



**HAL**  
open science

## Optimal speed scaling for multiclass fluid queues

Olivier Brun, Henda Ben Cheikh, Balakrishna Prabhu

► **To cite this version:**

Olivier Brun, Henda Ben Cheikh, Balakrishna Prabhu. Optimal speed scaling for multiclass fluid queues. 2014. hal-01094506v1

**HAL Id: hal-01094506**

**<https://hal.science/hal-01094506v1>**

Preprint submitted on 12 Dec 2014 (v1), last revised 13 Sep 2016 (v3)

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# Optimal speed scaling for multiclass fluid queues

O. Brun<sup>\*†</sup> H. Ben Cheikh <sup>†‡</sup> B.J. Prabhu<sup>\*†</sup>

<sup>\*</sup> CNRS, LAAS, 7 avenue du colonel Roche, F-31400 Toulouse, France

<sup>†</sup> Université de Toulouse, LAAS, F-31400 Toulouse, France

<sup>‡</sup> Université de Toulouse, INSA, F-31400 Toulouse, France

## Abstract

Inspired by speed-scaling techniques used in servers, we investigate the energy-delay tradeoff in bandwidth-sharing networks in which nodes can regulate their speed according to the load of the system. Assuming that the network is initially congested, we investigate the rate allocation to the classes that drains out the network with minimum total energy and delay cost. We formulate this optimal resource allocation problem as a Markov decision process which proves to be both analytically and computationally challenging. We propose to solve this stochastic problem using a deterministic fluid approximation. For the case of a linear network with two links, we provide numerical evidences in the support of the fluid model as a good approximation to the stochastic control problem. For a single link shared by an arbitrary number of classes, using Pontryagin's Maximum Principle we show that the optimal-fluid solution follows the well-known  $c\mu$  rule and give an explicit expression for the optimal speed.

## I. INTRODUCTION

Internet traffic is expected to increase to 1 zettabyte per year in 2015 according to [2]. In order to support such a large volume of traffic, network operators are forced to continuously increase network resources which consequently increases the overall network cost and energy consumption. Thus solutions (e.g., optimal share of resources) that achieve a good trade-off between energy consumption and network performance become of utmost importance.

One of the solutions for reducing energy consumption, called dynamic speed-scaling, involves regulating the speed of the processor as a function of the the load. The optimal speed-scaling policy taking into account the energy-delay trade-off was computed for a single node in [3], [7]. Our main contribution is to generalize the application of this energy-saving technique to bandwidth-sharing networks which were introduced in [19] to model the dynamics of network flows in the Internet. Our model is thus a transposition of speed-scaling to bandwidth-sharing networks in which nodes will be able to regulate their speed in order to minimize a given performance metric.

The main question that we investigate in this paper is the following: assuming that the network is initially in a congested state, how to share the available bandwidth in order to drain out the congestion with minimum total cost which comprises of the mean response times experienced by flows as well as the energy consumption of the network.

The optimal resource-allocation problem in a bandwidth-sharing network can be cast as a Markov decision process with dimension equal to the number of classes of flows in the network. Even without energy costs the computation of closed-form optimal policy has proven to be quite challenging for general networks. Optimal allocations are known only for a few special network topologies such as single node [22], [21], linear networks [25] and star networks [23]. In most cases, the optimal policy is a priority rule that serves classes that do not use competing resources.

When the option of speed-scaling is introduced in a bandwidth-sharing network, a second level of complexity gets added to the computation of the optimal policy. In addition to knowing which classes of traffic to serve, one needs to determine the speed at which each node must operate. This makes the speed-scaling policy for bandwidth-sharing networks even more challenging.

While, in theory, the optimal policy can be computed by value-iteration, in practice, the numerical computation becomes quickly prohibitive due to the exponential relation between the number of states and the number of classes. For example, a network with 5 classes and at most 100 flows per class will result in a Markov chain of with  $101^5$  states.

Our approach to this problem is to analyze an associated fluid model in the spirit of Avram, Bertsimas, and Richard [5] which can be interpreted as a deterministic approximation of the stochastic problem. This technique has been applied in various problems of optimal control of high-dimensional Markov chains [20]. In the context of the bandwidth-sharing networks, which is the focus of this paper, the utility of this technique was demonstrated in [25].

The fluid approach involves approximating the dynamics of the Markov chain using a set of differential equations, and then computing the optimal control policy for this set of differential equations. The utility of this technique lies its ability to reduce the complexity of the problem to one for which numerical solutions can be efficiently calculated. This feature will be illustrated later when we compare the times required to compute the stochastic as well as the fluid optimal policies.

In addition to the relative ease with which it can be computed, it has been shown in [8], [26], [14] that the fluid control is asymptotically optimal when the state of the system is large. These results are especially relevant in the context of this paper. The large state can be interpreted as a congested state of the network caused by a flash-crowd or a momentaneous arrival of a large number of flows. Thus, one can expect the fluid-optimal policy for draining a network from a congested state to be close to the stochastic-optimal policy.

One of the recurrent questions in this type of approximations is how to convert the fluid-optimal policy to the stochastic setting. We shall use a version of the Discrete-review method proposed in [18]. The idea is to periodically apply the fluid-optimal policy with the initial state as the state of the stochastic system sampled at review instants.

In addition to a network, we investigate the canonical case of a single link shared by multiple classes of flows, the so called multi-class queue. The aim is to characterize and compute the policy that minimizes the total cost in delay and energy consumption for the single link case. The single-link case is the basic model to study resource allocation problems and has been widely studied in several related works [21], [13], [4].

Without energy costs, for a single-node, the optimality of the well-known  $c\mu$  rule has been shown in various ways [22], [21], [11]. Speed-scaling policies for a single server and single class of traffic was analyzed in [3] but for multi-class queues such an analysis has been lacking.

#### A. Contributions and Organization of the paper

In the next section, we shall describe the bandwidth-sharing model with speed-scaling and formulate the related fluid optimal control problem. We shall first illustrate the utility of the fluid approximation on a two-node linear network by comparing the costs and the trajectories of the stochastic-optimal and the fluid-optimal policies with discrete-review type updates. It turns out that the analytical computation of the fluid-optimal policy even in the case of the two-node network becomes cumbersome due to the fact that one has to compute the speed of each of the nodes.

In section III, we shall focus on the single node case. The bulk of our the analytical results for the fluid-optimal policies are available only for this scenario. Using Pontryagin's Maximum Principle, we show the optimality of the well-known  $c\mu$  rule, that is, to serve the classes in decreasing order of priority according to the product  $c_i\mu_i$ , where  $c_i$  is the delay cost of class- $i$  and  $\mu_i$  denotes the mean rate of flows of class- $i$ . We also shall give analytical expressions for the optimal server speed as a function of time. The utility and the applicability of the fluid-optimal policies will be shown using several numerical examples.

Finally, in section IV, we present the conclusions and future work in this direction.

## II. MODEL DESCRIPTION

#### A. Stochastic model

We consider a network of a set of  $\mathcal{L}$  links shared by a set  $\mathcal{F}$  of flows with  $N$  being the number of flows (or classes). We denote by  $C = \{C_l\}_{l \in \mathcal{L}}$  the vector of links capacity. Class- $i$  job arrive according to a Poisson process at rate  $\lambda_i$  and have exponentially distributed sizes of mean  $1/\mu_i$ ,  $i \in \mathcal{F}$ . We denote by  $\rho_i = \lambda_i/\mu_i$  the traffic load of class  $i$ , and the total traffic load by  $\rho = \sum_{i=1}^N \rho_i$ . In the following, we let also  $X_i(t)$  be the number of class- $i$  jobs in the system at time  $t$ . We define  $\mathbf{X}(t) = (X_1(t), X_2(t), \dots, X_N(t))$  as the state of the system at time  $t$ , and note that the state belongs to  $\mathbb{N}^N$ . A holding cost rate of  $c_i$  is associated with flows of class- $i$ . Let  $p_i \subset \mathcal{L}$  denote the set of links used by flow  $i$ , and let  $P$  be the  $\mathcal{F} \times \mathcal{L}$  incidence matrix in which the entry  $(i, j)$  is 1 if flow  $i$  uses link  $j$  and is 0 otherwise.

Link  $l$ ,  $l \in \mathcal{L}$  can work at maximum speed  $C_l > \sum_i \rho_i p_{i,l}$ . We assume that its service rate can be selected in the interval  $[0, C_l]$  and can be changed when the number of jobs in the system changes, i.e., when the state  $\mathbf{x}(t)$  changes. In addition, we can decide the capacity allocated to each class of jobs. We let  $\mathbf{u}_i(\mathbf{x})$  be the capacity allocated to class  $i$  when the system is in state  $\mathbf{x}$ . Note that  $\mathbf{X}$  is a random variable whereas  $\mathbf{x}$  is a particular value that this random variable can take. We further assume that all class- $i$  jobs share the capacity allocated to the class according to the PS discipline. In the following, we denote by  $\mathbf{u}(\mathbf{x})$  the vector  $(\mathbf{u}_1(\mathbf{x}), \mathbf{u}_2(\mathbf{x}), \dots, \mathbf{u}_N(\mathbf{x}))$ . We note that this vector belongs to the following set

$$\mathcal{U} = \{\mathbf{u} \in \mathbb{R}_+^N : \mathbf{u}P \leq C\} \quad (1)$$

Assume that the power required to operate a link at rate  $x$  is proportional to  $x^\gamma$  where  $\gamma > 1$ . The total power requirement of the rate allocation vector  $\mathbf{u}$  is then

$$\sum_{l \in \mathcal{L}} \left( \sum_{j \in \mathcal{F}} u_l P_{l,j} \right)^\gamma \quad (2)$$

For any given stationary policy  $\mathbf{u}$ , the network state  $\mathbf{X}(t)$  is a Markov process (a multi-dimensional birth-and-death process) with transition rates

$$q(\mathbf{x}, \mathbf{y}) = \begin{cases} \lambda_i & \text{if } \mathbf{y} = \mathbf{x} + \mathbf{e}_i, i \in \mathcal{F} \\ \mu_i \mathbf{u}_i(\mathbf{x}) & \text{if } \mathbf{y} = \mathbf{x} - \mathbf{e}_i, i \in \mathcal{F} \end{cases}$$

Let  $\mathcal{S}$  denote the set of states which correspond to the ideal operating point of the network. This could be the set of states within which the network operator would like the number of flows to be so that each flow could receive a given quality of service. We shall assume that at time 0 the network finds itself in a state  $\mathbf{x}(0)$  which is far away from this  $\mathcal{S}$ . This could happen due to unseen circumstances such as flash-crowds or events which attract a large number of users at the same time (a major sports event, or some unexpected news event. (see Figure 1)

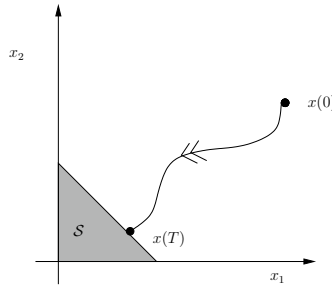


Fig. 1. Illustration of the optimal control problem with 2 classes.

The goal of the operator is to find the capacity allocation policy that will bring the network to a state within  $\mathcal{S}$  while minimizing the total cost incurred. This cost comprises of two conflicting components: one for the response time of tasks and the other for the energy consumption in the network. Equivalently, the goal of the operator is to bring the network state to a reasonable operating point with minimum total cost. Formally, we aim to find  $\mathbf{u}^* : \mathbb{N}^N \rightarrow \mathcal{U}$  solving the following problem:

$$\text{Minimize } \mathbb{E}_{\mathbf{x}(0)} \left\{ \int_0^T f(\mathbf{X}(t), \mathbf{u}(\mathbf{X}(t))) dt \right\}, \quad (3)$$

where  $T$  is the first time the process enters the set  $\mathcal{S}$ , that is

$$T = \inf\{t : \mathbf{X}(t) \in \mathcal{S} | \mathbf{X}(0) = \mathbf{x}(0)\},$$

and  $f(\mathbf{x}, \mathbf{u})$  represents the cost rate in state  $\mathbf{x}$  when a control  $\mathbf{u}$  is applied, that is,

$$f(\mathbf{x}, \mathbf{u}) = \sum_{i \in \mathcal{F}} c_i x_i + \kappa \sum_{j \in \mathcal{L}} \left( \sum_{i \in \mathcal{F}} u_i P_{i,j} \right)^\gamma. \quad (4)$$

In (4),  $\kappa$  is a parameter controlling the relative weights of energy consumption and delay, whereas  $\mathbf{c}$  is a vector giving the relative weights of the delays of the classes .

**Remark 1.** *The above problem is equivalent to a stochastic shortest path problem (cf. [9]) in which the cost in the states belonging to  $\mathcal{S}$  is 0 and the rate of transition from these states to any other state is 0 as well. The horizon of the problem now becomes infinite. We shall use this infinite horizon formulation from here on.*

**Remark 2.** *To simplify the analysis, we shall take  $\mathcal{S} = \{\mathbf{0}\}$ , that is the state where all classes have zero flows.*

Problem (3) can be transformed into an equivalent discrete-time problem by uniformizing this CTMC. The sojourn time in state  $\mathbf{x}$  is exponentially distributed with parameter  $\nu(\mathbf{x}) = \sum_i \lambda_i + \sum_i \mu_i u_i(\mathbf{x})$ . We note that  $\nu(\mathbf{x}) \leq \nu$  for all  $\mathbf{x} \in \mathbb{N}^N$ , where

$$\nu = \sum_i \lambda_i + \max_{l \in \mathcal{L}} (C_l) \sum_i \mu_i. \quad (5)$$

The uniformized version of the CTMC is a DTMC with state space  $\mathbb{N}^N$  and transition probabilities

$$p(\mathbf{x}, \mathbf{y}) = \begin{cases} \frac{\lambda_i}{\nu} & \text{if } \mathbf{y} = \mathbf{x} + \mathbf{e}_i, i \in \mathcal{F} \\ \frac{\mu_i}{\nu} u_i(\mathbf{x}) & \text{if } \mathbf{y} = \mathbf{x} - \mathbf{e}_i, i \in \mathcal{F} \\ \frac{1}{\nu} [\nu - \sum_i (\lambda_i + \mu_i u_i(\mathbf{x}))] & \text{if } \mathbf{y} = \mathbf{x}. \end{cases}$$

The equivalent discrete-time problem is as follows:

$$\text{Minimize } J_\phi(\mathbf{x}(0)) = \frac{1}{\nu} \mathbb{E} \left\{ \sum_{k=0}^{\infty} f(\mathbf{x}(k), \mathbf{u}(\mathbf{x}(k))) \right\}. \quad (6)$$

In theory, the optimal policy can then be computed using the value-iteration algorithm [9], that is by iterating over

$$J^{(k+1)}(\mathbf{x}) = \frac{1}{\nu} \min_{\mathbf{u} \in \mathcal{U}} \left[ f(\mathbf{x}, \mathbf{u}) + \sum_{\mathbf{y}} p(\mathbf{x}, \mathbf{y}) J^{(k)}(\mathbf{y}) \right], \quad (7)$$

for all  $\mathbf{x} \in \mathbb{N}^N$  and for  $k = 0, 1, 2, \dots$ . In practice however, this approach, that requires computing at each step  $k$  the optimal cost-to-go  $J^{(k)}(\mathbf{x})$  for each possible initial state  $\mathbf{x}$ , suffers from the curse of dimensionality. If truncation of the state space  $\mathbb{N}^N$  is feasible in light traffic regimes, the above technique is prohibitively expensive when either the number of flow classes or the traffic intensities get large.

### B. Fluid model

The stochastic control problem (3) is both analytically and computationally hard. Therefore, we shall study instead the fluid approximation of the original stochastic network, that is we shall approximate it by its fluid analog that has the following deterministic and continuous dynamics

$$\dot{\mathbf{x}}(t) = \boldsymbol{\lambda} - \boldsymbol{\mu} \cdot \mathbf{u},$$

where  $x_i(t)$  represents the quantity of fluid associated to class  $i$  at time  $t$ ,  $u_i(t)$  is the rate allocated to this class at that time and  $\cdot$  stands for componentwise multiplication. We describe below the associated fluid optimal control problem.

The cost of the policy  $\mathbf{u}$  starting from state  $\mathbf{x}_0$  on a horizon  $[0, T]$  is defined as the sum of the holding cost of the fluid and the energy cost for processing the fluid, and is given by

$$J(\mathbf{u}; \mathbf{x}_0) = \int_0^T f(\mathbf{x}(t), \mathbf{u}(t)) dt.$$

We note that the horizon  $T$  is not fixed, and it is also a part of the solution. It is the first time that the process enters the set  $\mathcal{S} = \{\mathbf{0}\}$ .

The fluid control problem is then to find the rate allocation that drains out the queue with minimum total cost. Formally, the problem can be stated as

$$\text{minimize } J(\mathbf{u}; \mathbf{x}_0) \tag{OPT}$$

subject to

$$\dot{x}_i(t) = \lambda_i - \mu_i u_i(t), \quad i \in \mathcal{F}, \tag{8}$$

$$-x_i(t) \leq 0, \quad i \in \mathcal{F}, \tag{9}$$

$$\mathbf{u}(t) \in \mathcal{U}, \tag{10}$$

$$\mathbf{x}(0) = \mathbf{x}_0, \tag{11}$$

$$\mathbf{x}(T) = \mathbf{0}. \tag{12}$$

### C. Implementation of the fluid policy

In this section, we discuss how to use the control policy  $\mathbf{u}(t)$  minimizing (OPT) in order to define an implementable rule for the stochastic network. Roughly speaking, the problem is that the stochastic process can significantly deviate from its fluid limit. Hence applying the open-loop fluid control law without taking into consideration the state of the network

is not an appropriate approach. In recent years, several approaches have been proposed for translating the solution of the associated fluid control problem into an implementable policy for the original stochastic network. In this work, we shall consider the family of so-called discrete-review policies [16], [18]. Those policies step through time in large intervals within which a deterministic control rule is employed. More precisely, network state is reviewed at discrete points in time. At each review instant, the controller formulates a plan for the period of time until the next review instant in order to best track what the fluid control policy would do starting at that point. In [18], Maglaras proposes a discrete-review policy for scheduling stochastic networks that achieves asymptotically optimal performances under fluid scaling, and guarantees stability if the traffic intensity is less than one at each station. See also [10], [6], [24], [15] for related approaches.

In our work, we have used discrete-review policies in their simplest form. Let  $t_k = k \Delta T$  for  $k = 0, 1, 2, \dots$  be the review instants. At time  $t_k$ , we compute the fluid control policy using as initial state the current state of the stochastic network. Dividing the review period in  $m$  steps, we allocate the rate  $u_i(t_k + \frac{n}{m} \Delta T)$  to class  $i$  in the time interval  $[t_k + \frac{n}{m} \Delta T, t_k + \frac{n+1}{m} \Delta T)$ .

In order to illustrate the benefits that can be expected from such an approach, consider the simple 2 links and 3 classes linear network shown in Figure 2. The link capacities are  $C_1 = 10$  and  $C_2 = 20$ . We assume that  $\kappa=1$  and  $\gamma=3$  and consider two scenarios:

- Scenario 1:  $\rho_0 = 0.2$ ,  $\rho_1 = 0.3$ ,  $\rho_2 = 0.1$ ,  $c_0 = 1$ ,  $c_1 = 2$  and  $c_2 = 2$ .
- Scenario 2:  $\rho_0 = 0.4$ ,  $\rho_1 = 0.1$ ,  $\rho_2 = 0.5$ ,  $c_0 = 2$ ,  $c_1 = 1$ ,  $c_2 = 6$ .

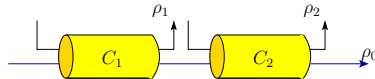


Fig. 2. Line Lot network.

Given a sample path for the flow arrival process as well as a sample path of the stochastic process governing flow sizes, one can compute the sample trajectories that would be obtained under the optimal stochastic policy and under the fluid policy reviewed at discrete times. The optimal stochastic policy is easily computed using well-known techniques such as value-iteration or policy-iteration [9]. In our case, we have used the former technique. The fluid control problem can be solved using an optimal control solver. In our case, we have used Bocop, an open source toolbox for optimal control problems [1]. We emphasize that, as clearly demonstrated by Tables II-C and II-C, the time required for solving the fluid control problem is several orders of magnitude lower than the time required to compute the optimal stochastic policy. In other words, whereas computing the optimal stochastic policy is out of reach for large networks, the computation of the optimal fluid control law requires a relatively modest effort.

Figure 3 compares sample trajectories obtained for the 3 classes under the optimal stochastic policy and under the optimal fluid control applied using the discrete review method in the first scenario and for  $\mathbf{x}(0) = (3, 5, 5)$  as initial state. Figure 4 provides a similar comparison



Scenario	Computing time (s)
1	780s
2	1020s

TABLE I  
TIME REQUIRED TO SOLVE FOR THE OPTIMAL STOCHASTIC POLICY USING VALUE ITERATION (RELATIVE PRECISION ON OPTIMAL COST-TO-GO IS  $10^{-3}$ ).

Scenario	Initial State	Computing time (s)
1	(3,5,5)	16 s
2	(4,5,3)	10 s

TABLE II  
TIME REQUIRED TO SOLVE FOR THE OPTIMAL FLUID POLICY USING BOCOP.

in the second scenario for  $x(0) = (4, 5, 3)$  as initial state. It is apparent from these two examples that the sample trajectories obtained under the fluid policy match very well those obtained under the optimal stochastic policy.

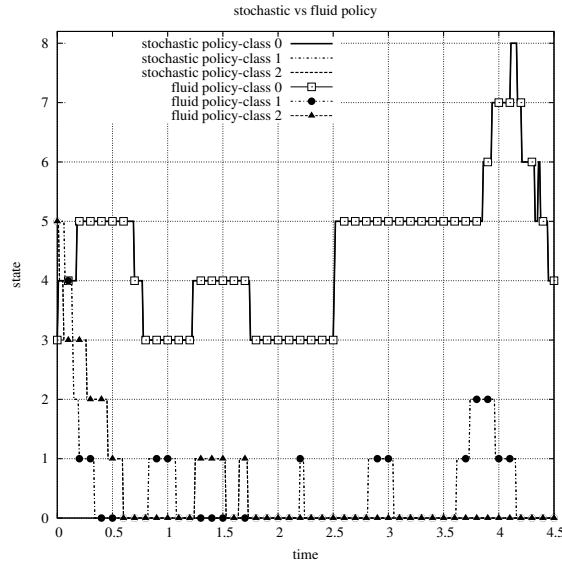


Fig. 3. Sample trajectories obtained under the optimal stochastic policy and under the fluid policy for  $x_0(0)=3$ ,  $x_1(0)=5$  and  $x_2(0)=5$ .

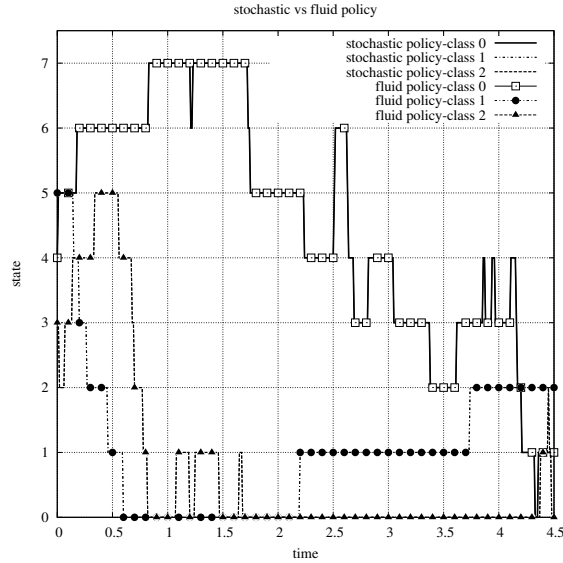


Fig. 4. Sample trajectories obtained under the optimal stochastic policy and under the fluid policy for  $x_0(0)=4$ ,  $x_1(0)=5$  and  $x_2(0)=3$ .

### III. THE SINGLE LINK CASE

In this section, we consider the case of a single node in the network with multiple classes. For this case  $P_{i,j} = 1, \forall i$ . We wish to compute the policy that minimizes the total cost and compute the optimal policy. Without energy costs, this problem was solved in [5]. It was shown that giving full priority to the non-empty class with the largest  $c\mu$  is optimal. We shall investigate how energy costs affect the priority rule and how the speed varies as a function of the fluid. We first show that in the original stochastic model, the optimal policy is a priority discipline. However given the complexity of the stochastic setting, we were not able to determine the priority criterion. Using the fluid model we show the optimality of the  $c\mu$  rule, that is to give priority to the class  $i$  with the highest  $\mu_i c_i$ .

Solving the optimal stochastic control problem numerically we observe the optimal policy coincides with the  $\mu c$  rule

#### A. Stochastic model

In the following, we denote by  $J(x)$  the optimal cost, that is the cost of an optimal policy  $u^* : \mathbb{N}^N \rightarrow \mathcal{U}$  minimizing (6). In the case of a single server, Bellman's equation takes the form

$$\begin{aligned}
J(\mathbf{x}) &= \frac{1}{\nu} \left[ \sum_i c_i x_i + \nu J(\mathbf{x}) + \sum_i \lambda_i \Delta_{\mathbf{x}+\mathbf{e}_i}(i) \right. \\
&\quad \left. + \min_{\mathbf{u} \in \mathcal{U}} \left[ \kappa \left( \sum_i u_i \right)^\gamma - \sum_i \mu_i u_i \Delta_{\mathbf{x}}(i) \right] \right], \tag{13}
\end{aligned}$$

where  $\Delta_{\mathbf{x}}(i) = J(\mathbf{x}) - J(\mathbf{x} - \mathbf{e}_i)$  is the optimal cost differential in state  $\mathbf{x}$ . An optimal policy is to use at state  $\mathbf{x}$  the vector  $\mathbf{u}$  that minimizes

$$\kappa \left( \sum_i u_i \right)^\gamma - \sum_i \mu_i u_i \Delta_{\mathbf{x}}(i)$$

Letting  $s$  be such that  $\sum_i u_i = s$ , the optimal choice is to maximize  $\sum_i \mu_i u_i \Delta_{\mathbf{x}}(i)$ , and the optimal solution is clearly  $u_k = s$  for  $k$  such that for all  $i$  we have  $\mu_k \Delta_{\mathbf{x}}(k) \geq \mu_i \Delta_{\mathbf{x}}(i)$ , whereas  $u_i = 0$  for  $i \neq k$ . We thus conclude that an optimal policy can be obtained by giving full priority to only one class.

It is difficult to obtain more precise results on the optimal speed and the priority-rule for the above described stochastic optimal control problem. Therefore, in the next section we propose to solve the related fluid control model as an approximation to the stochastic-optimal policy.

### B. Fluid model

For the single node case, the problem (OPT) takes the following form

$$\text{minimize } \int_0^T f(\mathbf{y}, \mathbf{u}) dt \tag{OPT-R}$$

subject to

$$\dot{y}_i(t) = \rho_i - u_i(t), \quad i \in \mathcal{F}, \tag{14}$$

$$-y_i(t) \leq 0, \quad i \in \mathcal{F}, \tag{15}$$

$$\mathbf{u}(t) \in \mathcal{U}, \tag{16}$$

$$\mathbf{y}(T) = \mathbf{y}_0, \tag{17}$$

$$\mathbf{y}(T) = \mathbf{0}, \tag{18}$$

where we have made the transformation  $y_i = \frac{x_i}{\mu_i}$  and  $f$  has been redefined as

$$f(\mathbf{x}, \mathbf{u}) = \sum_{i \in \mathcal{F}} \tilde{c}_i x_i + \kappa \sum_{j \in \mathcal{L}} \left( \sum_{i \in \mathcal{F}} u_i P_{i,j} \right)^\gamma, \tag{19}$$

with  $\tilde{c}_i = c_i \mu_i$  for all  $i \in \mathcal{F}$ .

The problem (OPT-R) has pure state constraints which makes the problem more difficult than the ones without state constraints. We will solve using the direct method or the direct adjoining approach [17] which consists of relaxing the pure state constraint and then finding the appropriate value of the Lagrange multiplier.

The Lagrangian for this problem is defined as

$$\mathcal{L}(\mathbf{x}, \mathbf{u}, \boldsymbol{\theta}, \boldsymbol{\eta}, \boldsymbol{\beta}) = f(\mathbf{y}, u) + \sum_i (\theta_i(\rho_i - u_i) - \eta_i y_i - \beta_i u_i),$$

where  $\theta_i$  is the adjoint variable for class  $i$ ,  $\eta_i$  corresponds to the pure state constraint  $x_i \geq 0$ , and  $\beta_i$  is for the constraint  $u_i \geq 0$ . We have ignored for the moment the constraint on the sum of the service rates.

The necessary conditions for the optimality of  $\mathbf{u}$  are

$$\dot{\theta}_i(t) = -\dot{\mathcal{L}}_{y_i} \quad (20)$$

$$\dot{y}_i(t) = \dot{\mathcal{L}}_{\theta_i} \quad (21)$$

$$\mathbf{u}(t) = \arg \max_{\mathbf{u} \in \Phi} \mathcal{L} \quad (22)$$

$$\eta_i y_i = 0, \quad \boldsymbol{\eta} \geq 0, \quad (23)$$

$$\beta_i u_i = 0, \quad \boldsymbol{\beta} \geq 0, \quad (24)$$

In addition, due to the pure state constraints there can be jumps in the trajectories of the adjoint variables at points on the boundary of the state space. Let  $\tau_k$  be the time when one of the state constraints is satisfied with equality. Then,

$$\theta(t_k+) = \theta(t_k-) - \nu_k$$

From [17], one infers that the adjoint variables are continuous if the boundary is reached non-tangentially, that is  $\dot{y}_i(t_k-) \neq 0$  when  $y_i$  reaches the boundary. We shall assume that this is indeed the case.

Ignoring for the moment the constraint on the sum of rates and using the exact form of the cost rate given in (4), the necessary condition of the optimality of the control (22) becomes

$$\dot{\mathcal{L}}_{\mathbf{u}} = 0,$$

that is,

$$\kappa\gamma \left( \sum_k u_k(t) \right)^{\gamma-1} - (\theta_k(t) + \beta_k(t)) = 0, \quad \forall k \in \mathcal{P}. \quad (25)$$

1) *Optimality of the  $c\mu$  rule:* Let  $\mathcal{P}(t) = \{i : y_i(t) > 0\}$ , i.e., it is the set of non-empty classes at time  $t$ .

**Theorem III.1.** *At any time  $t > 0$ ,*

- 1)  $u_i(t) > 0$  iff either  $y_i(t) = 0$  or  $i = \arg \max\{\tilde{c}_j : j \in \mathcal{P}(t)\}$ ;
- 2) if  $y_i(t) = 0$ , then  $u_i(t) = \rho_i$ .

The result says that amongst the classes with non-zero fluid, it is optimal to only serve the class with the largest  $\tilde{c}$ . This is the same as the  $c\mu$  rule when energy costs are not taken into account.

The proof of this theorem will be based on the following intermediate result which states that only one class in  $\mathcal{P}$  can have a non-zero service rate. The rest of the proof then consists of starting at the end time  $T$  and determining the class that receives service at time  $t$ .

**Proposition III.1.** *In any interval  $[t_1, t_2]$  such that  $\mathcal{P}(t) = \mathcal{P}(t_1)$ ,  $\forall t \in [t_1, t_2]$ ,  $u_i(t)u_j(t) = 0$ , except on a finite number of time instants.*

*Proof:* For any class  $k \in \mathcal{P}(t)$ , the complementarity conditions (23) imply that  $\eta_k(t) = 0$ . Combining this with (20), one obtains the dynamics of the adjoint variable as

$$\dot{\theta}_k(t) = -\tilde{c}_k, \forall k \in \mathcal{P}(t),$$

which upon integration gives

$$\theta_k(t) = -\tilde{c}_k t + d_k, \forall k \in \mathcal{P}(t), \forall t \in [t_1, t_2], \quad (26)$$

where  $d_k$  are constants of integration.

Let  $i$  and  $j$  be two classes in  $\mathcal{P}(t)$  which in addition also satisfy  $u_i(t) > 0$  and  $u_j(t) > 0$ . From the complementarity conditions (24), it then follows that  $\beta_i(t) = \beta_j(t) = 0$ . Substituting this in the necessary conditions for the optimal control (25), we get

$$\theta_i(t) = \theta_j(t).$$

From (26)  $\theta_i$  and  $\theta_j$  are linear functions on  $[t_1, t_2]$  with different slopes, and thus, they can intersect at most in one point within this interval. Therefore, any pair of classes can have non-zero rate simultaneously at most one point in  $[t_1, t_2]$ . Since there are finite number of pairs, the desired result follows. ■

**Corollary 1.** *There exists a  $\tau_N$  such that for all  $t \in [\tau_N, T]$  there is only one class in  $\mathcal{P}(t)$ .*

*Proof:* At  $T$  all classes reach 0, and before  $T$  two classes with non-zero fluid cannot be served simultaneously. Hence, the only way all the classes can be at 0 at  $T$  is that all classes except one were at 0 in some interval before  $T$ . ■

The above result says that in any interval with the same set of non-empty classes, only one class within this set is served at any given time except possibly for a finite number of time instants. Moreover, the interval  $[0, T]$  can be divided into intervals during which only one non-empty class is served.

*Proof of Theorem III.1:* Let the class be ordered so that  $\tilde{c}_1 > \tilde{c}_2 > \dots > \tilde{c}_N$ . We will start at time  $T$  and roll the time backwards. The proof will rely on the dynamics of the adjoint variables and the complementarity conditions.

From corollary 1, there exists a  $\tau_N$  such that only one class has non-zero fluid in  $[\tau_N, T]$ . Also, for any  $k \notin \mathcal{P}(t)$ ,  $u_i = \rho_i$ , so that all classes are served at a positive rate in  $[\tau_N, T]$ . From (25) and the complementarity conditions (24),

$$\theta_i(t) - \theta_j(t) = 0, \forall i, j \text{ and } \forall t \in [\tau_N, T], \quad (27)$$

that is

$$\dot{\theta}_i(t) - \dot{\theta}_j(t) = 0, \forall i, j \text{ and } \forall t \in (\tau_N, T), \quad (28)$$

From (20) and (23),

$$\dot{\theta}_i = -\tilde{c}_i + \eta_i(t), \quad i \notin \mathcal{P}(t), \text{ and } \eta_i \geq 0, \quad (29)$$

$$\dot{\theta}_i = -\tilde{c}_i, \quad i \in \mathcal{P}(t). \quad (30)$$

The only way (28), (29), and (30) can be satisfied is if  $\eta_N = 0$ . Thus, class  $N$ , which has the smallest  $\tilde{c}$ , is the last one to finish.

We can take  $\tau_N$  to be the last time one of the other  $N - 1$  was non-zero. Then there is some  $\tau_{N-1}$  such that there are two classes with non-zero fluid in the interval  $[\tau_{N-1}, \tau_N]$ . The conditions (29) and (30) will have to be satisfied which implies that the other class with non-zero fluid can only be class- $N - 1$ . From proposition III.1 only one of two classes can be served in this interval. Next, we shall show that only class- $N - 1$  can be served.

From (20) and (23),  $\forall t \in [\tau_{N-1}, \tau_N]$ ,

$$\theta_{N-1}(t) = \tilde{c}_N(\tau_{N-1} - t) + d_N, \quad (31)$$

$$\theta_N(t) = \tilde{c}_{N-1}(\tau_N - t) + d_{N-1}. \quad (32)$$

The constant  $d_N$  and  $d_{N-1}$  are equal which follows from (27). Thus, we have

$$\theta_{N-1}(t) - \theta_N(t) = (\tilde{c}_{N-1} - \tilde{c}_N)(\tau_N - t) \geq 0, \forall t \in [\tau_{N-1}, \tau_N].$$

Let  $i$  be the class that is served and  $j$  be the class that is idle. From (22) and the complementarity conditions (24),

$$\theta_i(t) = \theta_j(t) + \beta_j(t).$$

Since  $\beta_j(t) \geq 0$ , the above equality is possible only if  $i = N - 1$  and  $j = N$ . Thus, in the interval  $[\tau_{N-1}, \tau_N]$  only one class  $N - 1$  is served.

The above arguments can be recursively applied to show that within  $\mathcal{P}(t)$  the class with the largest  $\tilde{c}$  is the only one that receives non-zero service rate.  $\blacksquare$

2) *Optimal speed:* Theorem III.1 shows that is optimal to allocate  $u_i(t) = \rho_i$  for those classes  $i$  such that  $y_i(t) = 0$ , and to follow the  $c\mu$  rule for the classes with non-zero fluid, that is to serve the class  $k$  with the largest value of  $\tilde{c}_i$  among the classes  $i$  such that  $y_i(t) > 0$ . However, Theorem III.1 does not provide the complete solution of problem (OPT-R) since it does not tell the capacity allocated to class  $k = \arg \max\{\tilde{c}_j : j \in \mathcal{P}(t)\}$ . Observe that for that class  $k$  it hold that  $\beta_k(t) = 0$ , and hence from (25) we have

$$u_k(t) = \left( \frac{\theta_k(t)}{\kappa\gamma} \right)^{\frac{1}{\gamma-1}}, \quad (33)$$

so that, according to Theorem III.1, we can write the optimal rate allocation as follows

$$u_i(t) = \begin{cases} \left( \frac{\theta_k(t)}{\kappa^\gamma} \right)^{\frac{1}{\gamma-1}} & \text{if } i = k \\ 0 & \text{if } i \in \mathcal{P}(t) \setminus \{k\} \\ \rho_i & \text{if } i \notin \mathcal{P}(t) \end{cases} \quad (34)$$

In order to determine an optimal control law  $(\mathbf{u}(t), T)$  we therefore need to solve for the state trajectory  $\mathbf{y}(t)$ , the vector of adjoint variables  $\boldsymbol{\theta}(t)$  and the final time  $T$ . We observe that (20)-(21) provides  $2N$  differential equations, so that the vectors  $\mathbf{y}(t)$  and  $\boldsymbol{\theta}(t)$  are known at each time instant  $t$  up to a constant. Including the final time  $T$ , there are thus  $2N+1$  constants to be determined. Observe that if we can determine one of the constants, the values of the remaining constants can be inferred from the  $2N$  boundary conditions  $y_i(0)$  and  $y_i(T) = 0$ ,  $i = 1, \dots, N$ . The value of  $\theta_N(T)$  can be obtained from the following transversal condition [12]:

$$\mathcal{L}(\mathbf{y}(T), \mathbf{u}(T), \boldsymbol{\theta}(T), \boldsymbol{\eta}(T), \beta(T)) = 0 \quad (35)$$

Since  $y_i(T) = 0$  for all  $i$  and  $u_i(T) = \rho_i$  for all  $i \neq N$ , it yields

$$\kappa \left[ \sum_{i \neq N} \rho_i + u_N(T) \right]^\gamma + \theta_N(T) [\rho_N - u_N(T)] = 0. \quad (36)$$

Based on (36), Proposition III.2 gives  $u_N(T)$  as the unique root of a non-linear equation.

**Proposition III.2.** *For all  $x \geq 0$ , define  $g(x)$  as follows:*

$$g(x) = \kappa \left[ \sum_{i \neq N} \rho_i + x \right]^\gamma + \kappa \gamma x^{\gamma-1} [\rho_N - x]. \quad (37)$$

*Then  $u_N(T)$  is the unique root of  $g(x) = 0$  in  $[0, \infty)$ . Moreover,  $u_N(T) > \sum_i \rho_i / (\gamma - 1)$ .*

*Proof:* According to (33), equation (36) can be written in terms of  $u_N(T)$  only as  $g(u_N(T)) = 0$ . Obviously, we have  $u_N(T) > 0$ . It remains to show that  $u_N(T)$  is the only solution of the equation  $g(x) = 0$  in  $[0, \infty)$ . The proof relies on the following lemma.

**Lemma III.1.** *Let  $\mathcal{F}$  be the set of continuous function  $f$  such that  $f(0) > 0$ ,  $\lim_{x \rightarrow \infty} f(x) = -\infty$  and such that the equation  $f(x) = 0$  has a unique root in  $[0, \infty)$ . Then  $F(y) = a + \int_0^y f(x) dx$  belongs to  $\mathcal{F}$  for any  $a > 0$  and any  $f \in \mathcal{F}$ . Furthermore, if  $x^*$  and  $y^*$  are the unique roots in  $[0, \infty)$  of  $f(x) = 0$  and  $F(y) = 0$ , respectively, it holds that  $y^* > x^*$ .*

*Proof:* As an antiderivative of  $f$ ,  $F$  is continuous and it also satisfies  $F(0) = a > 0$ . Moreover, since  $f(x) \xrightarrow{x \rightarrow \infty} -\infty$ , there exists  $\hat{x}$  such that  $f(x) < -1$  for all  $x > \hat{x}$ . This implies that

$$F(y) < \left( a + \hat{x} + \int_0^{\hat{x}} f(x) dx \right) - y, \quad \forall y > \hat{x},$$

and hence that  $F(y) \xrightarrow{y \rightarrow \infty} -\infty$ . Finally, it follows from  $F'(y) = f(y)$  that  $F$  is increasing in  $[0, x^*)$  and strictly decreasing in  $[x^*, \infty)$ , implying that the equation  $F(y) = 0$  has a unique root  $y^*$ . Since  $F(x^*) \geq a$ , we necessarily have  $y^* > x^*$ . ■

Assuming that  $\gamma$  is integer, consider the  $(\gamma - 1)^{\text{th}}$  derivative of  $g(x)$ :

$$g^{(\gamma-1)}(x) = \kappa\gamma! \left[ \sum_{i=1}^N \rho_i + (1 - \gamma)x \right].$$

We have  $g^{(\gamma-1)}(0) = \kappa\gamma! \sum_{i=1}^N \rho_i > 0$  and  $\lim_{x \rightarrow \infty} g^{(\gamma-1)}(x) = -\infty$  since  $\gamma > 1$ . Moreover, the unique root of  $g^{(\gamma-1)}(x) = 0$  in  $[0, \infty)$  is  $x^* = \sum_i \rho_i / (\gamma - 1)$ . Since  $g^{(\gamma-1)}(x)$  is continuous over  $[0, \infty)$ , we can assert from Lemma III.1 that  $g^{(\gamma-2)}(x)$  satisfies the same properties. By repeated applications of Lemma III.1, we conclude that  $u_N(T)$  is the only solution of the equation  $g(x) = 0$  in  $[0, \infty)$ , and that  $u_N(T) > \sum_i \rho_i / (\gamma - 1)$ . ■

Once  $u_N(T)$  computed as described in Proposition III.2, we immediately obtain from (33) that  $\theta_N(T) = \kappa\gamma u_N(T)^{\gamma-1}$ . Assuming that the value of  $T$  is known, we are now in position to recursively compute the remaining constants as described in Proposition III.3.

**Proposition III.3.** *Assume that  $T$  is given and define for convenience  $\tau_{N+1} = T$  and  $\alpha = 1/(\gamma - 1)$ . For  $k = 0, \dots, N - 1$ , the service rate of class  $N - k$  in  $[\tau_{N-k}, \tau_{N-k+1}]$  is*

$$u_{N-k}(t) = \left( \frac{\tilde{c}_{N-k}(\tau_{N-k+1} - t) + \theta_{N-k}(\tau_{N-k+1})}{\kappa\gamma} \right)^{\frac{1}{\gamma-1}}, \quad (38)$$

where  $\tau_{N-k}$  is given by

$$\tau_{N-k} = \tau_{N-k+1} - \frac{\theta_{N-k}(\tau_{N-k+1})}{\tilde{c}_{N-k}} (V_{N-k} - 1), \quad (39)$$

with

$$V_{N-k} = \left[ 1 + \frac{(\alpha + 1)(\kappa\gamma)^\alpha \tilde{c}_{N-k}}{(\theta_{N-k}(\tau_{N-k+1}))^{\alpha+1}} (y_{N-k}(0) + \rho_{N-k} \tau_{N-k+1}) \right]^{\frac{1}{\alpha+1}}.$$

Moreover, for  $k = 1, \dots, N - 1$ , it hold that

$$\theta_{N-k}(\tau_{N-k+1}) = \theta_{N-k+1}(\tau_{N-k+2}) + \tilde{c}_{N-k+1}(\tau_{N-k+2} - \tau_{N-k+1}). \quad (40)$$

*Proof:* From (30) and (33), we know that



$$u_N(t) = \left( \frac{\tilde{c}_N(\tau_{N+1} - t) + \theta_N(\tau_{N+1})}{\kappa\gamma} \right)^{\frac{1}{\gamma-1}},$$

for  $t$  in  $[\tau_{N-k}, \tau_{N-k+1}]$ . In order to determine  $\tau_N$ , we can use the terminal condition  $y_N(\tau_{N+1}) = 0$  and (21) to obtain

$$y_N(\tau_{N+1}) = 0 = y_N(0) + \rho_N \tau_{N+1} - \int_{\tau_N}^{\tau_{N+1}} u_N(t) dt,$$

and, after some elementary algebra, it is readily verified that  $\tau_N$  is indeed given by (39) when  $k = 0$ . Once  $\tau_N$  is known, we are in position to compute  $\theta_{N-1}(\tau_N)$ . Indeed, we have from (30) that  $\theta_N(t) = \tilde{c}_N(\tau_{N+1} - t) + \theta_N(\tau_{N+1})$  for  $t$  in  $[\tau_{N-k}, \tau_{N-k+1}]$ . We also know from (27) that  $\theta_N(t) = \theta_{N-1}(t)$  in this time interval. We thus deduce that

$$\theta_{N-1}(\tau_N) = \theta_N(\tau_{N+1}) + \tilde{c}_N(\tau_{N+1} - \tau_N),$$

as stated in (40). The above arguments can be recursively applied to show that  $\tau_{N-k}$  and  $\theta_{N-k}(\tau_{N-k+1})$  are indeed given by (39) and (40), respectively, for all  $k = 1, \dots, N-1$ . ■

Lemma III.2 shows that finally the value of  $T$  can be computed as the unique value such that  $\tau_1 = 0$ .

**Lemma III.2.** *There is a unique value of  $T$  such that  $\tau_1 = 0$ .*

*Proof:* The proof relies on the fact that  $\tau_1$  is strictly increasing in  $T$ , and for  $T = 0$ ,  $\tau_1 < 0$ , while  $\tau_1 \rightarrow \infty$  for  $T \rightarrow \infty$ . ■

### C. Numerical results

We now describe the numerical results we have obtained in different scenarios. For each scenario, we compare the fluid policy with the optimal stochastic policy obtained by solving a stochastic shortest path problem with value iteration.

#### First Scenario

As a first scenario, we consider a single link shared by two classes. We assume that  $\kappa=0.5$  and  $\gamma=3$ . The parameters of the traffic classes are as follows:  $\rho_1=0.3$ ,  $\rho_2=0.5$ ,  $\mu_1 c_1=12$  and  $\mu_2 c_2=1$ . Table III shows the relative gap between the fluid-optimal cost and the stochastic-optimal cost according to the initial state. We observe that when the initial state gets sufficiently large, the relative gap becomes negligible.

Figure 5 shows the fluid trajectories of class 1 and 2, as well as the average state trajectories obtained for each class under the optimal stochastic policy when the initial state is  $\mathbf{x}(0) = (10, 10)$ . In Figure 6, we plot the optimal fluid rates allocated to each class. These rates are compared to those that would be obtained under the optimal stochastic policy, taking the quantity of fluid of each class as the state of the system<sup>1</sup>

<sup>1</sup>Since the optimal stochastic policy is defined only at integer points, we use bilinear interpolation to compute its value from the four neighboring integer points.

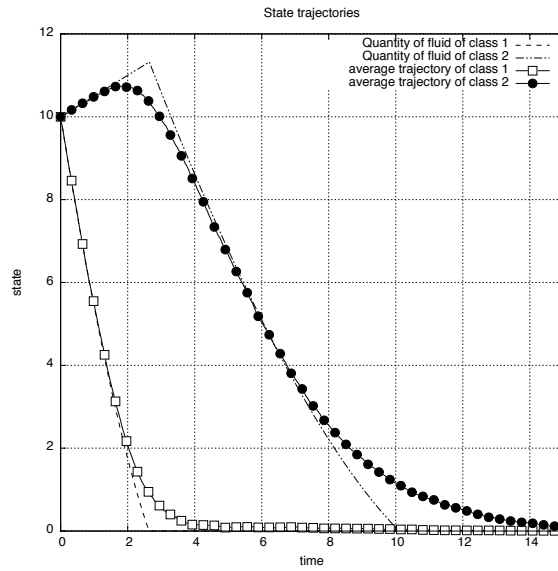


Fig. 5. Fluid trajectories vs. average state trajectories obtained in Scenario 1 under the optimal stochastic policy for  $x_1(0)=10$  and  $x_2(0)=10$ .

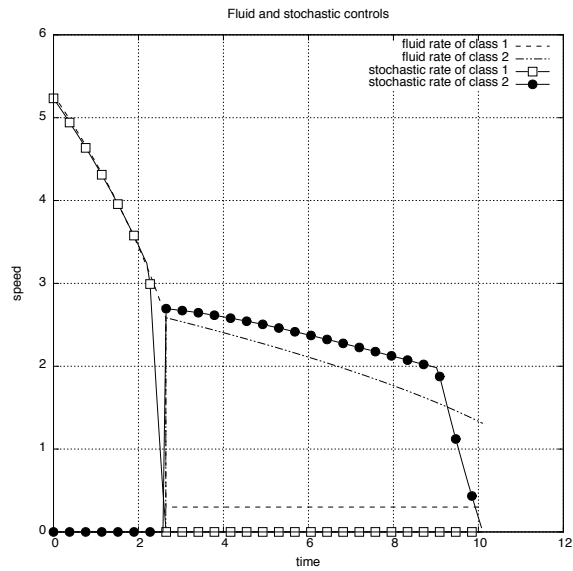


Fig. 6. Optimal fluid speed vs. optimal stochastic speed in Scenario 1 for  $x_1(0)=10$  and  $x_2(0)=10$ .

$x_1$	$x_2$			
	10	20	30	40
10	8.5	7.0	5.5	3.5
20	4.7	4.3	3.5	1.9
30	3.1	2.9	2.2	0.6
40	2.3	2.1	1.3	0.2

TABLE III  
RELATIVE GAP IN % BETWEEN THE FLUID-OPTIMAL COST AND THE STOCHASTIC-OPTIMAL COST ACCORDING TO THE INITIAL STATE.

We also compare sample trajectories obtained under the optimal stochastic policy and under the optimal fluid control applied using the discrete-review method in Figure 7 and 8. The initial state is  $\boldsymbol{x}(0) = (10, 10)$  and the instants at which the deterministic control policy is reviewed are indicated. It is clear that the sample trajectory for class 1 and 2 obtained under the fluid policy closely match those obtained under the optimal stochastic policy. It is also worth noticing that both policies follow the  $c\mu$  rule, that is they both give priority to class-1 which has a higher value of  $c_i\mu_i$  ( $c_1\mu_1 = 12 > 1 = c_2\mu_2$ ).

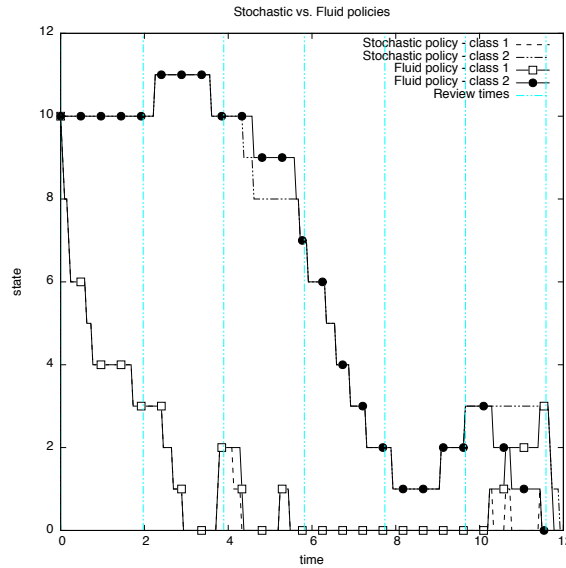


Fig. 7. Sample trajectories obtained under the optimal stochastic policy and under the fluid policy for  $x_1(0)=10$  and  $x_2(0)=10$ .

### Second Scenario

As a second scenario, we again consider the case of two classes, but this time with  $\gamma=1.5$  (the

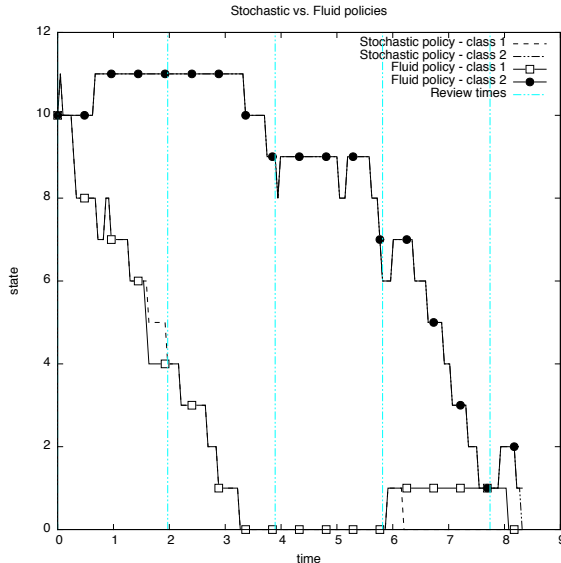


Fig. 8. Sample trajectories obtained under the optimal stochastic policy and under the fluid policy for  $x_1(0)=10$  and  $x_2(0)=10$ .

value of  $\kappa$  is still 0.5). The parameters of the traffic classes are as follows:  $\rho_1=0.4$ ,  $\rho_2=0.5$ ,  $\mu_1 c_1=4$  and  $\mu_2 c_2=1$ .

Figure 9 shows the fluid trajectories of class 1 and 2 when the initial state is  $\mathbf{x}(0) = (13, 15)$ . The fluid trajectory of each traffic class closely matches its average state trajectory under the optimal stochastic policy.

Figure 11, 12 show a sample path of class 1 and 2 under the optimal fluid control when the initial state is  $\mathbf{x}(0) = (13, 15)$ . Again, we note that these trajectories closely match those obtained under the optimal stochastic policy, and that both policies follow the  $c\mu$  rule.

#### IV. CONCLUSION

In this paper, we modeled the energy-delay tradeoff in bandwidth-sharing networks with speed-scaling. Our objective was to determine the policy minimizing the total cost to drain the network. For the case of a linear network with two links, it was shown via numerical experiments that the fluid model together with discrete-review provides a good approximation to the stochastic control problem. For a single link shared by an arbitrary number of classes, using Pontryagin's Maximum Principle we completely characterized the fluid optimal control policy. In particular, we established that the optimal fluid policy behaves according to the  $c\mu$  rule. Future work include the extension of the results obtained here for a single link to a linear network topology.

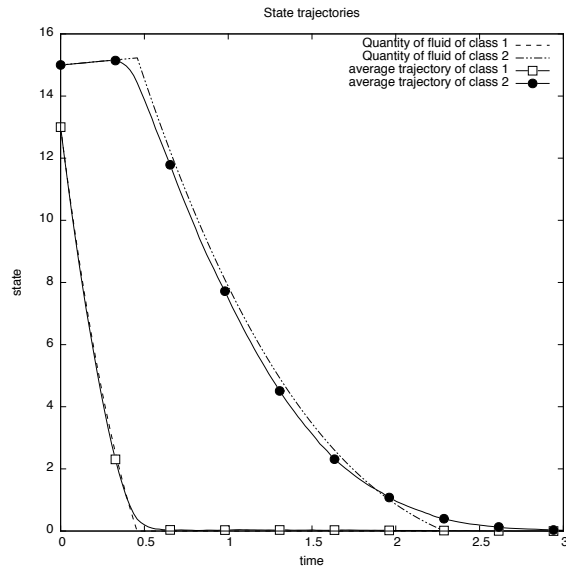


Fig. 9. Fluid trajectories vs. average state trajectories obtained in Scenario 2 under the optimal stochastic policy for  $x_1(0)=13$  and  $x_2(0)=15$ .

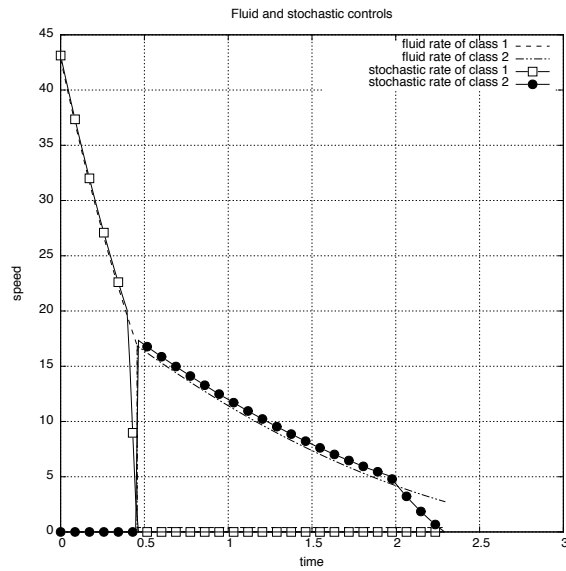


Fig. 10. Optimal fluid speed vs. optimal stochastic speed in Scenario 2 for  $x_1(0)=13$  and  $x_2(0)=15$ .

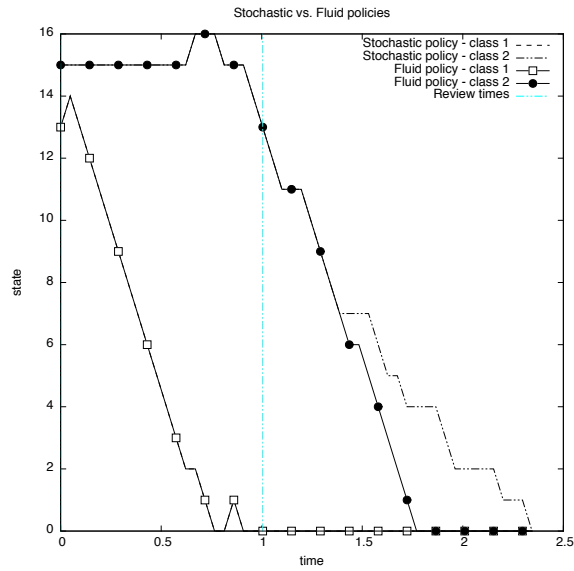


Fig. 11. Sample trajectories obtained in Scenario 2 under the optimal stochastic policy and under the fluid control for  $x_1(0)=13$  and  $x_2(0)=15$ .

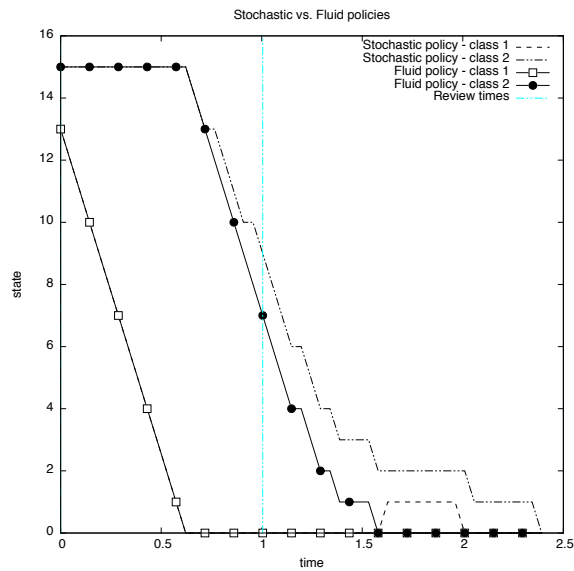


Fig. 12. Sample trajectories obtained under the optimal stochastic policy and under the fluid control for  $x_1(0)=13$  and  $x_2(0)=15$ .

## REFERENCES

- [1] Bocop - the optimal control solver, <http://bocop.saclay.inria.fr>.
- [2] Visual networking index: Global mobile data traffic forecast update, 2012-2017. *tech.rep Cisco.*, February 2013.
- [3] L.L.H. Andrew, A. Wierman, and A. Tang. Optimal speed scaling under arbitrary power functions. *SIGMETRICS Perform. Eval. Rev.*, 37(2):39–41, October 2009.
- [4] B. Ata and M. Tongarlak. On scheduling a multiclass queue with abandonments under general delay costs. *Queueing Systems*, pages 1–40, 2012.
- [5] F. Avram, D. Bertsimas, and M. Ricard. An optimal control approach to optimization of multiclass queueing networks. *Proceedings of Workshop on Queueing Networks*, 1994.
- [6] N. Bambos and J. Walrand. Scheduling and stability aspects of a general class of parallel processing systems. *Adv. in Appl. Probab.*, 25:176–202, 1993.
- [7] N. Bansal, H.L. Chan, and K. Pruhs. Speed scaling with an arbitrary power function. In *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '09, pages 693–701, Philadelphia, PA, USA, 2009. Society for Industrial and Applied Mathematics.
- [8] N. Bäuerle. Asymptotic optimality of tracking policies in stochastic networks. *Ann. Appl. Probab.*, 10:1065–1083, 2000.
- [9] D.P. Bertsekas. *Dynamic Programming and Optimal Control I*. Athena Scientific, Belmont, Massachusetts, 1995.
- [10] D. Bertsimas and G. Van Ryzin. A stochastic and dynamic vehicle routing problem in the euclidean plane. *Oper. Res.*, 39:601–615, 1991.
- [11] C. Buyukkoc, P. Varaiya, and J. Walrand. The  $\mu$ -rule revisited. *Adv. Appl. Prob.*, 17:237–238, 1985.
- [12] B. Chachuat. Nonlinear and dynamic optimization - from theory to practice. Technical report, Ecole Polytechnique Federale de Lausanne, 2006.
- [13] D.G. Down, G. Koole, and M.E. Lewis. Dynamic control of a single server system with abandonments. *Queueing Systems.*, 67(1):63–90, 2011.
- [14] A. Gajrat and A. Hordijk. Fluid approximation of a controlled multiclass tandem network. *Queueing Systems*, 35:349–380, 2000.
- [15] N. Gans and G. Van Ryzin. Optimal control of a muticlass, flexible queueing system. *Oper. Res.*, 45:677–693, 1997.
- [16] J.M. Harrison. The bigstep approach to flow management in stochastic processing networks. *Stochastic Networks: Theory and Applications (F. Kelly, S. Zachary and I. Ziedins eds)*, pages 57–90, 1996.
- [17] R.F. Hartl, S.P. Sethi, and R.G. Vickson. A survey of the maximum principles for optimal control problems with state constraints. *SIAM Review*, 37(2):181–218, June 1995.
- [18] C. Maglaras. Discrete-review policies for scheduling stochastic networks: trajectory tracking and fluid-scale asymptotic optimality. *The Annals of Applied Probability*, 10(3):897–929, 2000.
- [19] Laurent Massoulié and James Roberts. Bandwidth sharing: Objectives and algorithms. *IEEE/ACM Trans. Netw.*, 10(3):320–328, June 2002.
- [20] Sean Meyn. *Control Techniques for Complex Networks*. Cambridge University Press, 2007. Cambridge Books Online.
- [21] P. Nain and D. Towsley. Optimal scheduling in a machine with stochastic varying processing rate. *IEEE Transactions on Automatic Control*, 39:1853–1855, 1994.
- [22] R. Righter and J. G. Shanthikumar. Scheduling multiclass single server queueing systems to stochastically maximize the number of successful departures. *Probability in the Engineering and Informational Sciences*, pages 323–333, 1989.
- [23] B. Tan, L. Ying, and R. Srikan. Short-term fairness and long-term qos in the internet. *Performance Evaluation*, 67:406–414, 2010.
- [24] L. Tassiulas and S. Papavassiliou. Optimal anticipative scheduling with asynchronous transmission opportunities. *IEEE Trans. Automat. Control*, 40:2052–2062, 1995.
- [25] I.M Verloop and S.C. Borst. Heavy-traffic delay minimization in bandwidth-sharing networks. In: *Proc. IEEE Infocom*, 2007.
- [26] I.M. Verloop and R. Núñez-Queija. Asymptotically optimal parallel resource assignment with interference. *Queueing Systems*, 65:43–92, 2010.