Mellin Transforms for TCP Throughput with Applications to Cross Layer Optimization

François Baccelli† & David R. McDonald‡
†INRIA-ENS, {Francois.Baccelli@ens.fr}
‡University of Ottawa, {dmdsg@mathstat.uottawa.ca}

Abstract— **In this survey paper, we review recent results on a class of partial differential equations associated with the dynamics of TCP. We show how to use the closed form solutions obtained for them via Mellin transforms in order to optimize the parameters of a variety of communications channels when TCP is used.**

Keywords— **TCP, congestion control, flow control, additive increase–multiplicative decrease algorithm, IP traffic, HTTP, throughput, coding rate, processing gain, synchronization.**

I. Introduction

The most basic formula for predicting the performance of long lived TCP flows is the square root formula, see [8]. This formula shows that the mean window size is inversely proportional to the square root of the probability a packet is dropped. Since the transmission rate of a source is the window size divided by the round trip time this formula determines the mean bandwidth allocated to such a TCP flow. This formula (as well as this survey) assumes a regime with a constant drop probability. Such a regime might arise with certain Active Queue Management (AQM) schemes which stabilize the flows through losses at a congested buffer. It also arises in the context of a wireless or DSL channel where losses are due to transmission errors.

In the first part of the paper (§II), we review results that were obtained lately on the distribution of the rate of a TCP flow using a partial differential equation approach describing the AIMD dynamics. The first aim of the survey is to show the versatility of the approach, which allows one the analyze the long lived flow case as well as the on-off case, the cases with or without capacity limitations, the case of a single flow (as in all the models of the first part) as well as that of a collection of flows (see §III-B in the second part) and also various types of losses.

The closed form solutions obtained for these systems pave the way for various optimizations of channel parameters adapted to the fact that TCP is used and in particular to the fact that this protocol reacts quite strongly to losses. Three examples of such cross-layer optimizations are given in the second part (§III) of the paper.

II. The Dynamics of TCP Flows

Any finite measure Φ on the positive real line has an associated Mellin transform

$$\widehat{\phi}(u) = \int_0^\infty z^{u-1}\Phi(dz) \text{ where } u \geq 1. \tag{1}$$

If Φ has a density ϕ then $\widehat{\phi}(u) = \int_0^\infty \phi(z) z^{u-1} dz$. Note that the Euler Gamma function $\Gamma(u)$ is the Mellin transform of the function $\phi(z) = e^{-z}$.

A. Persistent Flows

A.1 Poisson Loss Model

Consider a long lived TCP-Reno flow subject to random losses. Its RTT will be denoted by R. The rate of the flow has a linear increase with slope $\frac{1}{R^2}$. At each loss epoch, the rate experiences a multiplicative decrease which consists of a division by 2. Let $s(z,t)$ denote the probability density that a flow has a rate of $z \in \mathbb{R}^+$ at time $t \in \mathbb{R}^+$ (we admit that such a density exists). In the case where losses are Poisson with rate λ (this is the case for impulse noise in e.g. DSL lines), one obtains (see e.g. [5]) that this function satisfies the partial differential equation (PDE)

$$\frac{\partial s}{\partial t}(z,t) + \frac{1}{R^2}\frac{\partial s}{\partial z}(z,t) = \lambda(2s(2z,t) - s(z,t)). \tag{2}$$

The stationary density $s(z)$ satisfies the equation

$$\frac{ds}{dz}(z) = \beta(2s(2z) - s(z)), \tag{3}$$

with $\beta = \lambda R^2$.

When multiplying both sides of (2) by z and integrating w.r.t. z, it is easy to check that the mean value at time t, $\overline{X}(t) = \int_0^\infty z\, s(z,t)dz$, satisfies the ODE

$$\frac{d\overline{X}(t)}{dt} = -\frac{\lambda}{2}\overline{X}(t) + \frac{1}{R^2}.$$

This immediately gives

$$\overline{X}(t) = \left(\overline{X}(0) - \frac{1}{R^2}\frac{2}{\lambda}\right) e^{-\frac{\lambda t}{2}} + \frac{1}{R^2}\frac{2}{\lambda}. \tag{4}$$

Hence $\overline{X}(\infty) = \frac{1}{R^2}\frac{2}{\lambda}$. In steady state, the number mean number of lost packets per unit of time is λ. The mean number of packets transmitted per unit of time is $\overline{X}(\infty)$. Hence the steady state packet loss probability is $p = \frac{\lambda}{\overline{X}(\infty)} = \frac{R^2\lambda^2}{2}$. Hence $\lambda = \frac{\sqrt{2p}}{R}$, so that $\overline{X}(\infty) = \frac{\sqrt{2}}{R\sqrt{p}}$.

The following theorem (see [5]) gives a more complete characterization of the steady state of this PDE:

Theorem 1 *The unique stationary distribution solution of (2) has for Mellin transform*

$$\widehat{s}(u) = \phi\Gamma(u)\beta^{-u}\prod_{k\geq 0}(1-2^{-u-k}) \tag{5}$$

with $\phi = \beta\left(\prod_{k\geq 1}(1-2^{-k})\right)^{-1}$.

A.2 Packet Error Rate Model

We now consider the case where the transmission error rate of the flow is px when the throughput of the flow is x (like e.g. in the AQM and many wireless cases). In this case (as in the other cases considered below), the mean rate cannot be obtained by a direct analysis as above. The use of Mellin transforms nevertheless allows one to determine the solution of the ODEs and to derive closed form solutions in particular for the mean values. The PDE associated with this case is

$$\frac{\partial s}{\partial t}(z,t) + \frac{1}{R^2}\frac{\partial s}{\partial z}(z,t) = p(4zs(2z,t) - zs(z,t)), \tag{6}$$

so that the stationary density $s(z)$ satisfies the ODE

$$\frac{ds}{dz}(z) = \beta(4zs(2z) - zs(z)), \tag{7}$$

with $\beta = pR^2$.

For all $l \geq 0$, define

$$\Pi_l(u) = \prod_{k=0}^{l}\left(1-2^{-u-2k}\right). \tag{8}$$

The following theorem gives the steady state solution of the PDE (6).

Theorem 2 *The unique stationary distribution solution of (6) has for Mellin transform*

$$\widehat{s}(u) = \phi\Gamma\left(\frac{u}{2}\right)\left(\frac{2}{\beta}\right)^{\frac{u}{2}}\Pi_\infty(u), \tag{9}$$

with $\phi = \left(\sqrt{\pi}\left(\frac{2}{\beta}\right)^{\frac{1}{2}}\Pi_\infty(1)\right)^{-1}$. *In particular, its mean is*

$$\overline{X} = \sqrt{\frac{2}{\beta}}\sqrt{\frac{1}{\pi}}\frac{\Pi_\infty(2)}{\Pi_\infty(1)}. \tag{10}$$

Hence $\overline{X} \sim \frac{1.309}{R\sqrt{p}}$.

The results of the last two theorems have been known for quite a while. The last theorem was for instance proved under a variety of equivalent forms in e.g. [8], [1], [7], [5] to quote a few. We now survey much more recent results.

A.3 Maximal Rate

Consider now the situation where there is an upper-bound C on the rate of a flow. We assume here that when the rate of the flow reaches C, there is a congestion and the rate is divided by 2; this is a natural assumption in the particular case of a link with a small buffer, where the fact that link capacity is reached triggers a loss with high probability. Other scenarios can of course be contemplated like e.g. the case when the rate remains constant at C until the next loss (where C then represents the presence of a maximal window). The loss model considered below is that of (6). For our model, the stationary density s satisfies the differential equation

$$\frac{ds(z)}{dz} = \begin{cases} -\beta zs(z) & \text{if } C/2 < z \leq C \\ -\beta zs(z) + 4\beta zs(2z) & \text{if } 0 \leq z < C/2 \end{cases}, \tag{11}$$

where $\beta = pR^2$ (see [2]). The density $s(x)$ has its support on $(0, C)$, is discontinuous at $x = C/2$ and such that

$$s((C/2)^+) - s((C/2)^-) = s(C^-). \tag{12}$$

The following result is proved in [2].

Theorem 3 *The Mellin transform of the probability density $s(x)$ is*

$$\widehat{s}(u) = \frac{\sum_{l\geq 0}\Pi_l(u)C^u\left(\frac{\beta C^2}{2}\right)^l\frac{\Gamma(u/2)}{\Gamma(u/2+l+1)}}{\sum_{l\geq 0}\Pi_l(1)C\left(\frac{\beta C^2}{2}\right)^l\frac{\Gamma(1/2)}{\Gamma(1/2+l+1)}}. \tag{13}$$

In particular, the mean is

$$\overline{X} = \frac{\sum_{l\geq 0}\Pi_l(2)C^2\left(\frac{\beta C^2}{2}\right)^l\frac{1}{(l+1)!}}{\sum_{l\geq 0}\Pi_l(1)C\left(\frac{\beta C^2}{2}\right)^l\frac{\sqrt{\pi}}{\Gamma(1/2+l+1)}}. \tag{14}$$

It is not difficult to see that when C tends to infinity, this mean value tends to that of Theorem 2, whereas when p tends to 0, this tends to 3/4 C which is the mean value of the classical saw tooth model. The distribution (13) actually encompasses all hybridations between these two extremes.

B. Non-Persistent Flow

In this section we review results on non-persistent flows. The model we consider is that introduced in [6], where the flow alternates between ON and OFF periods. A flow is silent for a random OFF time with exponential distribution of mean $1/\sigma$. After the OFF period, it switches to ON and transmits a file. The distribution of file sizes is exponential with a mean $1/\tau$ packets. The flow is subject to a constant packet loss probability p. During the OFF time, the transmission rate is set to zero. When a flow switches from OFF to ON, a slow start phase takes place; this phase is approximated by an instantaneous jump of the rate of some random size. The size of this jump is a random variable with law H with Mellin transform $\widehat{h}(u)$. The trajectory of a typical flow is depicted in Figure 1. The steady state

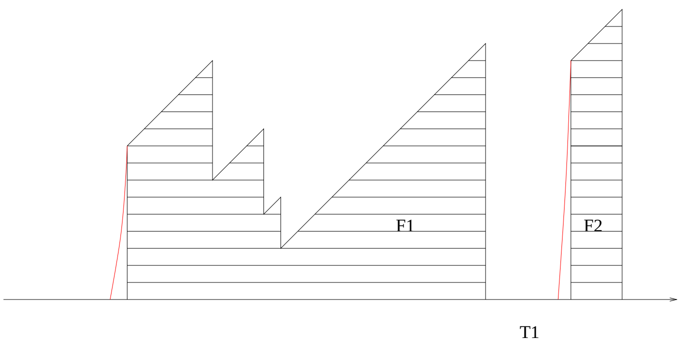

Fig. 1. Evolution of the rate of a non persistent flow with time. The exponential slow start phase is in red. The jump approximation in black.

probability that a flow is ON and has a rate z admits a density $s(z)$ which satisfies the ODE

$$\begin{aligned}\frac{ds}{dz}(z) &= -\tau R^2 zs(z) + \sigma\nu R^2 h(z) + 4zpR^2 s(2z) \\ &\quad -zpR^2 s(z) \qquad (15)\\ &= -\tau R^2 zs(z) + \tau(1-p)R^2 \int_0^\infty vs(v)dvh(z) \\ &\quad +4zpR^2 s(2z) - zpR^2 s(z), \qquad (16)\end{aligned}$$

where h is the density of the slow start jump. The stationary probability ν that this flow is OFF is such that $\int_0^\infty \tau(1-p)zs(z)dz = \sigma\nu$ and $\int_0^\infty s(z)dz + \nu = 1$. Let

$$\Psi_k(u) = \prod_{l=0}^{k-1}\left(1 - \frac{p}{p+\tau}2^{-u-2l}\right), \qquad (17)$$

for $k \geq 0$. The following result is proved in [5].

Theorem 4 *The mean time $\overline{T}$ (or mean latency) to transfer a file is:*

$$\begin{aligned}\overline{T} &= \frac{1}{1-p}\frac{\sqrt{\pi}}{\tau}\frac{\Psi_\infty(1)}{\Psi_\infty(2)}\sqrt{\frac{(p+\tau)R^2}{2}} \qquad (18)\\ &+\frac{\sqrt{\pi}R^2}{2(1-p)}\sum_{k=0}^{\infty}\left(\Psi_k(2)\frac{\Psi_\infty(1)}{\Psi_\infty(2)}\widehat{h}(2k+3)\frac{\left(\frac{(p+\tau)R^2}{2}\right)^{k+\frac{1}{2}}}{(k+1)!}\right. \\ &\left. -\Psi_k(1)\widehat{h}(2k+2)\frac{\left(\frac{(p+\tau)R^2}{2}\right)^{k}}{\Gamma(k+\frac{3}{2})}\right).\end{aligned}$$

The mean rate of a stationary flow is

$$\overline{X} = \frac{(\tau(1-p))^{-1}}{\sigma^{-1}+\overline{T}} \qquad (19)$$

and the probability that a flow is ON is

$$\nu = 1 - \overline{X}\tau(1-p)/\sigma. \qquad (20)$$

More generally, the Mellin transform of s is:

$$\begin{aligned}\widehat{s}(u) &= \widehat{s}(2)\Gamma(\frac{u}{2})\left(\frac{\Psi_\infty(u)}{\Psi_\infty(2)}\left(\frac{(p+\tau)R^2}{2}\right)^{1-\frac{u}{2}} + \right. \\ &\frac{\tau R^2}{2}\sum_{k=0}^{\infty}\left(\Psi_k(2)\frac{\Psi_\infty(u)}{\Psi_\infty(2)}\frac{\widehat{h}(2k+3)}{(k+1)!}\left(\frac{(p+\tau)R^2}{2}\right)^{k+1-\frac{u}{2}}\right. \\ &\left.\left. -\Psi_k(u)\frac{\widehat{h}(u+2k+1)}{\Gamma(\frac{u}{2}+k+1)}\left(\frac{(p+\tau)R^2}{2}\right)^{k}\right)\right). \qquad (21)\end{aligned}$$

with $\widehat{s}(2) = \overline{X}$.

When τ tends to zero, the last mean value for $\overline{X}$ tends to that of the persistent flow square root formula of Theorem 2. The case with no slow start is obtained when taking $\widehat{h}(u) = 0$ where we have:

$$\overline{X} = \frac{1}{\frac{\tau(1-p)}{\sigma} + \sqrt{\frac{\pi}{2}}R\frac{\prod_{l=1}^{\infty}\left(1-\frac{2p}{p+\tau}4^{-l}\right)}{\prod_{l=1}^{\infty}\left(1-\frac{p}{p+\tau}4^{-l}\right)}\sqrt{p+\tau}}. \qquad (22)$$

The last expression (derived in [6]) is the simplest exact formula extending the square root formula to the non persistent flow case.

C. Further Models

There are many potentially useful models along these lines when combining the options discussed above. For instance, what is the mean rate of a non-persistent flow with a capacity constraint? In additions, there are many interesting variants either when considering more general laws for the file sizes and the think times, or when considering more general laws for the loss processes (Poisson as in (2), with a stochastic intensity as in (6), or more general point processes, for instance controlled by a Markov chain as in Gilbert's model).

III. EXAMPLES OF CHANNEL PARAMETER OPTIMIZATIONS UNDER TCP

A. Optimal Processing Gain of a CDMA Flow using TCP

In many wireless models, one can trade a smaller loss rate against a smaller raw throughput. A typical example is that of a CDMA flow. When TCP runs on such a flow, what is the best compromise between a quite small loss rate and a bad raw rate or the opposite?

In [2] the following model was considered for analyzing this problem within the context of the the interaction of TCP and CDMA. A long lived CDMA flow transferring data is assumed to be given a target SINR of γ, a processing gain of m, and to send data in packets of size L.

Under these assumptions, the raw data rate is $C = 1/(mLT)$ packets per second, if one denotes by T the chip duration. When no FEC is used, a classical match filter model gives a packet error probability of

$$p = 1 - (1 - Q(\sqrt{m\gamma}))^L \approx LQ(\sqrt{m\gamma}), \qquad (23)$$

where $Q(\cdot)$ is the CDF of the zero mean unit variance Gaussian density, i.e. $Q(x) = (1/\sqrt{2\pi}) \int_x^\infty e^{-x^2/2} dx$.

When decreasing the processing gain, the raw data rate C increases proportionally to $1/m$ whereas the packet error probability p increases proportionally to $Q(\sqrt{m\gamma})$.

One can then use the results of Theorem 3 to determine the processing gain that optimizes the CDMA/TCP interaction in the case of small buffers.

In steady state the throughput of the TCP/CDMA flow, when expressed in bits/second, satisfies the ODE (11) with $\beta = pR^2/L^2$, where p is as defined in (23) and with C as above (see [2]).

Using the approximation $Q(x) \approx \frac{1}{\sqrt{2\pi}x} e^{-x^2/2}$ in (23), we get

$$p \sim L e^{-\frac{m\gamma}{2}} \frac{1}{\sqrt{2\pi m\gamma}}. \qquad (24)$$

If $\beta C^2/2$ is small, then the second order expansion of (14) gives

$$\overline{X} \approx \frac{3C}{4} - \frac{pR^2C^3}{L^2} \frac{11}{256}. \qquad (25)$$

which in turn gives the following expression for the mean rate:

$$\overline{X} \sim \frac{1}{4T} \left(\frac{3}{m} - \frac{1}{T^2} \frac{\xi}{m^3\sqrt{m}} e^{-\frac{m\gamma}{2}} \right), \qquad (26)$$

with $\xi = \frac{11R^2}{64L\sqrt{2\pi\gamma}}$. Differentiating with respect to m, we get that the optimal m solves

$$e^{-\frac{m\gamma}{2}} = \frac{6T^2}{\xi} \frac{m\sqrt{m}}{\gamma + \frac{7}{m}}. \qquad (27)$$

Here is a numerical example. Taking $L = 320$, $R = .1$, $\gamma = 3 \cdot 10^{-2}$, $T = 10^{-7}$, we get $\xi \sim 1.15 \cdot 10^{-6}$; so the optimal processing gain m satisfies

$$e^{-.015m} = 4.727 \cdot 10^{-9} \frac{m\sqrt{m}}{.03 + \frac{7}{m}}. \qquad (28)$$

The solution is $m = 459$. The associated value for $\beta C^2/2$ is $pC^2R^2/(2L^2) \sim 0.815$. One can justify the use of the above approximation by showing that the correction brought by the next order term is appr. 4/1000 of the value given by the proposed expansion (the actual optimum is around 450). It turns out that this value is quite different from (and of course much higher than) that optimizing the goodput of an unresponsive UDP flow, which is around 350. The reason for that should be clear: a TCP flow over-reacts to losses. In order for this flow to use the raw bandwidth offered to it correctly, the packet error probability should be small enough, and in any case much smaller than that of an unresponsive flow.

This approach can be used in a variety of related contexts like for instance that of the best coding rate and processing gain, or that of the best alphabet size in the adaptive coding case. For more on the matter, see [2].

B. Optimal Transmission Loss Rate for a Collection of TCP Flows Sharing a Link

When several persistent TCP flows share a common resource, a well known phenomenon, particularly strong when all flows have the same RTT, is that of packet loss synchronization: when the buffer fills in, many sources experience a loss at about the same time, and then react in a synchronous way by halving almost simultaneously their window. In case of small or moderate buffer, this may lead to a severe reduction of the total rate in the buffer, which may in turn lead to emptying the buffer and hence to some under-utilization of the link capacity.

In the model considered in this section, we analyze this phenomenon within the context of a large family of flows subject to transmission error losses and partially synchronized congestion losses. A typical example would be that of a family of N long lived DSL flows, all with the same RTT. Each flow is prone to transmission error losses. Two types of losses can be considered (separately or jointly): either impulse errors that occur according to a Poisson point process of intensity λ, or packet error transmissions that occur according to the model of e.g. (6) with packet error probability p. We will concentrate here on the Poisson impulse model, which was studied in [3] (quite similar results hold in the packet error rate case – see [4]).

These flows share a link of capacity $C \cdot N$. The assumption is that when the sum of the rates of the N TCP flows reaches C, then a proportion $0 < q \leq 1$ of these flows experience a congestion loss, all at the same time (q is referred to as the *synchronization rate* in the literature).

When letting N to infinity, one gets a so called mean-field interaction model which captures the dynamics of the limit

$$\overline{x}(t) = \frac{1}{N} \lim_{N\to\infty} \sum_{i=1}^{N} X_i(t), \tag{29}$$

where $X_i(t)$ denote the rate of flow i at time t. This limit may exhibit two types of behaviors as can be seen when using the results of §II-A.1. Let

$$\lambda^* = \frac{2}{C}\frac{1}{R^2}. \tag{30}$$

- If $\lambda > \lambda^*$, then an interaction-less regime is possible in view of (4) and Theorem 1: if the rate of each flow is sampled independently with an initial condition which is distributed according to (5), then each of them is in its steady state. The limit (29) exists for all t and is equal to $\frac{2}{\lambda}\frac{1}{R^2}$ as a result of the strong law of large numbers. Since $\frac{2}{\lambda}\frac{1}{R^2} < C$, the flows never interact in the limit. In other words, in this regime, transmission losses are sufficient to control the system without the need of congestion losses.
- In case $\lambda < \lambda^*$, (4) can be used to show that for all exchangeable initial conditions $X_i(0)$ such that the limit (29) holds at $t = 0$ with $\overline{x}(0)$ constant, there is an infinite number of congestion epochs. If a periodic regime with a congestion every τ seconds exists for $\overline{x}(t)$, one gets from (4) that τ is necessarily equal to

$$\tau = \frac{2}{\lambda}\log_e\left(1 + \frac{\lambda q}{2(\lambda^* - \lambda)}\right) \tag{31}$$

and that the mean aggregated rate $\overline{X} = \frac{\int_0^{\tau} \overline{x}(t)\, dt}{\tau}$ is then necessarily:

$$\overline{X} = C\frac{\lambda^*}{\lambda} - \frac{Cq}{2\log_e\left(1 + \frac{\lambda q}{2(\lambda^* - \lambda)}\right)}. \tag{32}$$

It is not difficult to show that $\overline{x}$ is an increasing function of λ for $\lambda \in [0, \lambda^*]$ and a decreasing function of λ for $\lambda \geq \lambda^*$; thus $\overline{x}$ achieves its optimal value, namely C, at $\lambda = \lambda^*$. In other words, some transmission losses need to be added in order to get the best average throughput within this context, as exemplified in Figure 2, which was obtained by simulation.

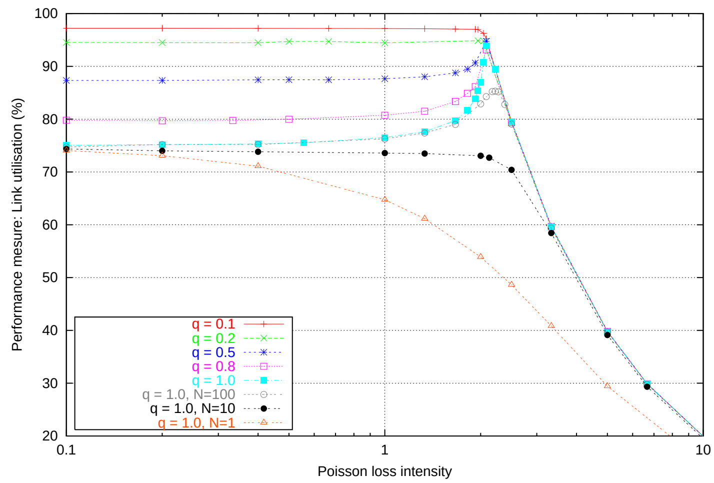

Fig. 2. Mean throughput vs Transmission loss rate

An intuitive explanation is that moderate transmission losses play the same role as RED and may improve performance by stretching congestion periods apart and hence decreasing synchronization.

C. Optimal Coding Rate of an HTTP Flow

When transferring a set of files via HTTP, is it worthwhile to use block codes that will reduce the probability of packet error but increase the file sizes? If so, what is the best coding rate?In this section (which is new), we show how to use the model of §II-A.2 and the results of Theorem 4 to answer this type of questions.

We consider some HTTP flow on a channel with a constant *bit error* rate p_b. File sizes and think times are assumed to be exponentially distributed. We assume that block codes are used for Forward Error Correction (FEC), with each packet of L bits encoded into N bits, with $N \geq L$. If a file has size F, its encoded version has size F/ρ where $\rho = L/N$ is the coding rate. The following model, which is based on the so called Gilbert-Varshamov bound ([9] p. 463) and used in [10], §III.A, and in [2], will be adopted below: for all coding rates ρ, there exists a block code that allows one to correct up to t bit errors in a packet, where t satisfies

$$\rho = 1 - h\left(\frac{2(t+1)}{N}\right), \tag{33}$$

with h the binary entropy function ($h(x) = -x\log_2(x) - (1-x)\log_2(1-x)$), provided $t+1 \leq N/4$. For $t = t(\rho)$ determined by (33), the packet error probability is hence

$$p = \sum_{j=t(\rho)+1}^{N} \binom{N}{j} p_b^j (1-p_b)^{N-j}.$$

All parameters of our HTTP model are now available: the file size has an exponential distribution with parameter $\tau\rho$ (when mesured in bits); the think time is exponential with parameter σ, and the packet error probability p is given by the last formula. The associated ODE giving the distribution of the throughput (expressed in bits per second) is

$$\frac{ds}{dz}(z) = -\tau L \frac{R^2}{N^2} zs(z) + \tau L(1-p)\frac{R^2}{N^2}\int_0^\infty vs(v)dvh(z) + p\frac{R^2}{N^2}\left(4zs(2z) - zs(z)\right), \quad (34)$$

where the probability distribution h is in bits. If one denotes by $\overline{X}$ the mean value of the distribution solution of the last equation, then the goodput is $\rho\overline{X}$ bits per seconds. We can then use the results of Theorem 4 to compute the mean goodput in function of the coding rate ρ.

Figure 3 gives a numerical example based on Equation (22). This figure displays the goodput in function of the coding rate. The parameters are the following: $p_b = 10^{-2}$; $R = 1s.$, the mean think time is 3.s, the mean number of bits in a file (before encoding) is 10^5 and each packet contains $N = 100$ bits (after encoding). Notice that this mean value is rather low as it takes into account the fact that the rate is 0 during the think times. There is numerical evidence that the optimal coding rate depends on the file size parameter.

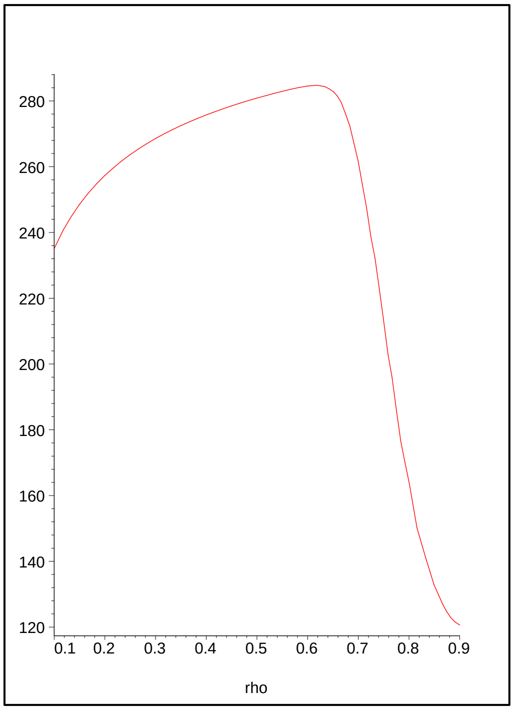

Fig. 3. Mean goodput vs coding rate

References

[1] Adjih, C., Jacquet, P., Vvedenskaya, N. (2001) Performance evaluation of a single queue under multi-user TCP/IP connections. *INRIA Research report #4141.*

[2] Baccelli, F., Cruz, R. and Nucci, A. (2005) Wireless Channel Parameters Maximizing TCP Throughput. To appear in the *Proceedings of the IMA workshop "Mathematics of Wireless Networks", June 05, Minneapolis*, G. Yin Ed. also available as INRIA Report 5652 at http://www.inria.fr/rrrt/rr-5652.html

[3] Baccelli, F. and Kim, K. B. (2004) TCP Throughput Analysis Under Transmission Error and Congestion Losses, *Proc. of IEEE Infocom 04*, Hong Kong.

[4] Baccelli, F., Kim, K. B. and De Vleeschauwer, D. (2005) Analysis of the Competition Between Wired, DSL and Wireless Users in an Access Network, *Proc. of IEEE Infocom 05*, Miami.

[5] Baccelli, F., Kim, K. B. and Mc Donald, D. (2005) Equilibria of a Class of Transport Equations Arising in Congestion Control, submitted for publication to *Questa*, also available as INRIA Report 5653 at http://www.inria.fr/rrrt/rr-5653.html

[6] Baccelli, F. and McDonald, D. (2005) A Square Root Formula for the Rate of Non-persistent TCP Flows. *Proceedings of the First EuroNGI Conference*, Roma, March 05.

[7] Baccelli, F., McDonald, D. R., Reynier, J. (2002) A mean-field model for multiple TCP connections through a buffer implementing RED. *Performance Evaluation Vol. 11, (2002) pp. 77-97.* Elsevier Science.

[8] Ott, T., Kemperman, J.H.B., and Mathis, M. (1996) The Stationary Behavior of Ideal TCP Congestion Avoidance. available at web.njit.edu/∼ott/Papers.

[9] Proakis, J.G. (1995) *Digital Communications*, McGraw-Hill, 3-rd Edition.

[10] Zhao, Q., Cosman, P. and Milstein, L. (2004) Optimal Allocation of Bandwidth for Source Coding, Channel Coding and Spreading in CDMA Systems, *IEEE Transactions on Communications*, vol. 52, No. 10, pp. 1797–1808.