}^h \log x_{jr}^h - x_{jr} \log x_{jr}) + \frac{1}{h} (q_j^h \log q_j^h - q_j \log q_j).
 \end{aligned}$$

In the inequality above, we use the fact that $\sigma_j^*(q)$ is suboptimal for a proportional fair problem with parameter choice q^h . Applying the same argument but instead observing that $\sigma_j^*(q^h)$ is suboptimal for parameter choice q , we can also derive the bound

$$\frac{H(x^h) - H(x)}{h} \leq \log \left(\frac{\sigma_j^*(q^h)}{\bar{a}_r} \right) + \frac{1}{h} (x_{jr}^h \log x_{jr}^h - x_{jr} \log x_{jr}) \quad (74)$$

$$+ \frac{1}{h} (q_j^h \log q_j^h - q_j \log q_j). \quad (75)$$

We remark that these same two bounds can also be derived in the case $h < 0$. The function $q \mapsto \sigma_j(q)$ is a continuous function for $q_j > 0$. This is proven in [21, Lemma A.3]. Thus, taking the limit as $h \rightarrow 0$, both of these bounds together imply

$$\frac{\partial H(x)}{\partial x_{jr}} = \log \left(\frac{\sigma_j^*(q)}{\bar{a}_r} \right) + \log x_{jr} - \log q_j.$$

as required. $\square$

References

1. Andrews, M., Kumaran, K., Ramanan, K., Stolyar, A., Vijayakumar, R., Whiting, P.: Scheduling in a queueing system with asynchronously varying service rates. *Probab. Eng. Inform. Sci.* **18**(2), 191–217 (2004)
2. Baskett, F., Chandy, K.M., Muntz, R.R., Palacios, F.G.: Open, closed, and mixed networks of queues with different classes of customers. *J. ACM* **22**(2), 248–260 (1975)
3. Bonald, T., Massoulié, L.: Impact of fairness on internet performance. *Proc. ACM Sigmetrics* **29**, 82–91 (2001)
4. Bramson, M.: Convergence to equilibria for fluid models of head-of-the-line proportional processor sharing queueing networks. *Queueing Syst.* **23**(1–4), 1–26 (1996)
5. Bramson, M.: Stability of queueing networks. *Probab. Surv.* **5**, 169–345 (2008)
6. Bui, L., Srikant, R., Stolyar, A.L.: A novel architecture for reduction of delay and queueing structure complexity in the back-pressure algorithm. *IEEE/ACM Trans. Netw.* **19**(6), 1597–1609 (2011)
7. Cover, T.M., Thomas, J.: *Elements of Information Theory*. Wiley, New York (1991)
8. Dai, J.: 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.: Stability of open multiclass queueing networks via fluid models. *IMA Vol. Math. Its Appl.* **71**, 71–71 (1995)
10. De Veciana, G., Lee, T., Konstantopoulos, T.: Stability and performance analysis of networks supporting elastic services. *IEEE/ACM Trans. Netw.* **9**(1), 2–14 (2001)
11. Dieker, A.B., Shin, J.: From local to global stability in stochastic processing networks through quadratic lyapunov functions. *Math. Oper. Res.* **38**(4), 638–664 (2013)
12. Eryilmaz, A., Srikant, R., Perkins, J.: Stable scheduling policies for fading wireless channels. *IEEE/ACM Trans. Netw.* **13**(2), 411–424 (2005)
13. Hazan, E., Agarwal, A., Kale, S.: Logarithmic regret algorithms for online convex optimization. *Mach. Learn.* **69**(2), 169–192 (2007)
14. Jackson, J.R.: Jobshop-like queueing systems. *Manag. Sci.* **10**(1), 131–142 (1963)
15. Jagannathan, K., Markakis, M., Modiano, E., Tsitsiklis, J.: Queue length asymptotics for generalized max-weight scheduling in the presence of heavy-tailed traffic. In: *INFOCOM, 2011 Proceedings IEEE*, pp. 2318–2326 (2011). doi:[10.1109/INFOCOM.2011.5935049](https://doi.org/10.1109/INFOCOM.2011.5935049)
16. Ji, B., Joo, C., Shroff, N.B.: Throughput-optimal scheduling in multihop wireless networks without per-flow information. *IEEE/ACM Trans. Netw.* **21**(2), 634–647 (2013). doi:[10.1109/TNET.2012.2205017](https://doi.org/10.1109/TNET.2012.2205017)
17. Jiang, L., Walrand, J.: A distributed csma algorithm for throughput and utility maximization in wireless networks. *IEEE/ACM Trans. Netw. (TON)* **18**(3), 960–972 (2010)
18. Jonckheere, M., Lopez, S.: Large deviations for the stationary measure of networks under proportional fair allocations. *Math. Oper. Res.* **39**, 418–431 (2013)
19. Kelly, F.: Charging and rate control for elastic traffic. *Eur. Trans. Telecommun.* **8**(1), 33–37 (1997)
20. Kelly, F.P.: *Reversibility and Stochastic Networks*. Cambridge University Press, Cambridge (2011)
21. Kelly, F.P., Williams, R.J.: Fluid model for a network operating under a fair bandwidth-sharing policy. *Ann. Appl. Probab.* **14**(3), 1055–1083 (2004)
22. Lu, S.H., Kumar, P.: Distributed scheduling based on due dates and buffer priorities. *IEEE Trans. Autom. Control* **36**(12), 1406–1416 (1991)
23. Massoulié, L.: Structural properties of proportional fairness: stability and insensitivity. *Ann. Appl. Probab.* **17**(3), 809–839 (2007)
24. Massoulié, L., Roberts, J.: Bandwidth sharing: objectives and algorithms. *IEEE/ACM Trans. Netw.* **10**(3), 320–328 (2002)
25. McKeown, N., Mekkittikul, A., Anantharam, V., Walrand, J.: Achieving 100% throughput in an input-queued switch. *IEEE Trans. Commun.* **47**(8), 1260–1267 (1999)
26. Mo, J., Walrand, J.: Fair end-to-end window-based congestion control. *IEEE/ACM Trans. Netw.* **8**(5), 556–567 (2000)
27. Robert, P.: *Stochastic Networks and Queues*. Springer, Berlin (2010)
28. Rybko, A., Stolyar, A.: Ergodicity of stochastic processes describing the operation of open queueing networks. *Problemy Peredachi Informatsii* **28**(3), 3–26 (1992)
29. Shah, D., Shin, J.: Randomized scheduling algorithm for queueing networks. *Ann. Appl. Probab.* **22**(1), 128–171 (2012)
30. Shah, D., Walton, N., Zhong, Y.: Optimal queue-size scaling in switched networks. *SIGMETRICS Perform. Eval. Rev.* **40**(1), 17–28 (2012)
31. Shah, D., Wischik, D.: Optimal scheduling algorithms for input-queued switches. In: *Proceedings of INFOCOM 2006, 25th IEEE International Conference on Computer Communications*, pp. 1–11 (2006). doi:[10.1109/INFOCOM.2006.238](https://doi.org/10.1109/INFOCOM.2006.238)
32. Shah, D., Wischik, D.: Fluid models of congestion collapse in overloaded switched networks. *Queueing Syst.* **69**(2), 121–143 (2011)
33. Shah, D., Wischik, D.: Switched networks with maximum weight policies: fluid approximation and multiplicative state space collapse. *Ann. Appl. Probab.* **22**(1), 70–127 (2012)
34. Stolyar, A.: Maxweight scheduling in a generalized switch: state space collapse and workload minimization in heavy traffic. *Ann. Appl. Probab.* **14**(1), 1–53 (2004)
35. Stolyar, A.L.: Large number of queues in tandem: scaling properties under back-pressure algorithm. *Queueing Syst. Theory Appl.* **67**(2), 111–126 (2011)
36. Tassiulas, L., Ephremides, A.: Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Trans. Autom. Control* **37**(12), 1936–1948 (1992)

37. Vlasiou, M., Zhang, J., Zwart, B.: Insensitivity of proportional fairness in critically loaded bandwidth sharing networks. (preprint, 2014)
38. Zhong, Y.: Resource allocation in stochastic processing networks: performance and scaling. Ph.D. Thesis, MIT (2012)
39. Zinkevich, M.: Online convex programming and generalized infinitesimal gradient ascent. In: ICML, pp. 928–936 (2003)