

Available online at www.sciencedirect.com

Computer Networks xxx (2007) xxx–xxx

**Computer
Networks**
www.elsevier.com/locate/comnet

A game theoretic comparison of TCP and digital fountain based protocols [☆]

Luis López ^a, Antonio Fernández ^a, Vicent Cholvi ^{b,*}

^a *Universidad Rey Juan Carlos, Spain*

^b *Departamento de Lenguajes y Sistemas Informáticos, Universitat Jaume I, Campus de Riu Sec, 12071 Castellón, Spain*

Received 22 November 2006; received in revised form 11 January 2007; accepted 31 January 2007

Responsible Editor: G.S. Kuo

Abstract

In this paper we analyze a novel paradigm of reliable communication which is not based on the traditional timeout-and-retransmit mechanism of TCP. Our approach, which we call Fountain Based Protocol (FBP), consists of using a digital fountain encoding which guarantees that duplicate packets are almost impossible. By using Game Theory, we analyze the behavior of TCP and FBP in the presence of congestion. We show that hosts using TCP have an incentive to switch to an FBP approach, obtaining a higher goodput. Furthermore, we also show that a Nash equilibrium occurs when all hosts use FBP (i.e., when FBP hosts act in an absolutely selfish manner injecting packets into the network as fast as they can and without any kind of congestion control approach). At this equilibrium, the performance of the network is similar to the performance obtained when all hosts comply with TCP. Regarding the interaction of hosts using FBP at different rates, our results show that the Nash equilibrium is reached when all hosts send at the highest possible rate, and, as before, that the performance of the network in such a case is similar to the obtained when all hosts comply with TCP.

© 2007 Published by Elsevier B.V.

Keywords: Protocol analysis; Digital fountain codes; Game theory

1. Introduction

Congestion control in communication systems has been an important and largely studied issue. Since many communication systems in our days are based on the principle of sharing common resources (e.g., routers, communication links) among different users, one of the main objectives of congestion control schemes is to establish rules to guarantee that the common resources are used optimally and shared fairly among users. However,

[☆] A preliminary version of this paper appeared in the Proceedings of the IEEE Symposium on Computers and Communications, ISCC 2005. This work was partially supported by the Spanish Ministry of Science and Technology under Grants No. TSI2006-07799, No. TSI2004-02940 and No. TIN2005-09198-C02-01, by Bancaixa under Grant No. P1-1B2003-37 and by the Comunidad de Madrid under Grant No. S-0505/TIC/0285.

* Corresponding author. Tel.: +34 964728332; fax: +34 964728435.

E-mail address: vcholvi@uji.es (V. Cholvi).

34 most of these schemes require end-users to behave
35 in a cooperative way. Users have to respect some
36 “socially responsible” rules. For instance, the TCP
37 (which is, by far, the most widely used protocol)
38 congestion control scheme is voluntary in nature
39 and critically depends on end-user cooperation.
40 Indeed, TCP congestion control algorithms [1–6]
41 voluntarily reduce the sending rate upon receiving
42 some congestion signal such as ECN [7], packet loss
43 [8–10], or source quench [11]. Such congestion con-
44 trol schemes are successful because all the end-users
45 cooperate and voluntarily reduce their sending rates
46 upon detection of congestion.

47 Nevertheless, it is currently impossible to guaran-
48 tee that end-users will not act in a selfish manner. If
49 they use TCP, this means that they will never reduce
50 their sending rates even in the presence of conges-
51 tion. As it has been shown in [12,13], if this happens
52 and users overload the network, the total through-
53 put of the network drops. This happens since most
54 Internet routers use a drop-tail FIFO (First In First
55 Out) scheduling discipline, and users can obtain
56 more network bandwidth by transmitting more
57 packets per unit of time. (With this policy, the more
58 packets a user sends the more resources it gets.)
59 Thus, the optimal strategy for each user is strongly
60 suboptimal for the network as a whole.

61 Among the different techniques that can be used
62 to evaluate the impact of selfish users, one of the
63 most popular is *Game Theory* [14,15]. Game theory
64 is a tool for analyzing the interaction of decision
65 makers with conflicting interests. Roughly speaking,
66 a *game* has three components: a set of players, a set
67 of possible actions for each player, and a set of util-
68 ity functions mapping action profiles into real num-
69 bers. In our case, the *game players* are the users and
70 the congestion control schemes establish the *game*
71 *rules*. Each player has a strategy, which establishes
72 the traffic that it injects into the network.

73 The behavior of the TCP protocol has already
74 been addressed with a game-theoretic approach by
75 several authors. Some of the most remarkable
76 works in this field are the ones carried out by Nagle
77 [12,16], and Garg et al. [17]. Both of them show that
78 evil (selfish) behavior leads to disaster and propose
79 solutions based on creating incentive structures in
80 the systems that discourage this behavior. Nagle
81 suggests replacing the single FIFO queue associated
82 to each outgoing link with multiple queues, one for
83 each source host, which are served in a round-robin
84 fashion. Garg et al. introduce a novel and sophis-
85 ticated scheduling discipline called rate inverse

scheduling (RIS) that punishes evil behavior and 86
rewards cooperation, in such a way that the result- 87
ing Nash equilibrium¹ leads to a fair allocation of 88
resources. Both solutions require a significant 89
(sometimes huge) per-packet processing, which 90
might be impractical in many realistic applications 91
(as in Internet core routers, for example). Another 92
interesting work based on slightly different ideas is 93
the one carried out by Akella et al. [13]. In this 94
paper, a combination of analysis and simulations 95
is carried out trying to characterize the performance 96
of TCP in the presence of selfish users. The study 97
covers different variations of TCP (Reno, SACK, 98
etc.) and buffer management policies (Drop Tail, 99
RED, etc.), showing that the most recent variations 100
of TCP may become very inefficient in the presence 101
of selfish behavior. Nevertheless, they show that a 102
novel stateless buffer scheduling discipline called 103
CHOKe [18], which does not require per-packet 104
processing, may be useful in restoring the Nash 105
equilibrium efficiency. There are other interesting 106
proposals related to problems similar to this [19– 107
22]. In all cases, these works show the potential 108
applications of Game Theory within the problem 109
of congestion control and routing in packet 110
networks. 111

112 The above mentioned problem has a closer ana-
113 logue with the, so called, Tragedy of the Commons
114 [23] problem in economics. In this problem, each
115 individual can improve her own position by using
116 more of a free resource, but the total amount of
117 the resource degrades as the number of users
118 increases. Historically, this analysis was applied to
119 the use of common grazing lands, but it also applies
120 to such diverse resources as air quality and time-
121 sharing systems. In general, experience indicates
122 that multiplayer systems with this type of instability
123 tend to go into serious trouble. To understand pre-
124 cisely what a Tragedy of the Commons is, we need
125 first to observe that, in the context of Game Theory,
126 players choose their strategy in a selfish way trying
127 to maximize their benefit. If the system gets into a
128 state in which no player has an incentive to unilat-
129 erally change its strategy we say that the system has
130 reached the Nash equilibrium. In this context, a
131 game is a Tragedy of the Commons when (i) there

¹ An important concept in game theory is the Nash equilibrium. In our context, a Nash equilibrium is a scenario where no selfish user has incentive to unilaterally deviate from its current state. Clearly, being in a Nash equilibrium means that we are in a stable state in the presence of selfish users.

132 is always an incentive for a new player to become
 133 evil (this guarantees that the Nash equilibrium is
 134 reached when all players are evil) and (ii) the final
 135 benefit for evil players in the Nash equilibrium is
 136 under the initial benefit of fair players when all play-
 137 ers collaborate. This definition guarantees the essen-
 138 tial ingredient of a Tragedy of the Commons: if
 139 players behave in a selfish way, the Nash equilib-
 140 rium will be reached, and hence, the benefit of the
 141 defectors will always be less than the initial reward
 142 of the fair players. Hence, all players lose. In the
 143 context of network protocols, it has been observed
 144 by several authors [12,13,24] that when hosts behave
 145 in a selfish manner and do not comply with the TCP
 146 congestion control mechanisms (for example, by
 147 using lower timeouts), a Tragedy of the Commons
 148 arises and the network throughput drops due to
 149 the presence of duplicate packets. This effect can
 150 be easily observed in Fig. 1, which presents a simu-
 151 lation of a system like the one shown in Fig. 2.




Fig. 2. The picture shows the interconnecting topology of the network we use in the proposed model. N hosts try to access a common communication link through a router. This router has a finite buffer which drops packets when it is full. All lines have the same capacity C .

In this paper we compare, from a game theoretic point of view, TCP with a protocol based on digital fountain codes [25,26], which we call Fountain Based Protocol (FBP). The Digital Fountain approach has already been proposed as an appropriate mechanism for TCP-like reliable data transfer in multicast environments [27]. Moreover, suitable congestion control algorithms have been proposed to make these flows work in a TCP-friendly manner [28]. In this paper, we dig into these concepts, providing the following additional contributions:

- We propose an FBP for one-to-one reliable data transfer. This protocol is similar to UDP but uses fountain codes to avoid the presence of duplicate packets. Because of this, it does not require any type of packet retransmission mechanism. Contrary to UDP, FBP guarantees that the original source data will be correctly delivered, regardless of whether there are packet losses or not.
- Then, we establish a theoretical framework suitable for the analysis of this interaction of FBP and TCP. Under this framework, we show that users always have an incentive to switch from TCP to FBP. Furthermore, we validate the theoretical framework and results through simulations.
- We show that the Nash equilibrium of a network with a mixture of hosts using TCP and hosts using FBP is reached when all hosts behave in a selfish manner (by using FBP instead of TCP), but that this does not drive the network to a collapse. Moreover, we demonstrate that, in general, it does not even lead to a Tragedy of the Commons, since the throughput of hosts, even in the

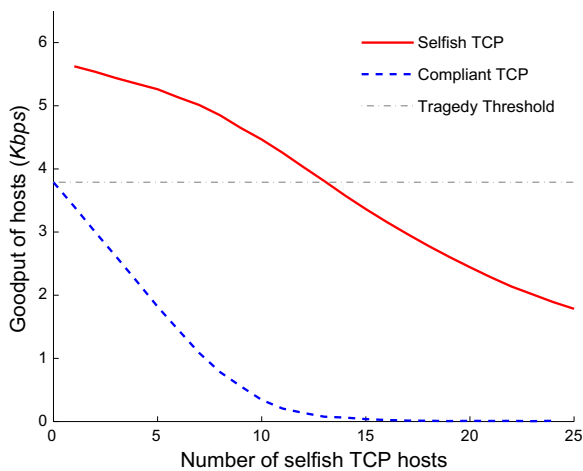


Fig. 1. This picture shows a typical Tragedy of the Commons scenario obtained by simulating with NS2 a single-bottleneck scenario with line capacities of $C = 100$ Kbps and a finite router buffer of 10 packets. Hosts can choose their strategy to be fair-cooperative (which complies with the TCP protocol) and evil-selfish (which implements a modified version of the TCP protocol in which the retransmission timeout is fixed to 0.4 s). The picture shows the throughput of fair and evil-selfish players as a function of the number N_e of selfish-evil hosts. As it can be observed, for any value of N_e , a given fair host always has an incentive to become evil. This implies that the Nash equilibrium is reached when all hosts are evil. Observe that the throughput in this equilibrium is remarkably smaller than the initial throughput of fair TCP hosts. Hence, the selfish strategy drives the game into a less efficient situation than the one obtained when all hosts cooperate. For this reason we say that a Tragedy of the Commons takes place.

186 case where all of them act in a selfishly way, is no
187 less than the throughput obtained when all host
188 comply with the TCP protocol.

189 • Finally, we also study the interaction of hosts
190 using FBP at different rates. Our results show
191 that the Nash equilibrium is reached when all
192 hosts send at the highest possible rate, and, as
193 before, that this does not lead to a Tragedy of
194 the Commons.

195 In the next section we present the details of the pro-
196 tocol FBP. In Section 3 we present the network
197 model we use, with some analytical results under
198 that model. In Section 5 we present simulations of
199 the same network and compare them with the previ-
200 ous analysis. In Section 6 we analyze systems with
201 only FBP hosts. Finally, in Section 7 we present
202 some concluding remarks.

203 2. Protocols based on digital fountain codes

204 The basic principle behind the use of *digital foun-
205 tain codes* [25,29,26] is conceptually simple. Roughly
206 speaking, it consists of generating a stream of differ-
207 ent encoded packets into the network, from which it
208 is possible to reconstruct the source data. The key
209 property is that the source data can be reconstructed
210 from any subset of the encoded packets of (roughly)
211 the same size as the source data. Such a concept is
212 similar to ideas found in the seminal works of
213 Maxemchuk [30] and Rabin [31].

214 A class of codes that satisfy the above mentioned
215 property are classical *erasure codes*. Erasure codes
216 generate additional redundant packets from the origi-
217 nal k packets of the source data. Then, they guar-
218 antee that the source data can be recovered from
219 any subset of $(1 + \varepsilon)k$ packets ($1 + \varepsilon$ is called the
220 *decoding inefficiency*). Hence, they allow to tolerate
221 packet losses during transmissions. For instance,
222 one can use Reed–Solomon erasure codes [32], since
223 they have the property that a decoder at the receiver
224 can reconstruct the original source data whenever it
225 receives any k of the transmitted packets (i.e., their
226 decoding inefficiency is 1). However, the encoding
227 and decoding processing times for such a class of
228 codes are prohibitive.

229 Digital fountain codes can be seen as a kind of
230 erasure codes with very fast encoding and decoding.
231 Furthermore, the number of encoded packets that
232 can be generated from the source data by using
233 these codes is potentially limitless and does not need
234 to be fixed ahead. That allows a digital fountain

235 code to take source data consisting of k packets
236 and produce as many encoding packets as needed
237 to meet the user demand. The only drawback is that
238 these codes have a decoding inefficiency a little lar-
239 ger than 1 (i.e., $\varepsilon > 0$).

240 *Fountain Based Protocols* use digital fountain
241 codes to appropriately encode data to be transferred.
242 Whenever a file has to be transmitted, a digital
243 fountain encoder is used to continuously generate
244 encoded packets. These packets are injected into
245 the network, by the sender, at a given rate. On its
246 turn, when the receiver has enough packets to
247 reconstruct the source data, it sends a `stop` mes-
248 sage to the sender. That is, the FBP does not require
249 any kind of congestion control mechanism. Further-
250 more, it does not make use of packet retransmis-
251 sions. The only “overhead” are the packets
252 injected in the time interval since the receiver sends
253 the `stop` message until the sender receives it. We
254 note that, in order to increase performance in real
255 scenarios, this simple protocol can be improved in
256 a number of ways (see [33] for an overview regard-
257 ing this issue).

258 For simplicity, in the next sections we will assume
259 that the decoding inefficiency of the used codes is 1.
260 In subsequent sections, we will analyze the effects
261 and consequences of having $\varepsilon > 0$. Current imple-
262 mentations of digital fountain codes can guarantee
263 an inefficiency of about 1.054 [29] and even less than
264 that [25,26] (up to 1.02). We will also assume that
265 the rate at which senders inject packets is constant
266 (i.e., it is CBR).

3. A model for the interaction between TCP and FBP 267

268 To understand the interaction between TCP and
269 FBP, we use the traditional single-bottleneck prob-
270 lem, in which a communication line is shared
271 between N different hosts, as depicted in Fig. 2.

272 In our analysis, we assume that time is discrete
273 and structured as a sequence of consecutive rounds,
274 where each round is a group of $S \geq N$ consecutive
275 slots. All communication lines are assumed to have
276 the same capacity, fixed to one packet per slot, and
277 all packets have the same size. Hosts are assumed to
278 be greedy (i.e., they always wish to send new packets
279 to the destination). The router is assumed to have a
280 finite buffer so that, when congestion occurs and the
281 buffer is full, new incoming packets are dropped.

282 Hence, we have a traditional Game Theory prob-
283 lem in which N different players (the hosts) compete
284 for a common resource (the shared line and the

router buffer) trying to obtain the maximum yield (the goodput). In this context, we assume that our hosts are free to choose between two different strategies when transmitting their packets. The first strategy is to comply with a given communication protocol suitable to solve the congestion problem of the single-bottleneck link. This protocol must be designed to fairly share the resources among the hosts. For this reason, following the traditional Game Theory notation, we say that players obeying this protocol are *fair*. On the other hand, hosts can adopt a different strategy consisting of sending packets as soon as they are ready and not complying with any given protocol designed to avoid congestion. These hosts will be called *evil* because they do not obey the established rules guaranteeing fairness in the game. Observe that, in a realistic communication environment, hosts using TCP could be considered as fair, while hosts using FBP would be evil because they do not take into account any congestion control mechanism. Thus, we say that TCP is an ordering protocol (in the sense that it enforces a set of fixed and known rules), while FBP is a disordered protocol (in the sense that it does not enforce any coordination).

Following, we describe the two protocols we will use:

The ordering protocol (TCP). The ordering protocol we use emulates the two main characteristics of TCP: resource sharing and congestion control. On one hand, it assigns one fixed exclusive slot to each host within each round, in which the host is allocated to send packets. Then, when all hosts use this ordering protocol, they transmit one packet per round, and this packet does not compete with any other to enter the router buffer. On the other hand, it implements a basic timeout-and-retransmit mechanism to control congestion. With this purpose, we introduce an acknowledgment scheme so that, whenever the destination receives a packet, it immediately generates an `ack`, which is sent back to the corresponding host in the subsequent time slot.

The disordered protocol (FBP). By using this protocol, hosts use some kind of digital fountain encoding which guarantees that duplicates are not possible and all packets reaching the destination are useful. As it has been said previously, a single `stop` message is sent at the end of the whole file transfer to indicate the sender that the transmission has ended. As a first approximation, we consider that the size of the files being exchanged is very large, and hence we disregard `stop` messages. We

consider that hosts that use FBP transmit on all slots of the round with a given probability p . For simplicity, we assume that the value of p is the same for all hosts.

Observation 1. Before continuing, we wish to remark that our model is not a totally realistic scenario where TCP and FBP could be competing. Actually, it is optimistic when estimating the fair (TCP based) yield, and pessimistic for the evaluation of the evil rates. The optimistic behavior occurs since our simplified ordering protocol does not react in any way when packets are lost, while current TCP implementations react to congestion by decreasing its offered load. In turn, the pessimistic behavior of FBP occurs since the decrease in the offered load of the TCP-based hosts would imply a higher probability for evil packets to get into the router buffer. Therefore, in our subsequent analysis, we will be using a scenario that penalizes FBP against TCP. In Section 5 this behaviour will be substantiated by means of experimental evaluation.

4. Analysis of the TCP–FBP Interaction

From the previous section, our communication scheme is based on rounds of S slots, with two kinds of slots. First, N_f fair slots (F-slots), where one fair host always transmits and N_e evil hosts transmit with probability p . Second, $S - N_f$ evil slots (E-slots) where N_e evil hosts transmit with probability p .

Before we proceed with the analysis, we note that, as it has been shown in [24], in scenarios where at least one of the hosts does not use any kind of congestion control mechanism, with high priority the router buffer is always full. Then, in that congested situation, only one packet can enter the buffer in each time slot, because only one packet gets out of it in that interval. Therefore, the probability of a given evil host with selfishness degree p to get a packet in the congested buffer in an E-slot can be easily calculated. To do so, just note that if we consider a particular evil host, the probability that the other $N_e - 1$ send $i - 1$ packets to the router is given by a binomial distribution of the form $\binom{N_e - 1}{i - 1} p^{i-1} (1 - p)^{N_e - i}$. As the considered host sends itself a packet with probability p , we have i packets trying to enter the router with probability $\binom{N_e - 1}{i - 1} p^i (1 - p)^{N_e - i}$. Given that the router admission policy is fair, if there are i packets trying

384 to occupy the single free buffer position, any of them
385 can get to it with probability $1/i$. Hence, summing
386 for all possible values of i , we have:
387

$$389 \quad p_E^e(N_e, p) = \sum_{i=1}^{N_e} \frac{1}{i} \binom{N_e - 1}{i - 1} p^i (1 - p)^{N_e - i}. \quad (1)$$

390 The probability of an evil host to get a packet into
391 the buffer in an F-slot can be calculated in the same
392 way, just noting that an additional fair host sends its
393 packet with probability 1
394

$$396 \quad p_F^e(N_e, p) = \sum_{i=1}^{N_e} \frac{1}{i + 1} \binom{N_e - 1}{i - 1} p^i (1 - p)^{N_e - i}. \quad (2)$$

397 With these results, we can evaluate the transmission
398 rate for evil hosts R_e . Since we assume evil hosts use
399 FBP, then all packets arriving to the destination (all
400 packets getting into the router buffer) are useful. So,
401

$$403 \quad R_e(N, N_e, p, S) = \frac{(S - N_f) p_E^e(N_e, p) + N_f p_F^e(N_e, p)}{S}. \quad (3)$$

404 For fair hosts the result is similar. First, the proba-
405 bility of a fair host to get its packet into the buffer in
406 its F-slot is

$$409 \quad p_F^f(N_e, p) = \sum_{i=0}^{N_e} \frac{1}{i + 1} \binom{N_e}{i} p^i (1 - p)^{N_e - i}. \quad (4)$$

410 Now, taking into account that fair hosts do not get
411 packets into the buffer in E-slots, we can evaluate
412 the transmission rate for fair hosts R_f . Namely,
413

$$415 \quad R_f(N, N_e, p, S) = \frac{p_F^f(N_e, p)}{S}. \quad (5)$$

416 From the analysis of the transmission rates for evil
417 and fair hosts (Eqs. (3) and (5)), we can derive some
418 interesting results.

419 **Property 1.** *An optimal protocol controlling the*
420 *congestion is just as good as letting all hosts to send*
421 *their FBP packets as fast as possible.*

422 **Proof.** To prove this property, we analyze the form
423 of the transmission rates for the two extreme situa-
424 tions. Namely, when all hosts are fair ($N_e = 0$) and
425 when all hosts are evil ($N_e = N$). The interesting fact
426 is to remark that if $p = 1$ then $R_e(N, N, 1, S) = \frac{1}{N} \geq$
427 $q R_f(N, 0, 0, S)$, which means that the best goodput
428 obtained when all hosts use an unordered protocol
429 is over the one obtained when they try to access
430 the common resource in an ordered way. \square

The key issue to understand why this happens is
to observe that, when using FBP, all packets arriv-
ing to the destination are useful and duplicates are
not possible. Many authors have remarked
[13,16,24] that when using a timeout-and-retransmit
based approach (as the one of TCP), if congestion
and flow control algorithms are not respected by
the hosts, the global throughput of the network
drops due to the presence of duplicates, which are
retransmitted when timeouts occur. Nevertheless,
when using FBP, no duplicates are present and no
timeouts are needed to ensure that the network is
not collapsed by them.

Property 2. *The Nash equilibrium of the game is*
reached when all hosts are evil.

Proof. For the proof, let us assume that $p > \frac{2}{N+1}$.
Then, it follows that $p > \frac{i+1}{Ni+n+1}$ for all $i \in \{1, \dots, N\}$
and for all $n \in \{0, \dots, N\}$. This can be seen by
assuming a worst case ($n = 0$), and by observing
that the inequality holds for $i = 1$ and that the
expression on the right strictly decreases with i .
The inequality can also be written as

$$455 \quad \frac{n+1}{i} + \frac{N-n-1}{i+1} > \frac{1}{ip}. \quad (6)$$

Now, we define $f_i = \binom{n}{i-1} p^i (1-p)^{n+1-i}$. Observe
that f_i is always positive. In this situation, we can
multiply Eq. (6) by f_i without changing the inequal-
ity. Hence, summing all the inequalities for all i
gives

$$462 \quad \sum_{i=1}^{n+1} \frac{n+1}{i} f_i + \frac{N-n-1}{i+1} f_i > \sum_{i=1}^{n+1} \frac{1}{ip} f_i.$$

Observe that substituting f_i , making a change of
variables in the second part of the inequality, divid-
ing by N and recovering the original expressions of
 p_E^e , p_F^e and p_F^f from Eqs. (1), (2) and (4), this can be
written as

$$469 \quad \frac{(n+1)p_E^e(n+1, p) + (N-n-1)p_F^e(n+1, p)}{N} > \frac{p_F^f(n, p)}{N}, \quad (7)$$

which using Eqs. (3) and (5) is equivalent to
 $R_e(N, n+1, p, S) > R_f(N, n, p, S)$. Then $R_e(N, N_e +$
 $1, p, S) > R_f(N, N_e, p, S)$ for all $N_e \in \{0, \dots, N-1\}$.
Therefore, in any given situation, a fair host always
has an incentive to become evil. \square

475 5. Simulations for the TCP–FBP interaction

476 The model we have just present allows under-
 477 standing some key issues in the interaction between
 478 TCP and FBP. Nevertheless, to gain a deeper
 479 insight into the TCP/FBP competition, we have car-
 480 ried out a number of simulations using a slightly
 481 modified version of the NS2 simulator. For these,
 482 we consider that all communication lines have a
 483 fixed capacity of $C = 100$ Kbps, with delays of 1
 484 ms and a router buffer of 10 packets. Fair hosts have
 485 been modeled using standard one-way TCP agents.
 486 FBP hosts have been implemented using modified
 487 UDP agents. In both cases, agents are greedy.

488 We use the goodput (including headers) as the
 489 measurement of the information transmitted by each
 490 player. The traffic of the TCP hosts has been imple-
 491 mented using the usual FTP application of NS2
 492 (which assumes that the file being transmitted is infi-
 493 nite). Fountain traffic has been implemented with
 494 CBR generators with the *random_* bit set (uniform
 495 distribution). This randomization is necessary to
 496 guarantee that the router does not benefit any of
 497 the hosts when dropping packets. (If a pure CBR is
 498 used, there may be time patters making some hosts
 499 more likely to introduce their packets into the rou-
 500 ter.) The buffer management policy is drop-tail and
 501 the scheduling discipline is FIFO. All the simulation
 502 results presented in this paper have been averaged for
 503 50 executions of the simulation scenario. Each execu-
 504 tion has been run for a simulated time of 30,000.

505 Note that it is possible to establish a direct paral-
 506 lelism between the TCP based hosts of the simula-
 507 tions and the fair hosts of the analytical model
 508 because both comply with a set of ordered rules
 509 which try of optimize the utilization of the shared
 510 resource avoiding congestion. In the same way,
 511 the evil hosts of the analytical model can be assim-
 512 ulated as the FBP (CBR-UDP) hosts of the simula-
 513 tions. In this case, the selfishness probability p can
 514 be easily calculated as the utilization of the corre-
 515 sponding line (the ratio between the offered load
 516 of the evil CBR source, λ_e , and the total capacity
 517 of the communication line C). For instance, since
 518 $C = 100$ Kbps, an evil host with $p = 0.5$ would cor-
 519 respond to an FBP agent using a CBR source of
 520 $\lambda_e = 50$ Kbps.

521 5.1. Optimal decoding inefficiency

522 The results of the simulations, as well as the pre-
 523 dictions of the simplified mathematical model when

considering optimal decoding inefficiency (presented
 above), have been depicted in Fig. 3.

526 The first thing we see is that our observation
 527 about the analytical model is correct. That is, the
 528 theoretical curve for fair hosts is optimistic and it
 529 remains always over the real goodput of the TCP
 530 hosts, and the one of evil hosts is pessimistic and
 531 stays all the time under the real FBP results. This
 532 confirms the validity of our arguments, in *Observa-*
 533 *tion 1*, about the analytical model.

534 Furthermore, it can be seen that an optimal pro-
 535 tocol controlling the congestion is just as good as
 536 letting all hosts to send their FBP packets as fast
 537 as possible in a selfish manner and without any kind
 538 of control. As we explained previously using the
 539 mathematical model (*Property 1*), this means that
 540 fair hosts always have an incentive to become evil,
 541 because in any possible situation the most rational
 542 strategy is to use FBP. Fig. 4 shows the *incentive*
 543 hosts have to become evil for different values of
 544 N_e , where incentive is defined as

$$\frac{R_e(N, N_e + 1, p, S) - R_f(N, N_e, p, S)}{R_f(N, N_e, p, S)}. \quad 546$$

547 Finally, the TCP (fair) rate when N_e hosts are evil
 548 (for any value of N_e) is always under the FBP (evil)
 549 rate when one more host becomes evil. This con-
 550 firms that, as explained using the mathematical
 551 model (*Property 2*), the Nash equilibrium is reached
 552 when all hosts are evil ($N_e = N$). This feature can be
 553 observed more clearly in Fig. 5, where we have rep-
 554 resented the simulated Nash equilibrium goodput
 555 and the simulated cooperative goodput for 10 differ-
 556 ent values of the load injected by the FBP hosts (10
 557 different values of p). This means that, in this partic-
 558 ular game, the selfish equilibrium is slightly more
 559 efficient than the global cooperation of TCP. Hence,
 560 we can claim that the Tragedy of the Commons is
 561 not present, at least under the assumptions we have
 562 accepted.

563 5.2. Suboptimal decoding inefficiency

564 For simplicity, in the previous sections it has
 565 been assumed that the decoding inefficiency of the
 566 used codes is 1. However, in a real situation, $\varepsilon > 0$,
 567 with typical values for ε in the range of [0.02,0.05].
 568 In this context, when the value of ε increases, the
 569 FBP (evil) goodput decreases in a factor of $1 + \varepsilon$
 570 with respect to the best case situation described pre-
 571 viously. The question that arises immediately is
 572 whether the same conclusions we described previ-




Fig. 3. The figure represents the goodput of evil (FBP) and fair (TCP) hosts as a function of the number of evil hosts N_e for four different values of evil selfishness. The shaded region represents the error between the theoretical model and the simulations. The horizontal dashed line indicates the simulated goodput of the TCP hosts when no evil players are present. All pictures have been calculated for $N = S = 25$.

ously would be obtained when the FBP hosts do not behave so optimally.

Here, we will study how the situation changes with ε . In Fig. 6 we have represented the goodput as a function of p for 4 different values of ε . The values have been normalized with respect to the TCP cooperative throughput, where there are not evil hosts. We can see that, for reasonable values of ε (smaller than approximately 0.05), the Nash equilibrium is slightly more efficient than the TCP solution, while for higher values of ε , it is slightly under the TCP throughput, and we have a (not very tragic) Tragedy of the Commons.

Hence, we show that it is possible to obtain a Nash equilibrium in the system that is less efficient than the TCP cooperative situation if the value of ε increases. That is, the system may fall in a Tragedy of the Commons. However, in contrast with the

results presented in [24,13,16], the tragedy is well bounded and the network would never collapse. The performance of the system is guaranteed to be very close to the value obtained in the TCP cooperative situation, at least for typical values of the parameter ε . This occurs since TCP does not have an efficiency of 100% in the utilization of the line.

5.3. Fast lines and high delays

In our analysis, we have implicitly accepted that the RTT of packets is low and that we do not have high speed communication lines. Observe that, in the mathematical model, large values of RTT or high speed communication lines decrease the throughput within a real TCP scenario because once the whole sender window has been transmitted, a host must wait either the arrival of an ack or a

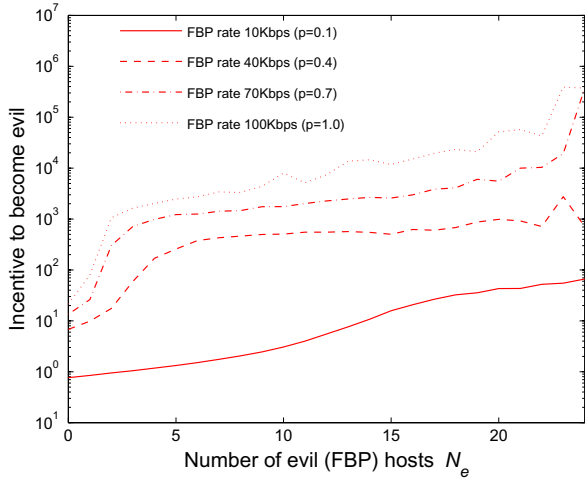


Fig. 4. Incentive to become evil as a function of the number of evil hosts N_e for different values of the load offered by FBP hosts. The simulation conditions are identical to the ones described in Fig. 3.

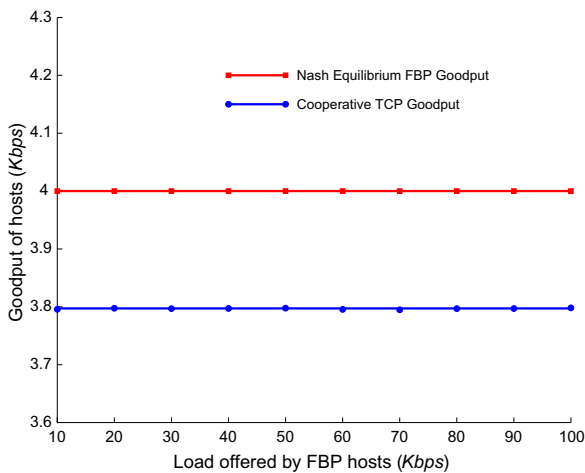


Fig. 5. Simulated FBP Nash equilibrium goodput (solid line with squares) and the simulated TCP cooperative goodput (solid line with circles) for 10 different values of the load offered by FBP hosts. The simulation conditions are identical to the ones described in Fig. 3.

607 timeout to be able to transmit something again.
 608 Therefore, this assumption implies that we have
 609 been considering the best TCP (fair) performance
 610 that can be found in a real scenario.

611 If we increase the RTT or the speed of the lines,
 612 the TCP goodput will fall. Hence, in all these cases,
 613 the Nash equilibrium will represent an even more
 614 efficient option than the all-TCP case. This can be
 615 easily observed in Fig. 7, where we show a situation




Fig. 6. Normalized goodputs for FBP as a function of the load offered by FBP hosts. A solid thick line with circles plots the normalized only-TCP goodput and a solid thin line with squares plots the normalized Nash equilibrium goodput for $\epsilon = 0$. The dotted lines represent the normalized Nash equilibrium goodput for $\epsilon > 0$: the stars indicate $\epsilon = 0.02$, the triangles $\epsilon = 0.05$, and the diamonds $\epsilon = 0.1$. The simulation conditions are identical to the ones described in Fig. 3.

616 similar to the one of Fig. 3, but where the speed of
 617 the communication lines has been increased to
 618 $C = 100$ Mbps. As it can be noticed, when the delay
 619 of the lines increases, the initial TCP goodput
 620 decreases, while the FBP Nash equilibrium remains
 621 constant. Note that for the same delay of 1 ms used
 622 in Fig. 3 the TCP goodput quickly degrades with the
 623 increase in the number of evil hosts.

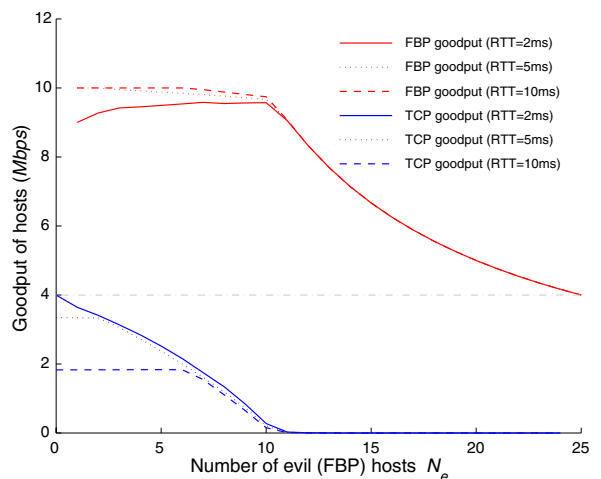


Fig. 7. Goodput for the FBP evil and TCP fair hosts as a function of the number of evil hosts N_e in the single-bottle-neck scenario of Fig. 2. The capacity of the lines has been fixed to $C = 100$ Mbps. The offered load of FBP hosts is 10 Mbps.

624 Taking into account the results obtained in this
625 section, the essential aspect that must be remarked
626 is that the introduction of FBP drives the system
627 to a Nash equilibrium where TCP disappears. Fur-
628 thermore, such an equilibrium has an efficiency that
629 can be slightly over or slightly under the one
630 obtained by using only TCP in most real situations,
631 but never drives the system into a collapse.

632 5.4. Effects of finite data files

633 Since the relationship between bandwidth (BW),
634 latency (measured as RTT) and size of target data
635 (D) is essential to evaluate the real goodput of the
636 FBP protocol, it is clear that, given the current def-
637 inition of the protocol, all packets arriving after the
638 last stop (ACK) signal is emitted by the receiver are
639 useless (basically because the data has fully been
640 decoded by that time). Given that, at Nash equilib-
641 rium, the sender emits packets at maximum rate, it
642 can be easily demonstrated that the product
643 $BW \cdot RTT$ corresponds to useless data. Hence, the
644 utility of the link can be evaluated as $U = D/(D +$
645 $BW \cdot RTT)$ being D the size of the original data
646 we want to transmit (we do not consider the decod-
647 ing inefficiency ϵ , which is discussed in another sec-
648 tion of the paper). With this equation in mind, it is
649 clear that the results provided in the paper are only
650 applicable when D is much larger than $BW \cdot RTT$.
651 In some cases, this could restrict the number of
652 applications for which FBP can be of real use, but
653 it is undoubtedly that there are many scenarios
654 where that condition is fulfilled; for example, in
655 the transmission of large video files (p2p applica-
656 tions, video on demand, etc), D is usually in the
657 range of some hundreds of megabytes, while the
658 $BW \cdot RTT$ product is rarely over one megabyte with
659 current Internet access capabilities (ADSL or
660 similar).

661 6. Congestion and fairness in FBPs

662 In the previous sections we have evaluated sys-
663 tems in which TCP and FBP hosts coexist. In this
664 section we analyze systems with only FBP hosts.
665 Our objective is to explore the situation when an
666 FBP host has a choice between sending packets at
667 a low rate and sending packets at a faster rate.
668 Hence, in our system we are going to have two clas-
669 ses of FBP hosts. *Slow* hosts will send packets at a
670 rate λ_{slow} (bits/second), while *fast* hosts will send
671 packets at a rate $\lambda_{fast} > \lambda_{slow}$. For simplicity we will

assume that $N\lambda_{fast} \geq C$, which implies that when all
hosts are fast the bottleneck link is fully used.

In order to analyze this system, we observe that
the behavior of the router can be approximated by
that of a queueing system $M/M/1/K$, where the buf-
fer of the router can hold $K - 1$ packets. The arrival
rate at this queue is $\lambda = N_e\lambda_{fast} + (N - N_e)\lambda_{slow}$ and
the service rate is $\mu = C$. Then, if we define $\rho = \frac{\lambda}{\mu}$,
using traditional queueing theory, we obtain that
the transmission rate of the bottleneck link is

$$\lambda' = \lambda(1 - p_K), \quad (8)$$

where

$$p_K = \begin{cases} \frac{1-\rho}{1-\rho^{K+1}} \rho^K & \text{when } \rho \neq 1 \\ 1/(K+1) & \text{when } \rho = 1 \end{cases} \quad (8)$$

Hence, the goodput for a slow host is

$$T_{slow} = \lambda_{slow}(1 - p_K), \quad (9)$$

while the goodput for a fast host is

$$T_{fast} = \lambda_{fast}(1 - p_K). \quad (10)$$

As we did in the previous section, we have used
NS2 to simulate a system with $N = 25$ hosts with
link capacities of $C = 100$ Kbps. In all the experi-
ments, hosts use, for different values of λ_{fast} ,
UDP-CBR packet generators with randomization.
Figs. 8 and 9 present the results of the simulations
compared with the queueing theory approach for
four different values of λ_{fast} , both when $\lambda_{slow} =$
 C/N (Fig. 8) and when $\lambda_{slow} < C/N$ (Fig. 9). As it
can be readily seen, the theoretical models fit very
nicely the results obtained by simulation. The small
differences have to do with the assumption that
packets have exponentially distributed lengths.

In these figures, it can be observed that, as in the
previous sections, the goodput of slow hosts when
there are N_e fast hosts is always smaller than the
goodput of fast hosts when there are $N_e + 1$ fast
hosts (for any $N_e < N$). Hence, slow hosts always
have an incentive to increase their sending rate. Fur-
thermore, this effect is more remarkable when
 $\lambda_{slow} < C/N$.

Another observation is that there is never a Tragic-
edy of the Commons. In the Nash equilibrium
(which is reached when $N_e = N$) all hosts evenly
share the resources like in the all-slow case, all
obtaining a goodput of C/N . This implies that if
the aggregation of slow rates fills the bottleneck link
(i.e., $\lambda_{slow} \geq C/N$), the Nash equilibrium yields the
same goodput as the all-slow case. However, if the




Fig. 8. Goodput of fast and slow hosts as a function of the number of fast hosts N_e for four different values of evil selfishness when $\lambda_{\text{slow}} = C/N = 4$ Kbps. The thick lines represent results from the simulations, the thin lines are predictions from the theoretical model based on queuing theory. The shaded region represents the error between the theoretical model and the simulations. The horizontal dashed line indicates the simulated throughput of the slow hosts when no fast hosts are present. All pictures have been calculated for $N = 25$ hosts.

722 slow rate is below C/N , the Nash equilibrium pre- 737
 723 sents a goodput larger than the all-slow case. 738

724 7. Concluding remarks 739

725 In this paper we have analyzed a novel paradigm 740
 726 of reliable communication which is not based on 741
 727 the traditional timeout-and-retransmit mechanism 742
 728 of TCP. Our approach, which we call Fountain 743
 729 Based Protocol (FBP), consists of using a digital 744
 730 fountain encoding which guarantees that all 745
 731 received packets are useful. Using Game Theory, 746
 732 we analyzed the behavior of TCP and FBP in the 747
 733 presence of congestion and show that two main 748
 734 characteristics arise. First, in this scenario, any 749
 735 given host using TCP has an incentive to switch 750
 736 to an FBP approach obtaining a higher through- 751

put. This guarantees the Nash equilibrium to be 737
 reached when all hosts use FBP. Second, we 738
 showed that, at this equilibrium, the performance 739
 of the network is similar (may be slightly over or 740
 slightly under) the performance obtained when all 741
 hosts comply with TCP. This latter claim holds 742
 even when FBP hosts act in an absolutely selfish 743
 manner injecting packets into the network as fast 744
 as they can and without any kind of congestion 745
 control mechanism. 746

The two above mentioned observations have 747
 direct implications in the context of the Internet. 748
 The first means that if FBP protocols are widely 749
 available for users, they will tend to employ them 750
 because they will obtain improved performance. 751
 Moreover, when more and more FBP hosts exist, 752
 the performance of the TCP players will decrease 753

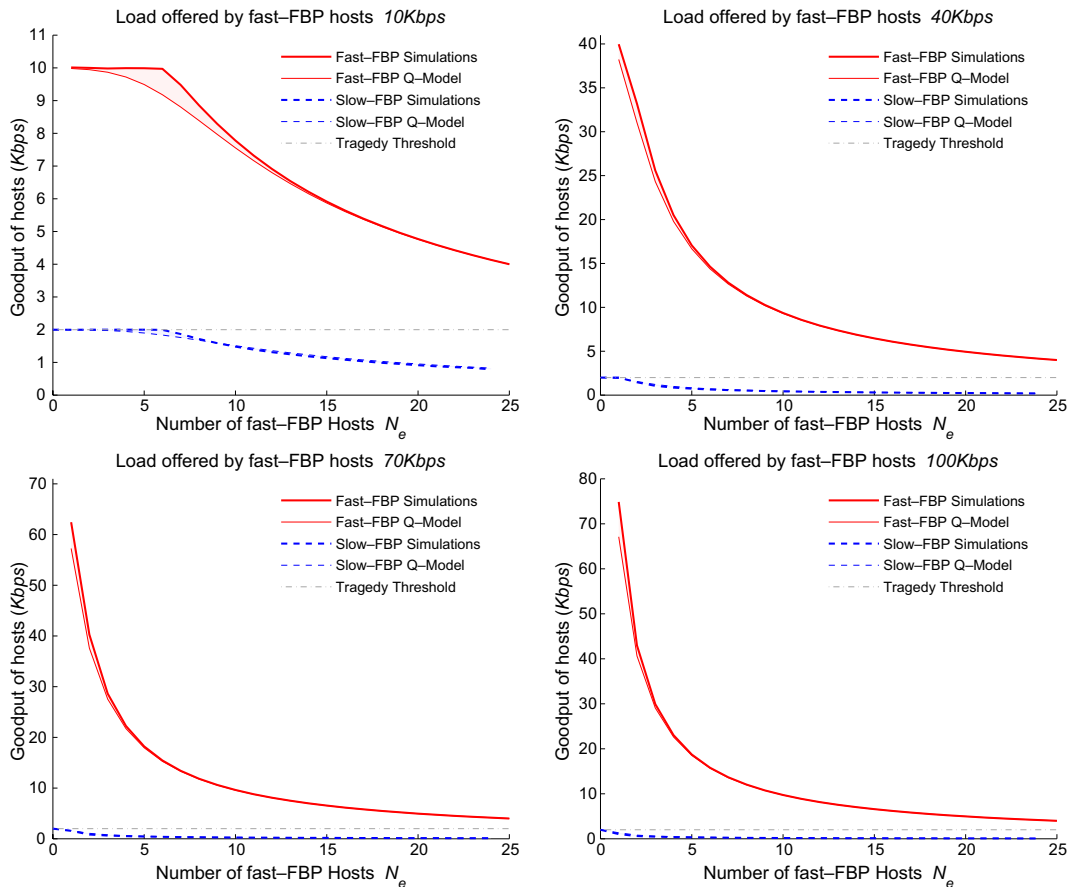


Fig. 9. Goodput of fast and slow hosts as a function of the number of fast hosts N_e for four different values of evil selfishness when $\lambda_{\text{slow}} = 2 \text{ Kbps} < C/N$. The thick lines represent results from the simulations, the thin lines are predictions from the theoretical model based on queuing theory. The horizontal dashed line indicates the simulated goodput of the slow hosts when no fast hosts are present. All pictures have been calculated for $N = 25$ hosts.

754 and the incentive to become evil will increase, possi-
 755 bly making that after some period of time all hosts
 756 become FBP. In this case, our second observation
 757 guarantees that the global performance of the net-
 758 work will not be under the original one which was
 759 obtained when only TCP hosts existed.

760 An aspect which merits some comments is the
 761 one relative to the architecture of current networks,
 762 which are designed to avoid congestion and to try to
 763 drop as few packets as possible. This fact could
 764 make the current Internet infrastructure to be seri-
 765 ously impaired by congestion if a large portion of
 766 users decides to switch to FBP. In this case, a new
 767 kind of routers would be necessary. This novel tech-
 768 nology should be designed to work under extremely
 769 congested scenarios with communication lines being
 770 saturated to nearly 100% of their capacities most of
 771 the time. In this new situation buffering could have a

772 limited utility, mainly contributing to increase net-
 773 work latencies.

774 Although these results seem promising, we wish
 775 to note that the FBP approach presents several
 776 aspects that should be taken into consideration.
 777 First, the analysis we have carried out has been
 778 done on the basis of large file transfers. However,
 779 this scenario can substantially change when consid-
 780 ering other kind of communications requiring more
 781 interaction between the sender and the receiver. For
 782 example, several real time or multimedia applica-
 783 tions (like Telnet) require small units to be contin-
 784 uously transferred. This scenario makes the FBP
 785 approach less practical, because, although duplicate
 786 packets cannot exist, it is possible that useless pack-
 787 ets not containing additional information could
 788 flood the network before the appropriate stop mes-
 789 sage issued by the receiver arrives to the sender.

790 Furthermore, in our analysis we have assumed that
 791 all packets arriving to the destination are useful. In
 792 reality, it could happen that a fast sender floods a
 793 slow receiver, which must drop packets. However,
 794 there are currently techniques that can be used to
 795 guarantee that receivers will not be saturated
 796 because of the fast sender rate (see for instance the
 797 mechanism used in [34]). Finally, another issue that
 798 deserves further attention is to analyze what hap-
 799 pens if we consider energy consumption issues in
 800 battery powered devices (which would waste a lot
 801 of energy). In those scenarios, the energy consump-
 802 tion is important and the use of FBP could be a
 803 problem. In these cases, it would be necessary to
 804 control the sending rate to avoid wasting a lot of
 805 energy due to the loss of many packets.

806 References

- 807 [1] S. Floyd, The NewReno modification to TCP's fast recovery
 808 algorithm, IETF Request for Comments 2582.
- 809 [2] L.S. Brakmo, S.W. O'Malley, L.L. Peterson, TCP Vegas:
 810 New techniques for congestion detection and avoidance, in:
 811 Proceedings of ACM SIGCOMM, London, UK, 1994, pp.
 812 24–35.
- 813 [3] M. Mathis, J. Mahdavi, Forward acknowledgement: refining
 814 TCP congestion control, in: Proceedings of ACM SIG-
 815 COMM, ACM Press, Palo Alto, California, United States,
 816 1996, pp. 281–291.
- 817 [4] M. Mathis, J. Mahdavi, S. Floyd, A. Romanow, TCP
 818 selective acknowledgment options, IETF Request for Com-
 819 ments 2018.
- 820 [5] J. Mo, R.J. La, V. Anantharam, J.C. Walrand, Analysis and
 821 comparison of TCP Reno and Vegas, in: Proceedings of the
 822 IEEE NFOCOM, 1999, pp. 1556–1563.
- 823 [6] J. Postel, Transmission control protocol, IETF Request for
 824 Comments 793.
- 825 [7] K. Ramakrishnan, S. Floyd, A proposal to add Explicit
 826 Congestion Notification (ECN) to IP, IETF Request for
 827 Comments 2481.
- 828 [8] S. Floyd, V. Jacobson, Random early detection gateways for
 829 congestion avoidance, IEEE/ACM Transactions on Net-
 830 working 1 (4) (1993) 397–413.
- 831 [9] V. Jacobson, Congestion avoidance and control, in: Pro-
 832 ceedings of ACM SIGCOMM, ACM Press, 1988, pp. 314–
 833 329.
- 834 [10] D. Lin, R. Morris, Dynamics of random early detection, in:
 835 Proceedings of ACM SIGCOMM, ACM Press, Cannes,
 836 France, 1997, pp. 127–137.
- 837 [11] W. Prue, J. Postel, Something a host could do with source
 838 quench: The source quench introduced delay (SQUID),
 839 IETF Request for Comments 1016.
- 840 [12] J. Nagle, On packet switches with infinite storage, IEEE
 841 Transactions on Communications 35 (4) (1987) 435–438.
- 842 [13] A. Akella, S. Seshan, R. Karp, S.S.C. Papadimitriou, Selfish
 843 behavior and stability of the internet: A game-theoretic
 844 analysis of TCP, in: Proceedings of ACM SIGCOMM,
 845 ACM Press, New York, 2002, pp. 117–132.
- [14] D. Fudenberg, J. Tirol, Game Theory, MIT Press, 1991. 846
- [15] G. Owen, Game Theory, Academic Press, 1995. 847
- [16] J. Nagle, Congestion control in IP/TCP internetworks, IETF 848
Request for Comments 896. 849
- [17] R. Garg, A. Kamra, V. Khurana, A game-theoretic 850
approach towards congestion control in communication 851
networks, Computer Communication Review 32 (3) (2002) 852
47–61. 853
- [18] R. Pan, B. Prabhakar, K. Psounis, CHOKe, A state- 854
less active queue management scheme for approximat- 855
ing fair bandwidth allocation, in: Proceedings of the 856
IEEE INFOCOM, IEEE, Los Alamitos, 2000, pp. 942– 857
951. 858
- [19] H. Yaïche, R.R.M.C. Rosenberg, A game theoretic frame- 859
work for bandwidth allocation and pricing in broadband 860
networks, IEEE/ACM Transactions on Networking 8 (5) 861
(2000) 667–678. 862
- [20] S.J. Shenker, Making greed work in networks: a game- 863
theoretic analysis of switch service disciplines, IEEE/ACM 864
Transactions on Networking 3 (6) (1995) 819–831. 865
- [21] T. Roughgarden, T. Éva, How bad is selfish routing? Journal 866
of the ACM 49 (2) (2002) 236–259. 867
- [22] C. Papadimitriou, Algorithms, games, and the internet, in: 868
Proceedings of the Thirty-third Annual ACM Symposium on 869
Theory of computing, ACM Press, Hersonissos, Greece, 870
2001, pp. 749–753. 871
- [23] G. Hardin, The tragedy of the commons, Science 162 (1968) 872
1243–1248. 873
- [24] L. López, G. Rey, A. Fernández, S. Paquelet, A mathemat- 874
ical model for the TCP tragedy of the commons, Theoretical 875
Computer Science 343 (1) (2005) 4–26. 876
- [25] M. Luby, LT codes, in: 43rd Annual IEEE Symposium on 877
Foundations of Computer Science, 2002. 878
- [26] A. Shokrollahi, Raptor codes, IEEE Transactions on Infor- 879
mation Theory, in press. 880
- [27] J. Byers, M. Luby, M. Mitzenmacher, A digital fountain 881
approach to asynchronous reliable multicast, IEEE Journal 882
on Selected Areas in Communications 20 (8) (2002) 1528– 883
1540. 884
- [28] J. Byers, G. Horn, M. Luby, M. Mitzenmacher, W. Shaver, 885
Flid-dl: Congestion control for layered multicast, IEEE 886
Journal on Selected Areas in Communications 20 (8) (2002) 887
1558–1570. 888
- [29] J. Byers, M. Luby, M. Mitzenmacher, A digital fountain 889
approach to the reliable distribution of bulk data, IEEE 890
Journal on Selected Areas in Communications 20 (8) (2002) 891
1528–1540. 892
- [30] N.F. Maxemchuk, Dispersivity routing, in: Proceedings of 893
ICC, San Francisco CA, 1975, pp. 41–10, 41–13. 894
- [31] M.O. Rabin, Efficient dispersal of information for secu- 895
rity, load balancing, and fault tolerance, Journal of the 896
Association for Computing Machinery 36 (2) (1989) 335– 897
348. 898
- [32] I.S. Reed, G. Solomon, Polynomial codes over certain finite 899
fields, Journal of the Society for Industrial and Applied 900
Mathematics 8 (2) (1960) 300–304. 901
- [33] B. Raghavan, A.C. Snoeren, Decongestion control, in: 902
Proceedings of the Fifth Workshop on Hot Topics in 903
Networks (HotNets-V), ACM Press, 2006, pp. 61–66. 904
- [34] H. Schulzrinne, S. Casner, R. Frederick, V. Jacobson, RTP: 905
A transport protocol for real-time applications, IETF 906
Request for Comments 3550. 907

912
908
914
915
916
917
918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935



Luis López obtained his Ph.D. in Computer Science at Universidad Rey Juan Carlos in 2003. He has a Telecommunications Engineering degree by the ETSIT-UPM, obtained in 1999, and a Telecommunications Engineering degree by the ENST – Télécom Paris, also in 1999. Before starting his academic career, he worked in several IT companies including Texas Instruments Inc (working as Hardware Engineer) and

Visual Tools S.A (working as Development Engineer).

He is currently a Professor at the DITTE department belonging to Universidad Rey Juan Carlos, where he coordinates/collaborates in several undergraduate and master courses in Computer Science and Telecommunications Engineering. His current research interests are concentrated on the applications of Complex Networks and Game Theory in the field of Computer Science in general, and IP networks in particular. He has coauthored more than 50 research publications in top scientific journals and conferences including Physica Review E, Theoretical Computer Science, IEEE Electronics Letters, etc.

938
939
940
941
942
943
944
945
946
947
948
949
950
951



Antonio Fernández is an associate professor at the Universidad Rey Juan Carlos in Madrid, Spain, since 1998. Previously, he was on the faculty of the Universidad Politécnica de Madrid. He graduated in Computer Science from the Universidad Politécnica de Madrid in 1991, and got a Ph.D. in Computer Science from the University of Southwestern Louisiana in 1994. His research interests include data communications,

computer networks, parallel and distributed processing, algorithms, and discrete and applied mathematics.



Vicent Cholvi graduated in Physics from the Universitat de València (Spain) and received his doctorate in Computer Science in 1994 from the Polytechnic University of Valencia (Spain). In 1995, he joined the Jaume I University in Castelló (Spain) where he is currently an associate professor. His interests are in distributed and communication systems.

954
955
956
957
958
959
960
961
962

953
963