Performance Analysis of Traffic Surges in Multi-class Communication Networks

M.T.S. Jonckheere*† R. Núñez-Queija‡§ and B.J. Prabhu¶‖
* CONICET, Departamento de Matemática, Facultad de Ciencias Exactas y Naturales, Universidad de Buenos Aires, Pabellón 1, Ciudad Universitaria, 1428 Buenos Aires, Argentina
†Eindhoven University of Technology, Postbus 513 5600 MB EINDHOVEN The Netherlands
Email: m.t.s.jonckheere@tue.nl
‡Operations Research, Faculty of Economics and Business, University of Amsterdam, The Netherlands
§CWI, P.O. Box 94079, 1090 GB Amsterdam, The Netherlands
Email: sindo@cwi.nl
¶CNRS; LAAS; 7 avenue du Colonel Roche, F-31077 Toulouse, France.
Email: bala@laas.fr
‖Université de Toulouse; UPS, INSA, INP, ISAE; LAAS; F-31077 Toulouse, France.

***Abstract*—In multi-class communication networks, traffic surges due to one class of users can significantly degrade the performance for other classes. During these transient periods, it is thus of crucial importance to implement priority mechanisms allowing the conservation of the quality of service experienced by the affected classes, while ensuring that the temporarily unstable class is not entirely neglected. In this paper, we examine – for a suitably-scaled set of parameters – the complex interaction occurring between several classes of traffic when an unstable class is penalized proportionally to its level of congestion. We characterize the evolution of the performance measures of the network from the moment the initial surge takes place until the system reaches its equilibrium. We show that, using a time-space-transition-scaling, the trajectories of the temporarily unstable class can be described by a differential equation, while those of the stable classes retain their stochastic nature. In particular, we show that the temporarily unstable class evolves at a time-scale which is much slower than that of the stable classes. Although the time-scales decouple, the dynamics of the temporarily unstable and the stable classes continue to influence one another. We further proceed to characterize the obtained differential equations for several simple network examples. In particular, the macroscopic asymptotic behavior of the unstable class allows us to gain important qualitative insights on how the bandwidth allocation affects performance.**

I. Introduction

The Internet is often confronted with traffic surges due to events that are of interest to a large number of users. For example, the number of visits to news websites more than tripled for the inauguration of President Obama [1]. Another example of traffic surges is the sudden increase in peer-to-peer downloads that coincide with the release date of a popular software or file [16].

On a smaller scale, traffic variations can result from the periodicity of user activity. In a data set collected over a period of two years from 67 ISPs (mostly in Europe and in North America) it was observed that web traffic in a geographical area reaches its peak late in the evening [20]. This surge is mainly attributed to peer-to-peer downloads and multi-platform games. On top of these easily explicable daily variations, web servers usually experience large weekly and monthly variations. Figure 1 [2] shows the variations of traffic on Wikipedia in December 2004.

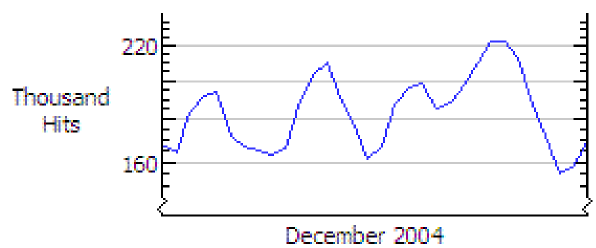

Fig. 1. Variations of traffic on Wikipedia

The impact of large-scale traffic surges, also known as slash-dot-crowds or flash-crowds, on web servers and content distribution networks has been the subject of several studies [29], [17], [12]. These mainly focus on designing mechanisms to make the content providers resilient to surges of a given type of traffic. However, in addition to overloading the content providers, a traffic surge can also negatively impact the performance of other flows in the network at that time. The temporarily unstable class can potentially starve the other classes from network capacity thereby subjecting them to unreasonable delays and packet losses. In such circumstances, in addition to protection mechanisms in web servers, it is crucial to implement bandwidth-sharing mechanisms inside the network that would protect the stable classes from the adverse effects of the surge. It seems natural that such mechanisms should penalize the temporarily unstable class more when the level of congestion it creates is larger, without actually throttling it. (Thus, the more significant the surge is, the smaller the bandwidth each flow in this class gets.) The consequences of traffic surges on the performance of the different classes in the presence of such bandwidth sharing mechanisms have not been explored much.

In this paper, we take a global view of the effects of a traffic surge in a multi-class communication network. Our aim is to present an analytic treatment of the complex interaction that takes place between the temporarily unstable class and the stable class during a traffic surge, when the temporarily unstable class is penalized *proportionally* to its level of congestion. Under appropriate scaling of the parameters of the bandwidth-sharing mechanism, we show that the dynamics of the temporarily unstable class can be described by a differential equation. On the other hand, the stable classes retain their stochastic nature. We show that a time-scale separation occurs – the temporarily unstable class evolves on a much slower time-scale compared to that of the stable classes.

The convergence of dynamics to a differential equation resembles the classical fluid limits of Markov processes (see for instance [25], [10]). There is however one important difference with the usual scalings since we do not only scale time and space, but also the transitions of the stochastic process in order to model both the surge of traffic and the network's reaction to it. In order to obtain a classical fluid limit as developed for Jackson networks in [25] or more generally for bandwidth-sharing networks in [14], all the classes are jointly scaled in time and in space. This yields a set of differential equations that govern the dynamics of all classes. Under additional assumptions on the drift δ of the considered Markov process, the differential equation is essentially of the form $\dot{x}(t) = \delta(x(t))$. In our case, we obtain two different types of dynamics – deterministic for the temporarily unstable class and stochastic for the stable classes. This difference in the dynamics is due to the scaling of the transitions of the process yielding a separation of time-scales of the two types of classes. However, even with this separation of time-scales, a strong coupling in the dynamics of the temporarily unstable and the stable classes remains. The dynamics of the temporarily unstable class, which we shall denote by class-1, is influenced by the stable classes through their stationary distribution, conditional on the level of congestion of class-1 flows being fixed to its present "macroscopic value". This macroscopic level thus depends on the dynamics of the stable classes. For class-1, the differential equation obtained is of the form $\dot{x}_1(t) = \bar{\delta}_1^t(x_1(t))$, where $\bar{\delta}_1^t$ is an average over the conditional distribution of the other classes, given $x_1(t)$.

Our contribution first consists in establishing convergence results for stochastic processes commonly used in the modeling and analysis of communications networks under the scaling considered. Second, we characterize the response of the network to the surge of traffic for different network scenarios. In particular, the limiting macroscopic state of class-1 gives valuable insight into the (macroscopic) stationary regime of the system.

The rest of the paper is organized as follows. We first discuss related work in the next section. In Section III, we introduce the stochastic processes describing the states of various communication networks. We give several examples of such models and then explain the specific time-space scaling we study. In Section IV, we present the main result. In Section V, we give numerical examples of applications of the main result to bandwidth sharing on several simple, but illustrative, network topologies. Finally, we conclude in Section VI.

II. Related work

The transient behavior of communications networks – though of crucial importance as underlined previously – has received much less attention than their stationary regimes (see for instance [18], [15], [27], [3]). There exists however a considerable body of literature based on fluid limits and ODEs (ordinary differential equations), both for general Markov processes and for communications networks [25], [10], [9], [14], [28]. In particular, the proof techniques that we use are inspired by the methodologies employed in [25], [10].

While there are similarities between our work and fluid limits and ODEs, the time-scale separation that we consider between a temporarily unstable and one or more stable classes is closer in spirit to the one observed in singularly perturbed and nearly decomposable Markov chains [30]. Most of the studies in that context, however, focus on characterization of the stationary behavior of the chain, see for example [7]. Notable exceptions are [8], [13], [24]. The study of fluid limits for queues in slowly changing environments [8] intimately relates to our results in the way that the time-scales are decoupled. In our setting, however, we do not consider a dynamically independent environment. Instead, our "slow" and "fast" processes influence each other both ways, which makes the analysis quite different. Independently, this phenomenon was recently also observed in [24].

Related transient time-scale separation phenomena have also extensively been studied in chemistry and biochemistry [26] where the kinetics of chemical reactions can be described by systems of ordinary differential equations. In this case, one of the dependent variables is assumed to be in steady-state with respect to the instantaneous values of the other dependent variables. Taking this time-scale separation as an assumption, an efficient approximation method called the quasi-steady-state is commonly used in that context.

III. Model description

We now proceed to describe the notation, the network model and the specific time-space scaling we shall be studying.

In the sequel, for $x \in \mathbb{Z}^N$, $|x|$ denotes the l_1-norm:

$$|x| = \sum_{i=1}^{N} |x_i|.$$

For $x, y \in \mathbb{Z}^N$, we also use the notation $x \leq y$ to denote the partial order $x_i \leq y_i$ for all $i = 1 \ldots N$.

A. The network model

Bandwidth-sharing network models [21], [4], [14] have become quite a standard modeling tool over the past decade for communication networks. Within each of the N traffic classes, resources are shared according to a processor-sharing service discipline. The service rates are state-dependent: they may depend on the number of flows within the same class, as well as on the numbers of flows in all the other classes. The service rates of the N traffic classes will be denoted by $\phi = (\phi_i(\cdot))_{i=1}^N$. Bandwidth-sharing networks have been used extensively to represent the flow level dynamics of data traffic in wireline and wireless networks [3] as well as for the integration of voice and data traffic [5]. Bandwidth-sharing networks generalize more traditional voice traffic models, e.g. [18]. Several examples are considered in the next section. Note that the service rate function ϕ captures the allocation of bandwidth which is determined by the specific network topology, the link rates, and the congestion control mechanisms.

We assume that class-i customers arrive according to a Poisson process of intensity λ_i and require exponentially distributed[1] service times of mean μ_i^{-1} for class i. The arrival processes of all classes are mutually independent. Our main results allow for time-varying arrival rates for the class exhibiting a traffic surge. When applicable, we reflect that in the notation by adding the time parameter to the arrival rates and then $\lambda_1(t)$ is the arrival rate of class 1 at time t. For ease of exposition, however, we restrict ourselves to constant arrival rates for all classes in this section and will formulate our results with time-varying arrival rates in Section IV.

Let $X = (X_1, \ldots, X_N) \in \mathbb{Z}^N$ be the N-dimensional stochastic process describing the numbers of flows (or calls) in progress. Thus, X_i represents the number of concurrent flows of class i. In the absence of additional priority mechanisms, X is a multi-dimensional birth-and-death process with transition rates:

$$\begin{aligned} q(x, x - e_i) &= \mu_i \phi_i(x), \\ q(x, x + e_i) &= \lambda_i, \end{aligned}$$

where the i-th coordinate of the vector $e_i \in \mathbb{Z}^N$ is 1 and all its others coordinates equal 0.
Assume now that priority mechanisms, parameterized by $r_1, \ldots, r_N$, are employed in the network such that the actual bandwidth allocation depends on $r_i x_i$ rather than on x_i alone. Hence, while x_i is a natural measure of the level of congestion of class i, a differentiation between classes can be enforced by giving different weights to the different classes. (In an actual implementation, such asymmetric capacity sharing can be enforced at the packet level through schedulers such as weighted deficit round robin.)

To avoid confusion, we emphasize that the dependence on the control parameter r_i is kept explicit on purpose, rather than defining an alternative allocation function $\tilde{\phi}(x) = \phi(r.x)$. This choice will be convenient as we shall investigate an adaptive choice of r_i, depending on the actual level of x_i. In practice, one can, for example, adapt r_i to averages of the number of flows over longer time intervals.

Additionally, our model can also represent scenarios in which each traffic class has a limited peak rate (because of access constraints, for instance). In order to meet the traffic demand, it may then be advantageous for providers to share capacity as a function of the demanded rates $r_i x_i$ rather than as a function of the actual number of flows of each class in the network. In both configurations, X can be described as a multi-dimensional birth-and-death process with transition rates:

$$\begin{aligned} q(x, x - e_i) &= \mu_i \phi_i(r.x), \\ q(x, x + e_i) &= \lambda_i, \end{aligned}$$

where $r.x = (r_i x_i)_{i=1 \ldots N}$ for some $r \in \mathbb{R}_+^N$.

[1] Such assumptions are certainly not necessary to obtain the results we are aiming at; however, a rigorous generalization would be technically very involved and is beyond the scope of the present paper.

B. Examples of topologies and traffic characteristics

1) Bandwidth sharing networks: Bandwidth sharing networks constitute a natural extension of a multi-class processor sharing queue, and have become a standard stochastic model for the flow level dynamics of Internet congestion control (they were introduced by Massoulié and Roberts in [21]).

Consider for example the tree network represented in Figure 2, with two traffic routes, each passing through a dedicated link, followed by a common link. If each dedicated link has a capacity $c_i \le 1$, $i = 1, 2$, and the common link has capacity 1, the flow on each route gets a capacity $\phi_i(x)$ that lies in the polyhedron $\mathcal{C}$:

$$\phi_i(x) \le c_i, \quad i = 1, 2, \tag{1}$$

$$\sum_{i=1}^{2} \phi_i(x) \le 1. \tag{2}$$

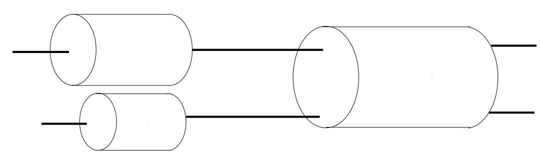

Fig. 2. Tree network

Another example of interest is the linear network represented in Figure 3 with 3 routes sharing two links. While the first route passes through both links, routes 2 and 3 only use one of the links (one each). This gives the following capacity constraints:

$$\phi_1(x) + \phi_2(x) \le c_1, \tag{3}$$

$$\phi_1(x) + \phi_3(x) \le c_2. \tag{4}$$

In general, like for the specific foregoing examples, the capacity constraints determine the space over which a network controller can choose a desired allocation function. It has been argued in [19] that a good approximation of current congestion control algorithms such as TCP (the Internet's predominant Transfer Control Protocol) can be obtained by using the weighted proportional fair allocation, which solves an optimization problem for each vector x of instantaneous numbers of flows. Specifically, the weighted proportional fair allocation $\eta(x)$ for state vector x maximizes

$$\sum_{i=1}^{N} w_i x_i \log(\eta_i), \eta \in \mathcal{C},$$

where the weights w_i are class-dependent control parameters.

Remark 1: By definition of this optimization program, if $\phi(\cdot) = \eta(\cdot)$ is the standard (unweighed) proportional fair allocation with $w_i \equiv 1$, then the allocation $\phi^r(x) = \phi(r.x)$ corresponds to the weighted proportional fair allocation with weights $w_i \equiv r_i$.

This framework has been generalized to so-called weighted α-fair allocations, which provide flexibility to model different levels of fairness in the network. Another important alternative is the balanced fair allocation [3], which allows a closed form expression for the stationary distribution of the numbers of flows in progress. In addition, the balanced fair allocation gives a good approximation of the proportional fair allocation while being easily evaluated, which is attractive for performance evaluation.

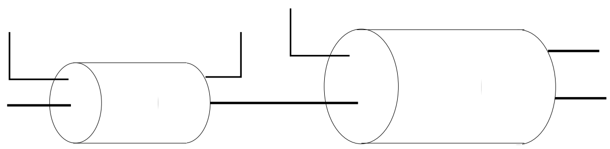

Fig. 3. Linear network

2) Integration of streaming and elastic traffic: Consider now a system where two intrinsically different types of traffic – "streaming" and "elastic" traffic – coexist and share a given link. Such models have been considered in [22], [11], [5]. It is natural to equip streaming traffic with a fixed required rate, say, c per flow. Giving priority to streaming traffic (class 2) the allocation of service may be chosen as:

$$\phi_1(x) = \max\left(\frac{r_1 x_1}{r_1 x_1 + c x_2}, 1 - c x_2\right),$$
$$\phi_2(x) = c x_2,$$

where the parameter r_1 quantifies the level of priority. The allocated capacity cannot exceed the total capacity. If the latter is normalized to 1, the state space must be restricted to states x_2 such that

$$\phi_1(x) + \phi_2(x) \leq 1.$$

Then, if the number of current streaming flows x_2 is such that $\phi_1(x+e_2) + \phi_2(x+e_2) > 1$, arriving streaming flows must be blocked from the network.

C. Modeling a traffic surge

In this paper, we look at the case where one class of flows undergoes an important and sudden rise in its activity level and the network reacts to that surge by penalizing this class. To be specific, in the sequel we study the case where:

1) the number of initial class-1 flows is of order K (we will investigate the system in the limit as $K \to \infty$),
2) we scale (accelerate) time by a factor K,
3) we scale class-1 states by a factor K,
4) the prioritization weight r_1 of class-1 is of order $1/K$.

The first scaling condition directly expresses that class-1 experiences a surge of traffic and that the number of class-1 customers is very large at time 0. Accelerating time together with re-scaling the first class (conditions 2 and 3) allow to 'zoom out' the process, just as for usual fluid limits and obtain a bird's-eye view of the large scale class-1 dynamics. Finally, condition 4 expresses that the priority weight of class-1 is very small so as to compensate for the large amount of traffic. This may describe several situations. For example, the surge of traffic may have been caused by an inappropriately small level of priority. Alternatively, if the surge is externally caused, perhaps due to a network attack, the network may be reacting to it by penalizing class-1 according to its level of congestion, so that other classes do not starve, see for instance [23] for practical considerations on the matter.

IV. Main result

A. The 1-dimensional case

To clarify the methodology, we begin by describing the fluid limit in the simplest case with only one class of traffic and constant service rate $\phi_1(\cdot)$ equal to φ. The process X is one-dimensional in this case. The following proposition is the usual fluid limit of an M/M/1, i.e., a functional law of large numbers (see [25]) and more generally of a $G/G/1$ queue with non-constant arrival rate (in that case we denote the arrival rate at time t with $\lambda(t)$). We will further attach a superscript ^{K} to all entities that correspond to the system that is initialized with the number of class-1 flows being of the order K. For example, the state vector is denoted by $X^K = (X_1^K, \ldots, X_N^K)$.

Proposition 4.1: Assume that:

$$\frac{1}{K}\int_0^{Kt} \lambda_1^K(s)ds \to a_1(t), \text{ as } K \to \infty.$$

Define $u^x(t)$ as the solution (assuming it is unique) of the differential equation:

$$\dot{u}^x(t) = (\dot{a}_1(t) - \varphi)\mathbf{1}_{\{u^x(t)>0\}},$$
$$u^x(0) = x.$$

Then, if $T^x = \inf\{t : u^x(t) = 0\}$, $\frac{X^K(Kt)}{K}$ converges in L^1 to $u(t)$, for all $t \leq T^x$, that is :

$$\forall t \leq T^x,\ E\left[\sup_{0\leq s\leq t}\left|\frac{X^K(Ks)}{K} - u(s)\right|\right] \to 0, K \to \infty.$$

B. The multidimensional case

We now consider a network with several classes of traffic and with class-1 going through a temporary surge of traffic. Recall that we focus on a regime where $r_1 \equiv \frac{1}{K}$ and $K \to \infty$. We further let Y^K denote the (scaled) process:

$$Y^K(t) = \left(\frac{X_1^K(Kt)}{K}, (X_i^K(Kt))_{i=2\ldots N}\right). \quad (5)$$

In the following we show that, as $K \to \infty$, Y^K converges to a stochastic process with a deterministic first coordinate, which is a solution of a differential equation which we describe in terms of an averaged rate $\bar{\phi}_1$. In the limit, the result implies a time scale separation between the first class and all other classes.

Define U^{z_1} to be an $N-1$ dimensional Markov birth-and-death process with arrival rates λ_i and death rates $\phi_i(z_1, \cdot)$, $i = 2 \ldots N$ ($z_1 \in \mathbb{R}^+$ being a fixed number here) and denote by $\pi^{z_1}(\cdot)$ its stationary probability. When we do not use a time index, we implicitly suppose that we consider stationary versions of the processes.

Let $u_1(t)$ be the solution (assuming it is unique) of the differential equation:

$$\dot{u}_1(t) = \dot{a}_1(t) - \bar{\phi}_1(u_1(t)),\ u_1(t) > 0, \qquad (6)$$
$$\dot{u}_1(t) = 0,\ u_1(t) = 0,$$

with $\bar{\phi}_1(z_1) = \sum_y \phi_1(z_1, y)\pi^{z_1}(y)$. To establish our main result, we shall make the following assumptions:

- (A_1) : ϕ is partially decreasing, i.e., $\phi_i(x)$ is decreasing in x_j, for $j \neq i$.
- (A_2) : $\phi_1(\cdot, x_2, \ldots, x_N)$ can be extended to a continuous function from $\mathbb{R}^+ \setminus \{0\}$ to $\mathbb{R}^+$.
- (A_3) : $\bar{\phi}(\cdot)$ can be extended to a Lipschitz-continuous function from $\mathbb{R}^+ \setminus \{0\}$ to $\mathbb{R}^+$.
- (A_4) : for all fixed z_1, the processes U^{z_1} are ergodic.
- (A_5) : $\frac{1}{K}\int_0^{Kt} \lambda_1^K(s)ds \to a_1(t)$.

We can now proceed to state our main result:

Theorem 4.2: Under the assumptions $(A_i)_{i=1\ldots5}$, the process $Y_1^K(t)$ converges in L^1, uniformly on compact intervals, to the deterministic trajectory $u_1(t)$, i.e.,

$$E[\sup_{0\leq s\leq t} |Y_1^K(s) - u_1(s)|] \to 0, \quad K \to \infty. \qquad (7)$$

Moreover, for all times t, and for all bounded Lipschitz-continuous functions f:

$$\lim_{s\to 0}\lim_{K\to\infty} E\left|\frac{1}{s}\int_t^{t+s} f(Y^K(h))dh - E[f(u_1(t), U^{u_1(t)})]\right| = 0. \qquad (8)$$

Remark 2: Note that the order of limits in (8) cannot be interchanged. Roughly speaking for a fixed s, the Ergodicity Theorem will make $\frac{1}{s}\int_t^{t+s} f(Y^K(h))dh = \frac{1}{Ks}\int_{Kt}^{K(t+s)} f(X^K(h))dh$ converge, while if s is too small, the first class will not have time to vary.

Proof: Let ϵ be given. Define the error estimate:

$$n_K(t) = \sup_{0\leq s\leq t} |Y_1^K(s) - u_1(s)|.$$

Using the classical martingale representation, for any function f:

$$f(X^K(t)) = f(X^K(0)) + M_K(t) + \int_0^t \Delta^K f(X^K(s))ds,$$

where M_K is a local martingale (which, in this case, is actually a martingale since the transitions are bounded), and where Δf is the function $\Delta^K f(x) = \sum_y q^K(x,y)(f(y) - f(x))$. Now, taking $f(x) = x_1$ and scaling time and space, we obtain:

$$\begin{aligned}\frac{X_1^K(Kt)}{K} &= x_1 + \frac{M_K(Kt)}{K} + \frac{1}{K}\int_0^{Kt} \lambda_1^K(s)ds \\ &\quad -\frac{1}{K}\int_0^{Kt} \phi_1(\frac{X_1^K(s)}{K}, \ldots, X_i^N(s))ds.\end{aligned}$$

For any martingale M, using Cauchy Schwartz and Doob's inequality [10], we get that:

$$\begin{aligned} E\left(|\sup_{0\leq s\leq t} M_s|\right)^2 &\leq E\left(\sup_{0\leq s\leq t} |M_s|\right)^2 \\ &\leq E\left(\sup_{0\leq s\leq t} M_s^2\right), \\ &\leq 4E\left(M_t^2\right). \end{aligned}$$

Since ϕ is bounded (independently of K), it follows that $E\left(M_K(t)^2\right) \leq At$ which implies that there exists a constant A' such that for K big enough:

$$E\left(|\sup_{0\leq s\leq t} \frac{M_K(Ks)}{K}|\right) \leq A'\sqrt{\frac{t}{K}} \leq \epsilon.$$

Define now the noise amplitude as:

$$\bar{M}_K(t) = \sup_{0\leq s\leq t} |\frac{M_K(Ks)}{K}|.$$

Using the convergence of the intensity of the arrival process,

$$\begin{aligned} &n_K(t) \leq \bar{M}_K(t) + \epsilon \\ &\quad + \sup_{s\leq t}\left|\frac{1}{K}\int_0^{Ks} \phi_1(\frac{X_1^K(z)}{K}, X_i^K(z))dz - \int_0^s \bar{\phi}_1(u(z))dz\right|. \end{aligned}$$

Let $s > 0$ be fixed. We decompose the interval $[0, s]$ into n sub-intervals $[t_i, t_{i+1}]$ of length $\delta_n = \frac{s}{n}$. We now write

$$\begin{aligned} &\frac{1}{K}\int_0^{Ks} \phi_1(\frac{X_1^K(z)}{K}, X_i^K(z))dz - \int_0^s \bar{\phi}_1(u(z))dz \\ &\quad = \sum_{i=1}^n \big(\frac{1}{K}\int_{Kt_i}^{K(t_i+\delta_n)} \phi_1(\frac{X_1^K(z)}{K}, X_i^K(z))dz \\ &\qquad - \int_{t_i}^{t_i+\delta_n} \bar{\phi}_1(u(z))dz\big). \end{aligned}$$

Consider an interval $[Kt_i, Kt_{i+1}]$ and write $\gamma_i = n_K(t_i + \delta_n)$. Define $x_i = \max_{s\in[t_i,t_{i+1}]} u(s)$. Using the definition of n_K, we have that for all $s \in [Kt_i, Kt_{i+1}]$,

$$\left|\frac{X^K(s)}{K} - x_i\right| \leq \gamma_i.$$

Using the monotonicity of the allocation function[2], there exists a coupling of the birth-and-death processes $U^{x+\gamma_i}$, $U^{x-\gamma_i}$ and $X_j^K,\ j \geq 2$ such that for $j \geq 2$ and $s \in [Kt_i, Kt_{i+1}]$:

$$U_j^{x+\gamma_i}(s) \leq X_j^K(s) \leq U_j^{x-\gamma_i}(s).$$

Using the previous inequalities:

$$\begin{aligned} &\frac{1}{K}\int_{Kt_i}^{K(t_i+\delta_n)} \phi_1(\frac{X_1^K(z)}{K}, X_2^K(z), \ldots, X_N^K(z))dz \leq \\ &\frac{1}{K}\int_{t_i}^{K(t_i+\delta_n)} \phi_1(x_i - \gamma_i, U_2^{x+\gamma_i}(z), \ldots, U_N^{x+\gamma_i}(z))dz, \end{aligned}$$

[2]For an exposition on stochastic comparisons for multidimensional birth-and-death processes, we refer to [6]

$$\frac{1}{K}\int_{Kt_i}^{K(t_i+\delta_n)} \phi_1(\frac{X_1^K(z)}{K}, X_2^K(z), \dots, X_N^K(z))dz \geq$$
$$\frac{1}{K}\int_{Kt_i}^{K(t_i+\delta_n)} \phi_1(x_i+\gamma_i, U_2^{x_i-\gamma_i}(z), \dots, U_N^{x_i-\gamma_i}(z))dz.$$

Now, applying the Ergodicity Theorem for the Markov processes $U^{x_i+\gamma_i}$ and $U^{x_i-\gamma_i}$ and using the Lipschitz continuity of $\bar{\phi}_1$, we get that there exists a function $g_{\delta_n,x_i,\gamma_i}(K)$ vanishing at infinity and a constant A such that:

$$\frac{\delta_n}{\delta_n K}\int_{Kt_i}^{K(t_i+\delta_n)} \phi_1(x+\gamma_i, U^{x-\gamma_i}(z))dz =$$
$$\delta_n(\bar{\phi}_1(x) + A\gamma_i + g_{n,x_i,\gamma_i}(K)).$$

Writing a similar equality for the upper bound we get that:

$$\left|\frac{1}{K}\int_{t_i}^{K(t_i+\delta_n)} \phi_1(\frac{X_1^K(z)}{K}, X_i^K(z))dz - \int_{t_i}^{t_i+\delta_n} \bar{\phi}_1(u(z))dz\right|$$
$$\leq 2A\delta_n\gamma_i + \delta_n g_{n,x_i,\gamma_i}(K).$$

There exists $K_i(n)$, such that $\forall K \geq K_i(n)$, $g_{n,x_i,\gamma_i} \leq \epsilon$. Hence, we obtain that for $K \geq \max_{i=1}^n K_i(n) = K(n)$:

$$\left|\frac{1}{K}\int_0^{Ks} \phi_1(\frac{X_1^K(z)}{K}, X_i^K(z))dz - \int_0^s \bar{\phi}_1(u(z))dz\right|$$
$$\leq \sum_{i=1}^n \delta_n(2An_K(t_i+\delta_n)+\epsilon).$$

Let $\eta_t^K = E[n_K(t)]$. Using that transitions rates are bounded, we have that $\eta_t^K \leq Bt$ and using further the Markov property, η_t^K is a continuous function of time and its modulus of continuity can be bounded independently of K. Hence, we can now choose n (independently of K) such that:

$$\left|\sum_{i=1}^n \delta_n \eta^K(t_i+\delta_n) - \int_0^t \eta_s^K ds\right| \leq \epsilon,$$

Then, for $K \geq K_n$,

$$E\left|\frac{1}{K}\int_0^{Ks} \phi_1(\frac{X_1^K(z)}{K}, X_i^K(z))dz - \int_0^s \bar{\phi}_1(u(z))dz\right|$$
$$\leq \sum_{i=1}^n \delta_n(2AE[n_K(t_i+\delta_n)]+\epsilon),$$
$$\leq 2A\int_0^t \eta_s^K ds + 3\epsilon.$$

Now using the bound on the supremum of the Martingale, we obtain: $\eta_t^K \leq 4\epsilon + 2A\int_0^t \eta_s^K ds$. Using Gronwall's lemma, we obtain that: $\eta_t^K \leq 4\epsilon e^{2At}$, which concludes the proof for the L^1 convergence of Y_1^K.

Note now that fixing t and replacing the function ϕ_1 by any Lipschitz-continuous function f, in the previous argument, we obtain that for n large enough and $K \geq K(n)$,

$$E\left|\frac{1}{K}\int_{Kt}^{Kt+K\delta_n} f(\frac{X_1^K(z)}{K}, X_i^K(z))dz - E[f(u_1(t), U^t)]\right|$$
$$\leq \delta_n 2A\eta_K(t) + \epsilon.$$

Hence for s small enough and $K \geq K(s)$:

$$E\left|\frac{1}{sK}\int_{Kt}^{Kt+Ks} f(\frac{X_1^K(z)}{K}, X_i^K(z))dz - E[f(u_1(t), U^t)]\right|$$
$$\leq \epsilon. \square$$

V. Qualitative behavior and numerical experiments

The theoretical results obtained in the previous section allow us to accurately describe the macroscopic behavior of the penalized unstable class, in conjunction with the finer dynamics of the other classes in the network. In this section, we again restrict to a fixed arrival rate $\lambda_1(t) \equiv \lambda_1$ for class-1. We shall compare the (deterministic) trajectories predicted by our theoretical results against simulation. Under the scaling considered in Theorem 4.2, we observe three types of qualitative behavior for the network responses:

1) The differential equation (6) governing the macroscopic dynamics of class-1 is unstable, in which case the traffic surge cannot be resolved and keeps building up. This may in turn lead to instability of other classes in the network as well.
2) The differential equation is asymptotically stable with a stable point $x_1 > 0$, in which case the network continues to see class-1 saturated (at a macroscopic state), for a relatively long (macroscopic) time period. The fixed point of the differential equation can be numerically evaluated by solving the equation:
$$\lambda_1 = \mu_1\bar{\phi}_1(z).$$
3) The differential equation is asymptotically stable with stable point 0, in which case the traffic surge will be resolved after a finite (macroscopic) time.

We now illustrate this trichotomy with a few examples.

A. Tree network

Let us consider the tree network shown in Figure 2 with $c_1 = 0.4$ and $c_2 = 0.8$. We shall assume the following bandwidth allocation : Define $\mathcal{S}_1 = \{(x_1,x_2) : (r_1x_1 + r_2x_2)c_1 < r_1x_1\}$. For $x_1 > 0$ and $x_2 > 0$,

$$\phi_1 = \begin{cases} c_1, & (x_1,x_2) \in \mathcal{S}_1, \\ \max\left(\frac{r_1x_1}{r_1x_1+r_2x_2}, 1-c_2\right), & (x_1,x_2) \in \mathcal{S}_1^c, \end{cases} \tag{9}$$

and $\phi_2 = 1 - \phi_1$.

For this network, the allocation becomes a strict priority allocation for class-2 when $r_1 = 0$, in which case class-1 gets capacity c_1 if there are no class-2 flows, and $1 - c_2$ otherwise. Thus, for a fixed value of ρ_2, class-1 is stable if $\rho_1 < \left(1 - \frac{\rho_2}{c_2}\right)c_1 + \frac{\rho_2}{c_2}(1-c_2)$. The stability regions for $r_1 = 0$ and $r_1 > 0$ are shown in Figure 4. Note that the stability region for the case $r_1 > 0$ (the entire shaded region) includes that of the case $r_1 = 0$ (darker shaded region only).

The dynamics of $u_1(t)$ for two different values of ρ_1 – one in each region – is plotted in figure 5, for which $\rho_2 = 0.5$.

For class 2, when the priority allocation is stable the dynamics of the average number of customers converges to

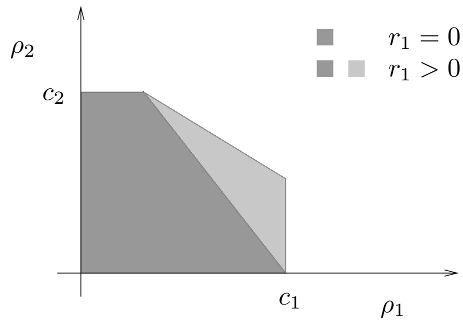

Fig. 4. Partitioning of the stability region for the tree network

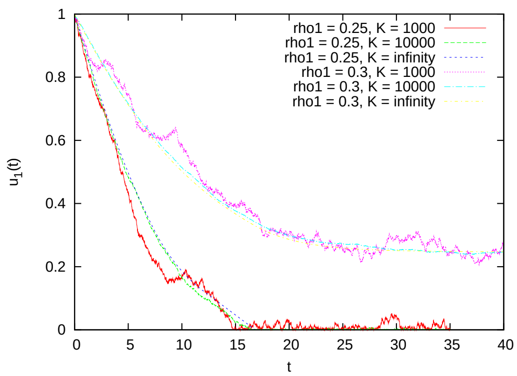

Fig. 5. Tree network : scaling of class-1

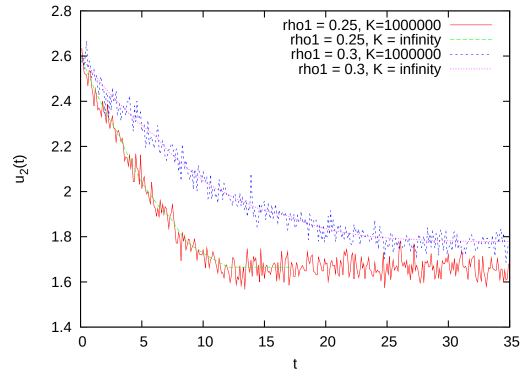

Fig. 6. Tree network: scaling of class-2

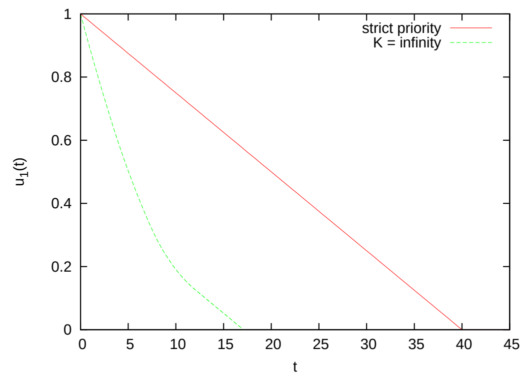

Fig. 7. Tree network: comparisons of trajectories of class-1 for a proportional fair allocation and a priority (to class-2) allocation

the one of priority, that is $\rho_2/(c_2 - \rho_2)$, as is illustrated in Figure 6.

In Figure 7, we show how class-1 is actually favored by asymptotically using the bandwidth of class-2, compared to the case where class-2 is given a strict priority.

B. Streaming and elastic traffic

The second example we consider is that of the integration of streaming and elastic traffic as described in Section III-B2. In this example, if accepted into the network, the capacity allocation to class-2 flows is cx_2 independently of the number of flows of class-1, for all values of $r_1 > 0$, whereas the capacity allocation of class-1 flows depends on the number of class-2 flows as follows

$$\phi_1(x_1, x_2) = r_1 x_1/(r_1 x_1 + c x_2).$$

However, class-2 flows are admitted only if there is sufficient capacity, that is, if $\frac{z_1}{z_1 + c(x_2+1)} + c(x_2+1) \leq 1$.

Let

$$\mathcal{S}_{z_1} = \{x_2 : \frac{z_1}{z_1 + cx_2} + cx_2 \leq 1\},$$

the state space of class-2 conditioned on $u_1(t) = z_1$. Define $\rho_2 = \frac{\lambda_2}{\mu_2 c}$. The process $U_2^{z_1}$ is a birth-death process with birth rate λ_2 and death rate $\mu_2 c x_2$, and whose stationary distribution is given by

$$\pi_2(x_2) = \frac{1}{\sum_{j \in \mathcal{S}_{z_1}} \rho_2^j/j!} \frac{\rho_2^{x_2}}{x_2!}.$$

For the priority allocation, class-1 is stable if and only if $\rho_1 < \pi_2(0)$, whereas for $r_1 > 0$, class-1 is stable if and only if $\rho_1 < 1 - \rho_2$. Thus, if $\rho_1 < \pi_2(0)$, then the limit point of $u_1(t)$ is 0, and if $\pi_2(0) < \rho_1 < 1$, then $u_1(t)$ tends to a positive limit point.

Note that under our scaling, for a fixed macroscopic state z_1, the state space depends both on z_1 and c. We can apply Theorem 4.2 with $\bar{\phi}$ being defined by:

$$\bar{\phi}(z) = \sum_{x_2 \in \mathcal{S}_{z_1}} \frac{z_1}{z_1 + cx_2} \frac{\rho^{x_2}}{x_2!} C(z_1),$$

where $C(z_1) = (\sum_{x_2 \in \mathcal{S}_{z_1}} \frac{\rho^{x_2}}{x_2!})^{-1}$. In the case that c is very small ($c << 1$), we might consider as a reasonable approximation a Poisson distribution for class-2, whatever the state of class-1. In that case, $\bar{\phi}$ takes a slightly simpler form. After simple calculations:

$$\bar{\phi}_1(cz_1) = H(z_1) = \frac{z_1 \int_0^{\rho_2} u^{z_1 - 1} \exp(u) du}{\rho_2^{z_1} \exp(\rho_2)}.$$

This allows a recursive evaluation for integer-valued z_1. Using simple calculus, for $n \in \mathbb{N}$:

$$H(n+1) = \frac{n+1}{\rho_2}(1 - H(n))$$

We can also evaluate H in terms of special functions:

$$H(n) = \frac{n! - n\Gamma(n, -1)}{(-\rho_2)^n \exp(\rho_2)},$$

where $\Gamma(n, -1)$ is the incomplete Γ function.

We can further argue, using the Poisson approximation, that when c is small the distribution of class-2 customers would be concentrated around its mean when $u_1(t)$ is close enough to its limit point. Thus, from the point of view of class-1, class-2 flows will appear to be take fixed capacity of $cx_2 = \rho_2$. We can compute the limit point using the equation

$$0 = \left(\lambda_1 - \frac{u}{u + \rho_2}\right),$$

from which we obtain the limit point to be $\frac{\rho_1}{1-\rho_1}\rho_2$. In Figure 8, we plot the $u_1(t)$ for $\rho_1 = 0.6$, $\rho_2 = 0.2$, and $c = 0.01$. Toward the limit point, u_1 gets close to 0.3 which is the value calculated using the approximation for small values of c.

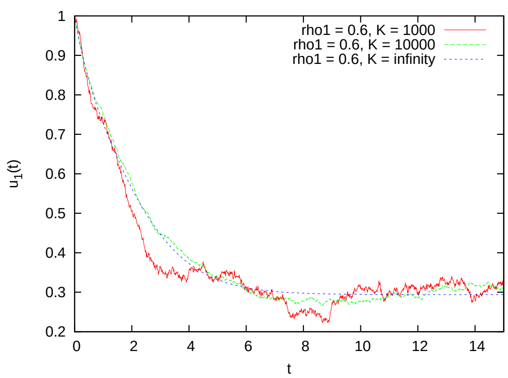

Fig. 8. Streaming and elastic traffic : scaling of the elastic traffic

VI. Conclusions

We analyzed the flow-level performance of multi-class communication networks when one of the classes undergoes a traffic surge. We showed that, under an appropriate scaling of space and time, the dynamics of the temporarily unstable class can be described by a deterministic differential equation in which the time derivative at a given point depends on the conditional stationary distribution of the other classes calculated at that point. We illustrated the behavior through several examples of network topologies and bandwidth allocations that are commonly used to model communication networks.

The time-space-transitions scaling that we considered raises several open questions which would give a better understanding of the network dynamics. We mention a few of these challenges. Since usual fluid limits are a powerful tool to study stability, one may wonder how the stochastic stability relates to the properties of the limiting processes of this type of scaling. In particular, finding necessary and sufficient conditions for the limit point of non work-conserving allocations to be zero would be extremely insightful. Finally, our work would benefit when equipped with error bound estimates, which are necessary for any reliable performance evaluation tool. Finally, as suggested by a reviewer, determining the diffusion process obtained by subtracting the fluid limit from the original process seems within reach.

References

[1] http://www.akamai.com/html/technology/nui/ news/index.html.

[2] http://en.wikipedia.org/wiki/Web_traffic.

[3] T. Bonald, L. Massoulié, A. Proutière, and J. Virtamo. A queueing analysis of max-min fairness, proportional fairness and balanced fairness. *Queueing Syst. Theory Appl.*, 53(1-2):65–84, 2006.

[4] T. Bonald and A. Proutière. Insensitive bandwidth sharing in data networks. *Queueing Syst. Theory Appl.*, 44(1):69–100, 2003.

[5] T. Bonald and A. Proutière. On performance bounds for the integration of elastic and adaptive streaming flows. In *SIGMETRICS*, pages 235–245, 2004.

[6] S. Borst, M. Jonckheere, and L. Leskelä. Stability of parallel queueing systems with coupled service rates. *Discrete Event Dyn. Syst.*, 18(4):447–472, 2008.

[7] S. Cheung, R. Boucherie, and R. Nunez-Queija. Quasi-stationary analysis for queues with temporary overload. In *ITC 22*, this volume.

[8] G. L. Choudhury, A. Mandelbaum, M. I. Reiman, and W. Whitt. Fluid and diffusion limits for queues in slowly changing environments. *Stoch. Mod*, 13:121–146, 1997.

[9] J. G. Dai. On positive harris recurrence of multiclass queueing networks: A unified approach via fluid limit models. *Annals of Applied Probability*, 5:49–77, 1995.

[10] R. W. R. Darling and J. R. Norris. Differential equation approximations for markov chains. *Probability Surveys*, 5:37, 2008.

[11] F. Delcoigne, A. Proutière, and G. Régnié. Modeling integration of streaming and data traffic. *Perform. Eval.*, 55(3-4):185–209, 2004.

[12] M. Deshpande, A. Amit, M. Chang, N. Venkatasubramanian, and S. Mehrotra. Flashback: A peer-to-peer web server for flash crowds. *Distributed Computing Systems, International Conference on*, page 15, 2007.

[13] N. Duffield and W. Whitt. Control and recovery from rare congestion events in a large multi-server system. *Queueing Syst. Theory Appl.*, 26:69–104, 1997.

[14] H. C. Gromoll and R. J. Williams. Fluid limits for networks with bandwidth sharing and general document size distributions. *Annals of Applied Probability*, 19:243, 2009.

[15] A. Hordijk and N. M. van Dijk. Stationary probabilities for networks of queues. *Applied Probability-Computer Science: The Interface*, 2:423–451, 1982.

[16] M. Izal, G. Urvoy-Keller, E. W. Biersack, P. Felber, A. A. Hamra, and L. Garcés-Erice. Dissecting BitTorrent: Five months in a Torrent's lifetime. In *PAM*, pages 1–11, April 2004.

[17] S. Kandula, D. Katabi, M. Jacob, and A. W. Berger. Botz-4-Sale: Surviving Organized DDoS Attacks That Mimic Flash Crowds. In *2nd Symposium on Networked Systems Design and Implementation (NSDI)*, Boston, MA, May 2005.

[18] F. Kelly. *Reversibility and Stochastic Networks*. Wiley, 1979.

[19] F. P. Kelly, A. K. Maulloo, and D. K. H. Tan. Rate control for communication networks: Shadow prices, proportional fairness and stability. *The Journal of the Operational Research Society*, 49(3):237–252, 1998.

[20] C. Labovitz, D. McPherson, M. Hollyman, and S. Iekel-Johnson. Internet traffic trends - a view from 67 ISPs, June 2008.

[21] L. Massoulié and J. Roberts. Bandwidth sharing: objectives and algorithms. *IEEE/ACM Trans. Netw.*, 10(3):320–328, 2002.

[22] R. Nunez-Queija, J. van den Berg, and M. Mandjes. Performance evaluation of strategies for integration of elastic and stream traffic. In *ITC 16*, 1999.

[23] T. Peng, C. Lecki, and K. Ramamohanarao. Survey of network-based defense mechanisms countering the DoS and DDoS problems. *ACM Comput. Surv.*, 39(1), 2007.

[24] O. Perry and W. Whitt. Responding to unexpected overloads in large-scale service systems. *Management Science*, 55(8):1353–1367, 2009.

[25] P. Robert. *Stochastic Networks and Queues*. Springer, 2003.

[26] L. A. Segel and M. Slemrod. The quasi-steady-state assumption: A case study in perturbation. *SIAM Review*, 31(3):446–477, 1989.

[27] R. Serfozo. *Introduction to Stochastic Networks*. Springer, 1999.

[28] S.Meyn. *Control techniques for complex networks*. Cambridge University Press, 2008.

[29] A. Stavrou, D. Rubenstein, and S. Sahu. A lightweight, robust p2p system to handle flash crowds. *Selected Areas in Communications, IEEE Journal on*, 22(1):6–17, 2004.

[30] G. Yin and Q. Zhang. *Continuous-time Markov Chains and Applications: A Singular Perturbation Approach*. Springer, 1998.