

Fair Internet traffic integration: network flow models and analysis

Peter Key, Laurent Massoulié

Microsoft Research, 7 J J Thompson Avenue, Cambridge CB3 0FB, UK

Alan Bain, Frank Kelly

Statistical Laboratory, University of Cambridge, Cambridge CB3 0WB, UK

Abstract

We use flow-level models to study the integration of two types of Internet traffic, elastic file transfers and streaming traffic. Previous studies have concentrated on just one type of traffic, such as the flow level models of Internet congestion control, where network capacity is dynamically shared between elastic file transfers, with a randomly varying number of such flows. We consider the addition of streaming traffic in two cases, under a fairness assumption that includes TCP-friendliness as a special case, and under certain admission control schemes. We establish sufficient conditions for stability, using a fluid model of the system. We also assess the impact of each traffic type on the other: file transfers are seen by streaming traffic as reducing the available capacity, whereas for file transfers the presence of streaming traffic amounts to replacing sharp capacity constraints by relaxed constraints. Simulation results suggest that the integration of streaming traffic and file transfers has a stabilizing effect on the variability of the number of flows present in the system.

Keywords: Internet, Quality of Service, Fairness, Elastic Traffic, Streaming Traffic, Fluid Models, Flow Level

Intégration équitable du trafic dans l'Internet: modèles fluides de flots et leur analyse.

Résumé

Nous proposons des modèles de flots représentant l'intégration de deux types de trafic Internet: les transferts de données, ou trafic élastique, et le trafic temps-réel. Les travaux antérieurs ont principalement traité un seul type de trafic, comme les modèles de flots du contrôle de congestion dans l'Internet, où la capacité du réseau est partagée dynamiquement entre les transferts de données en cours, dont le nombre évolue dans le temps. Nous considérons deux scénarios d'intégration, l'un reposant sur une hypothèse d'équité générale, dont un cas particulier est la compatibilité avec TCP (ou "TCP-friendliness"), et l'autre reposant sur une politique de contrôle d'admission des flots temps-réel. Nous considérons une renormalisation des processus décrivant l'état du réseau. Nous donnons des conditions suffisantes de stabilité pour ces processus renormalisés. Nous évaluons aussi l'impact qu'a chaque type de trafic sur l'autre: le trafic élastique a pour seul effet de réduire la capacité offerte au trafic temps-réel, alors que l'effet du trafic temps-réel sur le trafic élastique est de changer des contraintes strictes de capacité en des contraintes pénalisées. Des résultats de simulation suggèrent que l'intégration des deux types de trafic réduit la variabilité du nombre de flots présents dans le système.

Mots clés: Internet, Qualité de Service, Equité, Trafic élastique, Trafic temps-réel, Modèles fluides, modèles de flots.

I INTRODUCTION

The motivation for this paper arises from the need to understand and model the integration of different types of traffic within the Internet. At the transport layer, the current Internet is dominated by flows which use TCP. The percentage of TCP traffic is variable, and may depend on time of day and the particular route chosen; however typical measurements on a backbone [17] show that upwards of 70% of flows use TCP, rising to over 90% by volume, with UDP the main alternative protocol (up to 20% of packets, or 10% of bytes). Prevailing applications can change rapidly: whereas Web traffic used to be the dominant application type for TCP traffic, at the time of writing file-sharing applications can dominate and may account for 40% of the traffic on backbone links. The current volume of streaming traffic carried by UDP is small (less than 10%), but the rapid increase in peer-to-peer traffic illustrates how quickly the status-quo can change, and we would like to predict behaviour in different scenarios.

How TCP and UDP should co-exist is a vexed question, and many regard UDP-related traffic as inherently problematic. Some authors have proposed that streaming traffic should be TCP-friendly, so that it can share network resources fairly with the dominant form of existing traffic [11]. Indeed some streaming applications use TCP as the transport protocol. Applications that use UDP often need some form of quality of service to function adequately, which has led some researchers to consider distributed or end-point admission control [14, 6, 3, 16].

The need to model such situations requires modelling the heterogeneous traffic streams, with their different characteristics. Previous work in this area has focused on analysing occupancy distributions of single resource systems, via either exact or approximate techniques, see e.g. [2],[1],[21] and [23]. In contrast we look at arbitrary network topologies. We consider two types of traffic, which we label ‘file transfers’ and ‘streaming’ traffic. A flow carrying a file transfer must transfer a given volume: the volume may be random, but is independent of the level of congestion experienced. An admitted flow carrying streaming traffic remains present for a holding time: the holding time may be random, but is independent of the level of congestion experienced.

Our analysis is a flow-level analysis, that generalises the flow level model of Internet congestion control of Massoulié and Roberts [18] by incorporating streaming traffic. They considered a network where a randomly varying number of flows is present, and capacity is dynamically shared between elastic file transfers using different sharing mechanisms.

The analysis of streaming traffic on its own gives rise to a product-form solution under certain reasonable assumptions, a form which is preserved under certain types of call admission control [14]. Moreover the limiting behaviour as the size of the system grows leads naturally to a non-degenerate limit for the (scaled) number of connections. In contrast, a similar scaling applied to just file transfer traffic results in numbers of competing flows either increasing to infinity or decreasing to zero; it has been suggested [8] that such a model is flawed, lacking any self-limiting behaviour. We shall see that this criticism is avoided when the two types of traffic are mixed, and that the presence of even a small amount of streaming traffic has a stabilising effect.

The organization of this paper is as follows. In Section 2 we describe the sharing policy between flows where we assume a generalized form of TCP-friendliness for the streaming traffic. The generalisation is based on the so-called α -fair allocations [20]. In Section 3 we describe the flow level stochastic model of a network, a generalization of [18]. File transfers are characterized by a random Poisson arrival process, with exponentially distributed file sizes, whereas streaming traffic has Poisson arrival rates but an exponentially distributed holding time. In Section 4 we establish appropriate stability conditions, for a fluid model of the system, through the construction of an appropriate Lyapunov function. We also characterize and interpret the network state in equilibrium. In Section 5 we consider extensions where we relax the sharing assumptions between the two types of traffic. In particular, we discuss admission control strategies for the case where the streaming traffic enters at a fixed rate, and show how a particular admission control strategy previously considered in [13] can be considered TCP-fair. In Section 6 we discuss simulations of the flow level model for a star network, and explore the impact of streaming traffic on the variability of flow allocations. We conclude in Section 7.

II FAIRNESS ASSUMPTIONS

Consider a network with resources labelled by $j \in J$. Let a route r identify a non-empty subset of J (interpreted as the set of resources used by a flow on route r). Write R for the set of possible routes. Set $A_{jr} = 1$ if resource j lies on route r (i.e. $j \in r$), and set $A_{jr} = 0$ otherwise. We

assume positive finite capacities $(C_j, j \in J)$.

Let N_r be the number of flows on route r . Given a fixed parameter $\alpha \in (0, \infty)$ and strictly positive weights $(w_r, r \in R)$, we suppose that the bandwidth allocation to each of the N_r flows on route r is x_r , where $x = (x_r, r \in R)$ is a solution to the following optimization problem:

$$\text{maximize} \quad \sum_{r \in R} w_r N_r \frac{x_r^{1-\alpha}}{1-\alpha} \quad (1)$$

$$\text{subject to} \quad \sum_r A_{jr} N_r x_r \leq C_j, \quad j \in J \quad (2)$$

$$\text{over} \quad x_r \geq 0, \quad r \in R. \quad (3)$$

Call the resulting allocation a weighted α -fair allocation [20].

The form of a solution to the problem (1–3) can be given in terms of Lagrange multipliers $(p_j, j \in J)$, one for each of the capacity constraints (2), as

$$x_r = \left(\frac{w_r}{\sum_j p_j A_{jr}} \right)^{1/\alpha}, \quad (4)$$

where

$$p_j \geq 0, \quad p_j \left(C_j - \sum_r A_{jr} N_r x_r \right) = 0 \quad j \in J. \quad (5)$$

The strict concavity of the objective function (1) as a function of $(x_r, r : N_r > 0)$ ensures that the component x_r is unique if N_r is positive. When $w_r = 1, r \in R$, the cases $\alpha \rightarrow 0$, $\alpha \rightarrow 1$ and $\alpha \rightarrow \infty$ correspond respectively to an allocation which achieves maximum throughput, is *proportionally fair* or is *max-min fair* [5, 20]. Weighted α -fair allocations provide a tractable theoretical abstraction of decentralized packet-based congestion control algorithms such as TCP.

If $\alpha = 2$ and w_r is the reciprocal of the square of the round trip time on route r , then the formula (4) is a version of the *inverse square root law* familiar from studies of the throughput of TCP connections [10, 19, 22]. A flow carrying streaming traffic is termed *TCP-friendly* if, *inter alia*, it adapts its rate to correspond with the steady-state rate of a TCP connection, usually characterized in terms of a version of the inverse square root law [11].

The relations (2–5), and more refined versions of these relations, can be solved to give predictions of throughput, given the numbers of flows N present [7, 12, 25]. Given N , network performance along different routes can be predicted. But what determines the behaviour of N ? One aim of this paper is to better understand how the behaviour of N is influenced by the mix of traffic types present.

III FLOW LEVEL STOCHASTIC MODEL

We now describe our model of how flows arrive and depart. Our aim is to generalize the stochastic model for file transfers introduced in [18] to include streaming flows.

Let N_r be the number of document transfers on route r , and let M_r be the number of streaming flows on route r . Define the indicator function $I[r = s] = 1$ if $r = s$, $I[r = s] = 0$ otherwise. Let $T_s N = (N_r + I[r = s], r \in R)$, with inverse $T_s^{-1} N = (N_r - I[r = s], r \in R)$. We suppose that $(N, M) = (N_r, r \in R; M_r, r \in R)$ is a Markov process, with state space $\mathbb{Z}_+^J \times \mathbb{Z}_+^J$ and non-trivial transition rates

$$q((N, M), (T_r N, M)) = \nu_r, \quad q((N, M), (T_r^{-1} N, M)) = \mu_r N_r x_r(N + M), \quad r \in R$$

$$q((N, M), (N, T_r M)) = \kappa_r, \quad q((N, M), (N, T_r^{-1} M)) = M_r \eta_r, \quad r \in R$$

for $(N, M) \in \mathbb{Z}_+^J \times \mathbb{Z}_+^J$, where $x(N)$ is a solution to the optimization problem (1–3). This corresponds to a model where new file transfers arrive on route r as a Poisson process of rate ν_r , new streaming flows arrive on route r as a Poisson process of rate κ_r , and $x_r(N + M)$ is the bandwidth allocated to each flow on route r , whether it is a file transfer or streaming flow. A file transfer on route r transfers a file whose size is exponentially distributed with parameter μ_r , and a streaming flow on route r has an exponentially distributed holding time with parameter η_r .

If $\kappa_r = 0, r \in R$, then this model reduces to the model introduced by Massoulié and Roberts [18], in which there are no streaming flows, only file transfers. For this case, De Veciana, Lee and Konstantopoulos [9] and Bonald and Massoulié [5] have shown that a sufficient condition for the Markov chain $(N(t), t \geq 0)$ to be positive recurrent is that

$$\sum_r A_{jr} \rho_r < C_j, \quad j \in J, \quad (6)$$

where $\rho_r = \nu_r / \mu_r$; this condition is also necessary [15]. The condition is natural: ρ_r is the load on route r , and we can identify the ratio of the two sides of the inequality (6) as T the *traffic intensity* at resource j . Kelly and Williams [15] have explored the behaviour of a fluid model for this case in heavy traffic, when the inequalities (6) are close to being tight, which is a key step towards proving state space collapse. The papers [5, 9, 15] all make use of a fluid model of the Markov process, an approach which we shall use for our analysis of the extended model.

We shall henceforth assume that $\kappa_r > 0, r \in R$, and that condition (6) is satisfied. Define the *reduced capacities*

$$\tilde{C}_j = C_j - \sum_r A_{jr} \rho_r, \quad j \in J. \quad (7)$$

Thus the reduced capacity \tilde{C}_j on resource j is just the amount by which inequality (6) fails to be tight. The reduced capacities will determine the capacity available to streaming flows in a sense that will be made precise in the next section.

IV STABILITY OF FLUID MODELS

Next we describe a fluid model, which can be thought of as a formal law of large numbers approximation under the scaling

$$(n, m)(t) = \left(\frac{N_c(t)}{c}, \frac{M_c(t)}{c} \right) \quad c \rightarrow \infty,$$

where $(N_c(t), M_c(t))$ is the model of the previous Section but with $C_j, j \in J$, and $\nu_r, \kappa_r, r \in R$, replaced by $cC_j, j \in J$, and $c\nu_r, c\kappa_r, r \in R$, respectively. The fluid model is an approximation appropriate for the case where $C_j, j \in J$, and $\nu_r, \kappa_r, r \in R$, are all large, an important case in applications.

The fluid model for the Markov process of the last Section takes the form

$$\frac{d}{dt} n_r(t) = \nu_r - \mu_r n_r(t) x_r (n(t) + m(t)), \quad r \in R \quad (8)$$

$$\frac{d}{dt} m_r(t) = \kappa_r - \eta_r m_r(t), \quad r \in R. \quad (9)$$

Note that our assumption that $\kappa_r > 0, r \in R$, implies that $m_r(t) > 0, r \in R, t > 0$.

Proposition 1. *Provided the condition (6) is satisfied, the differential equations (8,9) have a unique invariant point, (\hat{n}_r, \hat{m}_r) . It takes the form $\hat{m}_r = \kappa_r / \eta_r$ and*

$$\hat{n}_r = \frac{\nu_r}{\mu_r} \left(\frac{\sum_{j \in J} p_j A_{jr}}{w_r} \right)^{1/\alpha} \quad r \in R, \quad (10)$$

for some $p \in \mathbb{R}_+^J$. At the invariant point the bandwidth allocation to each flow on route r is

$$x_r = \left(\frac{w_r}{\sum_j p_j A_{jr}} \right)^{1/\alpha}. \quad (11)$$

The pair (x, p) forms a solution of equation (11) and the conditions

$$p_j \geq 0, \quad p_j \left(\tilde{C}_j - \sum_r A_{jr} \hat{m}_r x_r \right) = 0 \quad j \in J, \quad (12)$$

and together these relations determine x uniquely.

Proof. At an invariant point $m_r(t) = \hat{m}_r$, from equation (9). Further,

$$\hat{n}_r x_r (\hat{n} + \hat{m}) = \rho_r, \quad (13)$$

from equation (8). Now at any time t ,

$$x_r(n(t) + m(t)) = \left(\frac{w_r}{\sum_j p_j(t) A_{jr}} \right)^{1/\alpha}$$

where

$$p_j(t) \geq 0, \quad p_j(t) \left(C_j - \sum_r A_{jr} (n_r(t) + m_r(t)) x_r(n(t) + m(t)) \right) = 0 \quad j \in J,$$

from the characterization of x as a solution to an optimization problem of the form (1–3). Thus, at an invariant point,

$$p_j \geq 0, \quad p_j \left(\tilde{C}_j - \sum_r A_{jr} \hat{n}_r \left(\frac{w_r}{\sum_j p_j A_{jr}} \right)^{1/\alpha} \right) = 0 \quad j \in J,$$

using equation (13) and the definition (7). Thus x , given by (11), is the unique optimum to a problem of the form (1–3), with C replaced by \tilde{C} and N replaced by \hat{n} . \square

Equation (10) describes the vector \hat{n} , of dimension $|R|$, in terms of p , a vector which may have a much smaller dimension, $|J|$, a phenomenon first noted in the balanced fluid model of [15].

Remark 2. *The invariant point can be interpreted as follows. File transfers place an irreducible load $\sum_r A_{jr} \rho_r$ on resource j for each $j \in J$. The reduced capacities $(\tilde{C}_j, j \in J)$ that remain after this load is satisfied are available to be shared amongst streaming traffic, and determine the bandwidth allocation to flows on route r for both types of traffic.*

When $\kappa_r = 0, r \in R$, the unique invariant point of the fluid model is $\hat{n} = 0$ [9, 5]. It is notable that the inclusion of streaming traffic within the fluid model forces the components of \hat{n} to be positive.

We now discuss convergence to the equilibrium point of the above dynamics. In order to do so, it is convenient to introduce a modification for the dynamics of file transfers. This is naturally described in terms of the quantities λ_r , which represent the total capacity allocated to type r file transfers, and thus with the previous notation, $\lambda_r = n_r x_r$. Let the function $\psi(\lambda)$ be a penalty function. Then the modified dynamics are as follows:

$$\frac{d}{dt} n_r(t) = \nu_r - \mu_r \lambda_r(n(t)), \quad r \in R, \quad (14)$$

where the vector λ of service rates λ_r is defined as the solution to the optimisation problem

$$\text{maximize} \quad \phi(\lambda) := \sum_{r \in R} w_r n_r^\alpha \frac{\lambda_r^{1-\alpha}}{1-\alpha} + \psi(\lambda) \quad (15)$$

$$\text{subject to} \quad \sum_r A_{jr} \lambda_r \leq C_j, \quad j \in J \quad (16)$$

$$\text{over} \quad \lambda_r \geq 0, \quad r \in R. \quad (17)$$

In the case where ψ is identically zero, this reduces to the previous dynamics for the file transfers in the absence of streaming traffic. The function ψ is assumed to be concave and strictly monotonic decreasing in each coordinate on the domain of the optimisation problem. This latter condition implies that the rate λ_r goes to zero as n_r goes to zero, and hence the trajectories n_r stay away from the boundary of the orthant \mathbb{R}_+^R . Let us prove stability of the above dynamics.

Theorem 3. *Under the stability conditions (6), the function $L(n)$ defined by*

$$L(n) = \sum_r \frac{1}{\mu_r} \left\{ w_r \frac{n_r^{1+\alpha}}{(1+\alpha)\rho_r^\alpha} + n_r \psi'_r(\rho) \right\}, \quad (18)$$

where $\psi'_r(\rho)$ stands for the r -th partial derivative $\frac{\partial \psi}{\partial \lambda_r}$ evaluated at the vector of loads ρ_r , is a Lyapunov function for the dynamics (14–17). Hence these dynamics converge to the unique minimiser of L on the orthant \mathbb{R}_+^R , that is

$$\hat{n}_r = \rho_r \left(\frac{-\psi'_r(\rho)}{w_r} \right)^{1/\alpha}. \quad (19)$$

Proof. Under the condition (6), the vector $\rho = (\rho_r, r \in R)$ lies in the interior of the domain (16–17) of the optimisation problem defining the vector λ . The function ϕ is strictly concave on this domain, since both terms in its definition (15) are concave, with strict concavity of the first term. Hence

$$\sum_r \phi'_r(\rho)(\rho_r - \lambda_r) \leq 0,$$

and this inequality is strict unless $\lambda = \rho$. The left-hand side also reads

$$\sum_r \left\{ w_r \left(\frac{n_r}{\rho_r} \right)^\alpha + \psi'_r(\rho) \right\} (\rho_r - \lambda_r),$$

and is thus equal to

$$\sum_r \frac{\partial L}{\partial n_r}(n(t)) \frac{d}{dt} n_r(t) = \frac{d}{dt} L(n(t)).$$

Thus the value of $L(n)$ decreases strictly along the trajectories of the system, except at the equilibrium point specified by (19), which is the only point for which the corresponding rate vector λ equals the load vector ρ . \square

Remark 4. *If the concave function ψ fails to be differentiable at ρ , by adapting the above proof it can be shown that the dynamics (14–17) converge to the set of points \hat{n} satisfying (19), where the vector $(-\psi'_r(\rho), r \in R)$ spans the set of sub-gradients of the convex function $-\psi$ at ρ . We refer the reader to [24], p.214 for a definition and basic properties of sub-gradients of convex functions.*

We now apply this result to establish stability of the dynamics (8–9).

Corollary 5. *Under the stability condition (6), the dynamics (8–9) are asymptotically stable.*

Proof. We shall only treat the special case where the m_r have already converged to their equilibrium values, \hat{m}_r . As the convergence of $m(t)$ to \hat{m} does not depend on the evolution of $n(t)$, the general case can be deduced by continuity arguments. We now show that the n_r evolve according to (14–17) for some suitable choice of a penalty function ψ . Indeed, (14) holds, with the service rates λ_r solving

$$\begin{aligned} & \text{maximize} && \phi(\lambda, \gamma) := \sum_{r \in R} w_r \left\{ n_r^\alpha \frac{\lambda_r^{1-\alpha}}{1-\alpha} + \hat{m}_r^\alpha \frac{\gamma_r^{1-\alpha}}{1-\alpha} \right\} \\ & \text{subject to} && \sum_r A_{jr}(\lambda_r + \gamma_r) \leq C_j, \quad j \in J \\ & \text{over} && \lambda_r, \gamma_r \geq 0, \quad r \in R. \end{aligned}$$

Performing the optimisation over the γ_r first, this is of the form (15–17), with

$$\begin{aligned} \psi(\lambda) & := \sup \left\{ \sum_r w_r \hat{m}_r^\alpha \frac{\gamma_r^{1-\alpha}}{1-\alpha} \right\}, \\ \text{over } \gamma \in S(\lambda) & := \left\{ \gamma \in \mathbb{R}_+^R, \sum_r A_{jr} \gamma_r \leq C_j - \sum_r A_{jr} \lambda_r, j \in J \right\}. \end{aligned} \tag{20}$$

It is readily seen that ψ is decreasing in each coordinate: given λ, λ' , such that $\lambda'_r \leq \lambda_r$ for all r , the inequality being strict for some r , any vector γ in $S(\lambda)$ is such that $\gamma' := (\gamma_r + \lambda_r - \lambda'_r)$ is in $S(\lambda')$, so that $\psi(\lambda) < \psi(\lambda')$. Concavity of ψ also holds: given λ, λ' and ϵ in $[0, 1]$, denote by γ and γ' the maximising vectors in the definition of $\psi(\lambda), \psi(\lambda')$ respectively. Then $\epsilon\gamma + (1-\epsilon)\gamma'$ lies in $S(\epsilon\lambda + (1-\epsilon)\lambda')$, and hence

$$\psi(\epsilon\lambda + (1-\epsilon)\lambda') \geq \sum_r w_r \hat{m}_r^\alpha \frac{(\epsilon\gamma_r + (1-\epsilon)\gamma'_r)^{1-\alpha}}{1-\alpha} \geq \epsilon\psi(\lambda) + (1-\epsilon)\psi(\lambda'),$$

where concavity of the function maximised in the definition of ψ gives the second inequality. \square

Remark 6. Under the particular choice (20) of penalty function, and comparing equations (10) and (19), we deduce that $\sum_{j \in J} p_j A_{jr} = -\psi'_r(\rho)$. Notice the identification between the sensitivity of the penalty function ψ with respect to the load ρ_r and the sum of the Lagrange multipliers along route r .

V EXTENSIONS: PACKET MODELS AND ADMISSION CONTROL

V.1 Constraint relaxation

The formulation (14–17) is also useful to model situations where the hard capacity constraints described by the intersection of half-spaces (2) are relaxed. If the optimization problem (1–3) is

replaced by

$$\begin{aligned} & \text{maximize} && \sum_r w_r N_r \frac{x_r^{1-\alpha}}{1-\alpha} - \sum_j C_j \left(\sum_r A_{jr} N_r x_r \right) \\ & \text{over} && x \geq 0, \end{aligned}$$

where $C_j(\cdot), j \in J$, are convex, strictly increasing, differentiable functions, then an optimum is again given by equation (4), but where now $p_j, j \in J$, satisfy

$$p_j = C'_j \left(\sum_r A_{jr} N_r x_r \right).$$

This formulation arises naturally from packet level models, with x_r the mean rate of a stochastic packet generation process. For example, if the resources j correspond to output ports of routers, then there is a limited amount of buffering available, and packets will be dropped if the capacity is exceeded, or more generally marked according to some active queue management technique. We may interpret $p_j(y_j)$ as the probability of dropping (or marking) a packet at resource j when the load on the resource is y_j .

Stability of the corresponding fluid model can be deduced from the formulation (14–17), by setting

$$\psi(\lambda) = - \sum_j C_j \left(\sum_r A_{jr} \lambda_r \right).$$

V.2 Admission controlled traffic

Streaming may need some minimal non-zero rate for the application to function adequately. For example in the case of streaming multimedia, even with adaptive codecs, some minimal transmission rate is often required for acceptable performance. Suppose that type r streaming traffic only enters if $x_r \geq x_r^{\min}$: then in both the flow level stochastic model and in the fluid limit, κ_r is replaced by $\kappa_r I[x_r \geq x_r^{\min}]$. At an invariant point, either $m_r > 0$ and $x_r \geq x_r^{\min}$ or $m_r = 0$. The condition $x_r \geq x_r^{\min}$ is equivalent to

$$\sum_j p_j A_{jr} \leq \frac{w_r}{(x_r^{\min})^\alpha} \quad r \in R. \quad (21)$$

If the parameters $p_j, j \in J$, satisfy the linear constraints (21) with strict inequality, then the fluid model predicts there will be no call admission blocking.

A more extreme case is when real-time streaming traffic cannot adapt its rate at all. We now describe a sharing model relevant for such a scenario, according to which streaming flows either proceed at their target rate, or are rejected, in such a way that the equilibrium points are the same as for the previous model (at least in a situation of interest).

More specifically, type r streaming flows have a target rate π_r . At a given time t , they measure the current TCP-friendly rate they would get under the previous sharing model, say λ_r ; they then proceed at full target rate π_r with probability $\min(1, \lambda_r/\pi_r)$, and are rejected with the complementary probability. Once started, they no longer adapt to the network state. Such an approach has been proposed by Karlsson [13].

Keeping the same notations as in the previous sections, the fluid equations describing the evolution of the numbers of flows are now

$$\frac{d}{dt} n_r = \nu_r - \mu_r n_r x_r(n; C - A\pi m), \quad r \in R \quad (22)$$

$$\frac{d}{dt} m_r = \kappa_r \min\left(1, \frac{x_r(n; C - A\pi m)}{\pi_r}\right) - \eta_r m_r, \quad r \in R. \quad (23)$$

In the above we denote by $x_r(n, C)$ the solution to the optimisation problem (1–3), where we have made explicit both the numbers of flows n and the capacity constraints C . Note that the capacity allocations are now defined based on the numbers of file transfers n , and the reduced capacities $C - A\pi m$, where A is the link-flow incidence matrix (as before), π is the diagonal matrix with diagonal entries π_r , and $m = (m_r)$. We now characterize the equilibrium points under these dynamics.

Proposition 7. *Provided the condition (6) is satisfied, any invariant point (\hat{n}_r, \hat{m}_r) of the differential equations (22,23) takes the form $\hat{m}_r = \kappa_r/\eta_r \min(1, x_r/\pi_r)$, $\hat{n}_r = (\nu_r/\mu_r)x_r^{-1}$, where x_r satisfies (11), for some $p \in \mathbb{R}_+^J$, and is the equilibrium bandwidth allocation to each flow on route r . The pair (x, p) forms a solution of equation (11) and the conditions*

$$p_j \geq 0, \quad p_j \left(\tilde{C}_j - \sum_r A_{jr} \hat{m}_r \pi_r \right) = 0 \quad j \in J. \quad (24)$$

The quantities $y_r := \min(\pi_r, x_r)$ solve the optimisation problem

$$\text{maximize} \quad \sum_{r \in R} w_r \frac{\kappa_r}{\eta_r} \frac{y_r^{1-\alpha}}{1-\alpha} \quad (25)$$

$$\text{subject to} \quad \sum_r A_{jr} \frac{\kappa_r}{\eta_r} y_r \leq \tilde{C}_j, \quad j \in J, \quad y_r \leq \pi_r, \quad r \in R, \quad (26)$$

$$\text{over} \quad y_r \geq 0, \quad r \in R, \quad (27)$$

and the p_j 's constitute a set of Lagrange multipliers associated with the capacity constraint \tilde{C}_j in the above. The y_r 's are thus uniquely determined. The x_r 's are not necessarily uniquely determined.

Proof. The expressions of the quantities of interest at any invariant point are obtained exactly as in the proof of Proposition 1. Rewriting (24) as

$$p_j \geq 0, \quad p_j \left(\tilde{C}_j - \sum_r A_{jr} \frac{\kappa_r}{\eta_r} \min(\pi_r, x_r) \right) = 0 \quad j \in J,$$

we can readily interpret the quantities $y_r := \min(\pi_r, x_r)$ as the solutions to the optimisation problem (25–27). By strict concavity of the function (25) being maximised, y is indeed uniquely defined. That x is in general not uniquely defined can be on the following counter-example: Consider a network with two links $\{1, 2\}$ of equal capacities C , and three routes $\{1\}$, $\{2\}$, $\{1, 2\}$. Routes $\{1\}$ and $\{2\}$ carry traffic with exactly the same characteristics. It is then easy to select parameters such that streaming traffic along route $\{1, 2\}$ experiences admission control, while streaming traffic along the two other routes is always accepted. In that case, only the sum of multipliers $p_1 + p_2$ is determined, not the individual multipliers. As a result the corresponding rates $x_{\{1\}}$, $x_{\{2\}}$ are not uniquely defined either. \square

We can again provide an interpretation for invariant points, in particular when for all r it holds that $y_r < \pi_r$. This is the case where admission control is active along each route. Note that in this specific situation, x is now uniquely determined and coincides with y .

Remark 7. *In the case the admission control is active along each route, i.e. for all r it holds that $y_r < \pi_r$, the equilibrium allocation x is the same under the present admission control mechanism as under the previous rate adaptation scheme, the invariant numbers of file transfers \hat{n}_r are also unchanged, and the load used up by streaming traffic, here $\hat{n}_r \pi_r$, is also unchanged. In this sense, the present admission control mechanism may be deemed “TCP-fair” at the level of detail captured by the present fluid flow models.*

Remark 8. *Consider the following modified admission rule: instead of choosing to proceed with probability $\min(1, x_r/\pi_r)$, type r streaming flows will instead chose to proceed with probability $\min(1, \epsilon x_r/\pi_r)$, where $\epsilon > 0$ is some fixed parameter. Denote by x_r^ϵ , m_r^ϵ and p_j^ϵ the corresponding equilibrium variables. It is readily seen that a valid solution is provided by choosing $\epsilon x_r^\epsilon = x_r^1$, $m_r^\epsilon = m_r^1$, and $p_j^\epsilon = \epsilon^\alpha p_j^1$. The interpretation is as follows: by basing their admission decision on the scaled down TCP-friendly rate ϵx_r rather than x_r , the real-time flows ensure that the rate obtained by file transfers in equilibrium is scaled up by ϵ^{-1} , but this comes at no cost for them, as they have the same admission probability in equilibrium.*

This stems from the fact that in the current model, file transfers contribute an incompressible load on the system, independent of their performance. If we were to consider an Engset-like model for file transfers, this property would no longer hold.

We expect the dynamical system (22–23) to be stable under the stability condition (6), however we have not proven this yet.




Figure 1: Impact of streaming traffic on file transfers. The substantial amount of streaming traffic present in mix (a) relative to mix (b) has a stabilizing effect on the number of flows in progress. Impact du trafic temps-réel sur les transferts de fichiers. La quantité importante de trafic temps-réel dans (a) relativement à (b) un effet stabilisant sur le nombre de flots en cours.

VI EXAMPLE: A STAR NETWORK

As a concrete example, consider a star network of 10 links connected to a core. This example is motivated by the current Internet, where the back-bone is relatively uncongested, and congestion occurs mainly on the access links. Flows use two links, with traffic spread randomly across links.

For the example, $J = \{1, 2, \dots, 10\}$, $R = \{(i, j) : i < j, i, j \in J\}$. The capacity of each link C_j was chosen equivalent to a T3 link (45 Mbit/s), for $j \in J$. The mean holding time of streaming traffic ($1/\eta_r$, $r \in R$) was taken to be 200 seconds, corresponding to voice traffic, with the mean file size ($1/\mu_r$, $r \in R$) taken to be 600kB. The arrival rates for the two types of traffic (ν_r , and κ_r) were chosen to be identical, giving a file-transfer traffic intensity of 0.5 on each link, and such that in equilibrium each flow has rate 25kbit/s (x_r). Under this regime the equilibrium number of flows of each type is 100 ($\hat{m}_r = \hat{n}_r = 100$) per route r , giving 900 flows of each type on each link j .

Figure 1(a) shows the evolution of the number of each type $\sum_r A_{jr} n_r$ and $\sum_r A_{jr} m_r$ on a typical *link*, obtained by simulation of the Markov chain of Section 3. Note that the two curves look very similar and have a mean of 900. The number of streaming flows in progress has a standard deviation of 30, while the number of file transfers has a slightly higher standard deviation of just over 40.

We now alter the offered load of each type of traffic, to keep the nominal quality (x_r) seen by the flows fixed while significantly altering the proportions of the two types of traffic. We make the file-transfer traffic intensity 0.995 on each link, with a very small amount of streaming traffic. The load was such that the equilibrium of the fluid model has $\hat{n}_r = 199$, $\hat{m}_r = 1$. (With so little streaming traffic we do not expect our fluid model to be a good approximation; as the amount of

streaming traffic decreases to zero, we expect the behaviour of the system to be better described by the Brownian model of [15].) In Figure 1(b) we plot the behaviour of $\sum_r A_{jr} n_r$ on two typical links: observe the different vertical scale in this figure, and the marked variability of the number of flows in progress. Comparing the two figures, we see that the substantial proportion of streaming traffic present, in Figure 1(a), has the effect of reducing the variability of the number of flows in progress. Of course Figure 1(b) concerns a fairly extreme case where there is a very small amount of streaming traffic. More generally, the larger the proportion of streaming traffic (for a given nominal quality) the lower the variability of the number of flows in progress, and hence the lower the variability of the bandwidth received by flows.

VII CONCLUSION

We have studied a flow level model of Internet congestion control, that represents the randomly varying number of flows present in a network. Bandwidth was assumed to be dynamically shared between file transfers and streaming traffic, according to a fairness criterion that includes TCP friendliness as a special case. Through the construction of an appropriate Lyapunov function we have established stability, under conditions, for a fluid model of the system. The presence of fair-sharing streaming traffic results in a non-degenerate fluid model. Analysis of the model suggests that file transfers are seen by streaming traffic as reducing the available capacity, whereas for file transfers the presence of streaming traffic amounts to replacing sharp capacity constraints by relaxed constraints. Simulations show that the integration of streaming traffic and file transfers has a stabilizing effect on the variability of the number of flows present in the system.

Acknowledgement Alan Bain and Frank Kelly were partially supported by EPSRC grant GR/S86266/01.

References

- [1] E. Altman, D. Artiges and K. Traore, On the Integration of Best-Effort and Guaranteed Performance Services, INRIA Research Report No. 3222, July 1997.
- [2] Andersen (S.), Blaabjerg (S.), Fodor (G.), Telek (M.), A Partially-blocking queueing system with CBR/VBR and ABR/UBR arrival streams. *Tel. Syst. J.* 19 (1) (2002) 75-99.
- [3] Bain (A.), Key (P. B.), Modelling the performance of distributed admission control for adaptive applications. *Performance Evaluation Review*, December 2001.
- [4] Ben Fredj (S.), Bonald (T.), Proutière (A.), Regnié (G.), Roberts (J.), Statistical bandwidth sharing: a study of congestion at flow level. In *Proceedings of SIGCOMM 2001*.
- [5] Bonald (T.), Massoulié (L.), Impact of fairness on Internet performance. In *Proceedings of ACM SIGMETRICS 2001*.
- [6] Breslau (L.), Knightly (E. W.), Shenker (S.), Stoica (I.), Zhang (H.), Endpoint admission control: Architectural issues and performance. In *Proceedings of SIGCOMM 2000*, 57-69.

- [7] Bu (T.), Towsley (D.), Fixed Point Approximation for TCP behavior in an AQM Network. In *Proceedings of ACM SIGMETRICS 2001*.
- [8] Courcoubetis (C.), Dimakis (A.), Reiman (M. I.), Providing bandwidth guarantees over a best-effort network: call admission and pricing. *IEEE INFOCOM 2001*, 459–467.
- [9] de Veciana (G.), Lee (T.-J.), Konstantopoulos (T.), Stability and performance analysis of networks supporting elastic services, *IEEE/ACM Trans. on Networking* (2001) **9**, 2–14.
- [10] Floyd (S.), Fall (K.), Promoting the use of end-to-end congestion control in the Internet. *IEEE/ACM Transactions on Networking* (1999) **7**, 458–472.
- [11] Floyd (S.), Handley (M.), Padhye (J.), Widmer (J.), Equation-based congestion control for unicast applications. In *Proc. ACM SIGCOMM 2000*, pages 43–54, Stockholm.
- [12] Gibbens (R. J.), Sargood (S. K.), Van Eijl (C.), Kelly (F. P.), Azmoodeh (H.), Macfadyen (R. N.), Macfadyen (N. W.), Fixed-point models for the end-to-end performance analysis of IP networks. 13th ITC Specialist Seminar: IP Traffic Measurement, Modeling and Management, Sept 2000, Monterey, California.
- [13] Karlsson (G.), Más (I.), Lundqvist (H.), “QOS PDQ”. Transactions of the Royal Institute of Technology (TRITA-IMIT-LCN R 03:06), ISRN KTH/IMIT/LCN/R-03/06–SE, 2003.
- [14] Kelly (F. P.), Key (P. B.), Zachary (S.), Distributed admission control. *IEEE Journal on Selected Areas in Communications* (2000), **18**, 2617–2628.
- [15] Kelly (F. P.), Williams (R. J.), Fluid model for a network operating under a fair bandwidth-sharing policy. *Annals of Applied Probability* (2004) **14**, 1055–1083.
- [16] Key (P. B.), Massoulié (L.), Probing strategies for distributed admission control in large and small scale systems. *IEEE Infocom* (2003).
- [17] IP monitoring project, Sprint labs. <http://ipmon.sprintlabs.com>
- [18] Massoulié (L.), Roberts (J.), Bandwidth sharing and admission control for elastic traffic. *Telecommunication Systems* (2000) **15**, 185–201.
- [19] Mathis (M.), Semke (J.), Mahdavi (J.), Ott (T.), The macroscopic behaviour of the TCP congestion avoidance algorithm. *Computer Communication Review*(1997) **27**, 67–82.
- [20] Mo (J.), Walrand (J.), Fair end-to-end window-based congestion control. *IEEE/ACM Transactions on Networking* (2000) **8**, 556–567.
- [21] Núñez-Queija (R.), van den Berg (H.), Mandjes (M.), Performance Evaluation of Strategies for Integration of Elastic and Stream Traffic. 16th International Teletraffic Congress, Edinburgh 1999.
- [22] Padhye (J.), Firoiu (V.), Towsley (D.), Kurose (J.), Modeling TCP Reno performance: a simple model and its empirical validation. *IEEE/ACM Transactions on Networking* (2000) **8**, 133–145.
- [23] RÁCZ (S.), Gerö (B.), Fodor (G.), Flow level performance analysis of a multi-service system supporting elastic and adaptive services. *Performance Evaluation* (2002) **49**, 451–469.
- [24] Rockafellar (T.), *Convex Analysis*. (1970) Princeton University Press.
- [25] Roughan (M.), Erramilli (A.), Veitch (D.), Network performance for TCP Networks, Part I: persistent sources. In *Proceedings of ITC’17 Brasil*, September 2001.