

EFFORT / GAINS DYNAMICS IN HETEROGENEOUS NETWORKS

L. Mamatas and V. Tsaoussidis

Dept. Of Electrical and Computer Engineering
Demokritos University of Thrace, Greece
Email: emamatas@ee.duth.gr, vtsaousi@ee.duth.gr

Abstract

We investigate the behavior of TCP(α, β) protocols in the presence of wireless networks. We seek an answer to strategic issues of maximizing energy and bandwidth exploitation, without damaging the dynamics of multiple-flow equilibrium. We take a fresh perspective on protocol design: What is the return of the effort that a protocol expends? Can we achieve more gains with less effort? We study first the design assumptions of TCP(α, β) protocols and discuss the impact of equation-based modulation of α and β on protocol efficiency. We introduce two new metrics to capture protocol behavior: The “Extra Energy Expenditure” and the “Unexploited Available Resource Index”. We confirm experimentally that, in general, smoothness and responsiveness constitute a tradeoff; however, we show that this tradeoff does not graft its dynamics into a conservative/aggressive behavior, as it is traditionally believed. We uncover patterns of unjustified tactics; our results suggest that an adaptive congestion control algorithm is needed to integrate the dynamics of heterogeneous networks into protocol behavior.

1. INTRODUCTION

Transmission control of reliable protocols, as exemplified by TCP [1], is based on somewhat “blind” increase/decrease window mechanism that exploits the bandwidth availability dynamically and, meanwhile, avoids persistent congestion. The adjustments are modeled on the Additive Increase/Multiplicative Decrease algorithm from the perspective of fair resource allocation and efficient resource utilization [5]. AIMD is the core algorithm of standard TCP and is becoming the core algorithm of all transport protocols that support congestion control functions [6].

Several different mechanisms / protocols have been proposed regarding the transport layer. A thorough analysis of the different approaches to congestion control in transport protocols can be found in [11]. For example, TFRC [7] calculates throughput via a throughput equation that incorporates the loss event rate, round-trip time and packet size. TCP-Vegas [3] estimates the level of congestion using throughput-based measurements. TCP-Vegas demonstrates that measurement-based window adjustments is a viable mechanism, however, the corresponding estimators can be improved. In TCP-Westwood [4], the sender continuously measures the effective bandwidth used by monitoring the rate of returned ACKs. TCP-Real [15] uses wave patterns: a wave consists of a number of fixed-sized data segments sent back-to-back, matching the inherent characteristic of TCP to send packets back-to-back. The protocol computes the data-receiving rate of a wave, which reflects the level of contention at the bottleneck

link. Bimodal congestion avoidance and control mechanism [2] computes the fair-share of the total bandwidth that should be allocated for each flow, at any point, during the system's execution. TCP-Jersey [16] operates based on an "available bandwidth" estimator to optimize the window size when network congestion is detected.

However, a concern is how efficiently do protocols administer the network resources; that is, is the (whatever) gain proportional to the expended effort? Also, how can we measure effectively the effort / gain relative performance of a transport protocol? The evaluation of the proposed transport protocols needs to focus on the effort/gain dynamics of their corresponding mechanisms too.

The problems of standard TCP have been mainly investigated from two different perspectives, namely the application requirements and the characteristics of the underlying networks. The former expounds the impact of the transmission gaps caused by halving the transmission rate during congestion on the quality of delay-sensitive applications. Authors in [7, 8, 17, 18] propose TCP-friendly protocols that satisfy two fundamental goals: (i) To achieve smooth window adjustments. This is done by reducing the window decrease ratio during congestion. (ii) To compete fairly with TCP flows. This is approached by reducing the window increase factor according to a steady-state TCP throughput equation. It has been effectively established that TCP can achieve application-oriented improvements by favoring smoothness using a gentle backward adjustment upon congestion, at the cost of lesser responsiveness (i.e., speed to approach an equilibrium) - through moderated upward adjustments. The latter perspective unfolds the need for error detection and classification that would permit a responsive strategy, oriented by the nature of the error detected (congestion in wired networks versus transient random errors in wireless networks) [14]. As we show, implementation of such strategy requires occasionally a more responsive TCP. Our approach, however, is dominated by the distinctive characteristics and requirements of wireless networks: we address issues of energy and wireless error recovery, through a parallel study of a smooth/responsive protocol design and an aggressive/conservative outcome. Note that the conservative-through-to-aggressive behavioral spectrum reflects the effort a protocol expends. The real issue, therefore, is how much this effort is invested into efficient transmission.

TCP(α, β) protocols parameterize the congestion window increase value α and decrease ratio β , where the sender's window size is increased by α if there is no packet loss in a round-trip time, and the window is decreased to β times the current value if there is a loss indication. We discuss the impact of the smoothness/responsiveness tradeoff on protocol performance, assuming that it follows strictly the friendliness-oriented α/β tradeoff. A natural question is therefore "under what network conditions can we achieve efficiency; and how do we define efficiency". Having shown in previous work [13] that a protocol for wireless networks may need to be occasionally more conservative and occasionally more aggressive, we attempt to explore how this tradeoff is shaped by the responsive or smooth protocol strategy. In our discussion below, we refer to three classes of TCP(α, β) protocols: (i) Standard New Reno TCP(1, $\frac{1}{2}$); (ii) Responsive TCP(α, β), with *relatively* low β value and high α value; and (iii) Smooth TCP(α, β), with *relatively* high β value and low α value.

We compare the performance of our TCP(α, β) versions in heterogeneous (wired and wireless) networks and in static and dynamic¹ environments. Based on the assumptions of equation-based congestion control and on experimental data, we arrive at the conclusion that protocols, which are based entirely on the α/β tradeoff may be adequate for specific applications, networks and scenarios; however, they are inappropriate for several other occasions.

We organized the paper as follows: we give an overview of TCP(α, β) protocols in section 2 and we discuss their inherent assumptions. In section 3 we define new performance metrics. In section 4 we present our testing methodology and in section 5 we analyze the results of our experiments. Finally, in section 6 we highlight our conclusions.

2. TRADING α FOR β

A throughput equation for standard TCP is first introduced in [12]. GAIMD [18] extends the equation to include parameters α and β :

$$T_{\alpha, \beta}(p, RTT, T_0, b) = \frac{1}{RTT \sqrt{\frac{2b(1-\beta)}{\alpha(1+\beta)}} p + T_0 \min\left(1, 3\sqrt{\frac{(1-\beta^2)b}{2\alpha}} p\right) p(1+32p^2)} \quad (1)$$

where p is the loss rate; T_0 is the retransmission timeout value; b is the number of packets acknowledged by each ACK. The overall throughput of TCP-Friendly (α, β) protocols is bounded by the average throughput of standard TCP ($\alpha=1, \beta=0.5$), which means that equation (2), which is derived from (1) (see [18]) could provide a rough guide to achieve friendliness.

$$T_{\alpha, \beta}(p, RTT, T_0, b) = T_{1, 0.5}(p, RTT, T_0, b) \quad (2)$$

Authors of [18] derive from (1) and (2) a simple relationship for α and β :

$$\alpha = 4(1 - \beta^2) / 3 \quad (3)$$

Based on experiments, they propose a $\beta = 7/8$ as the appropriate value for the reduced the window (i.e. less rapidly than TCP does). For $\beta = 7/8$, (3) gives an increase value $\alpha = 0.31$.

The observations of the window dynamics and event losses are frequently assumed within a time period of a *congestion epoch* [7], which reflects the *uninterrupted growing lifetime of congestion window*. More precisely, a congestion epoch begins with βW packets, increased by α packets per RTT and reaching a congestion window of W

¹ From the perspective of the participating flows with criterion whether their number is fixed or not.

packets, when a packet is dropped. The congestion window is then decreased to βW . Hence, a congestion epoch involves

$$n = (1-\beta) * W / \alpha + 1 \text{ RTTs} \quad (4)$$

Assuming that the capacity of the bottleneck link is B packets per second and the number of active flows going through the bottleneck router is N , and assuming a control system as in [5], we further calculate that:

$$W = B * RTT / N \quad (5)$$

We can easily observe that *it takes several RTTs for a small α to pay back the bandwidth credit of a high β .*

Equation (1) is modeled by calculating the average throughput over a congestion epoch, which is associated with several RTTs. Since equation (1) gives the *steady state* TCP throughput, in a dynamic network where conditions changing rapidly, friendliness might not be attained. More precisely, based on (4) we conclude that (1) and (2) can be achieved at a time n RTTs or later since multiple drops will extend further the time of convergence. Based on (4) and (5) we further conclude that the time period required for (1) and (2) to hold is in reverse proportion to contention within a fixed bandwidth channel; the smaller the number of flows, the larger the window and therefore the longer the convergence time. By the same token, the fact that a responsive protocol can exploit bandwidth better suggests that lower contention is a favorable case for such protocols.

This analysis implies that, smooth protocols may be more aggressive (since they consume temporarily more bandwidth) in the presence of transient errors, while they may behave more conservatively, due to their low increasing rate, when multiple drops force the multiplicative decrease factor to adjust the congestion window back to its initial value. This can be justified by a hidden assumption behind (3): when packet drops occur at the end of the congestion epoch, the window decreasing by a factor of $(1-\beta)$ is applied only once. However, multiple packet drops could cause the window size to be decreased multiple times, or they could also cause the retransmission timer to expire. At the end, it is possible that the window size and the *ssthresh* could be decreased down to 2 segments, even with smooth backward adjustments. Under such scenarios, the performance of applications (including real-time applications) is not affected by how slowly the sender reduces its sending rate, but rather by how fast it can recover from the error and restore its sending rate. Note that our scenario is not unrealistic. For example, in mobile networks, burst correlated errors and handoffs generate this kind of error pattern. The aggressiveness of responsive TCP may be a desirable behavior. We confirm our statements experimentally in section 5.

3. METRICS FOR EVALUATING EFFORT / GAIN DYNAMICS

For a proper evaluation of the effort / gain dynamics, we propose new metrics in order to monitor:

- the expended effort of a protocol.
- the effective utilization of available resources.
- the achieved gain of the effort from the applications viewpoint.

Additionally, the new metrics need to be combined with traditional metrics:

The system goodput is used to measure the overall system efficiency in bandwidth utilization. The system Goodput is defined as:

$$\text{Goodput} = \text{Original_Data} / \text{Connection_time}$$

where Original_Data is the number of bytes delivered to the high-level protocol at the receiver (i.e. excluding retransmitted packets and overhead) and Connection_time is the amount of time required for the data delivery.

Fairness is measured by the Fairness Index, derived from the formula given in [5] and defined as:

$$\text{Fairness} = \frac{(\sum_{i=0}^n \text{Throughput}_i)^2}{n(\sum_{i=0}^n \text{Throughput}_i^2)}$$

where Throughput_i is the Throughput of the i_{th} flow and n the flow number. This Fairness Index provides a sort of “average-case” analysis used by most researchers. In order to conduct a “worst-case” analysis and provide a tight bound on fairness, we propose the *Worst-Case Fairness* as:

$$\text{WorstCaseFairness} = \frac{\min_{1 \leq i \leq n} \text{throughput}_i}{\max_{1 \leq i \leq n} \text{throughput}_i}$$

The range of worst-case fairness is also in [0, 1], with 1 representing the greatest fairness. As an example demonstrating why worst-case fairness is introduced, consider a scenario of 6 flows, the throughputs of which are 9 Mbps, 9.5 Mbps, 8.5 Mbps, 9 Mbps, 9 Mbps, and 6 Mbps, respectively. The traditional “average-case” fairness index is 0.982, while the worst-case fairness is 0.667. Compare this scenario with a perfectly fair case in which all flows achieve 9.5 Mbps, and both the “average-case” fairness index and worst-case fairness index are 1.0. The difference between the first scenario and the ideal case cannot be obviously distinguished by the “average-case” fairness index. In the first scenario, the system is fair in general, but is particularly unfair to the 6th flow. This unfairness to a very small fraction of flows can only be captured by the worst-case fairness.

In order to validate the efficiency of the protocols in real-time / multimedia applications, we assume an application, which demands every 100ms to receive at least one packet. Because of the sending window fluctuation and the transmission gaps of TCP(α , β), there are instances when data is unavailable to the application (because the packets were delayed more than 100ms). We use the *Application Success Index* in order to measure the protocol’s real-time performance:

$$\text{ApplicationSuccessIndex} = \frac{\text{PacketsDeliveredinTime}}{\text{PacketsDelivered}}$$

where $PacketsDeliveredinTime$ is the number of packets which have been received by the application in time and $PacketsDelivered$ is the total number of packets received by the application.

In order to capture the amount of *extra* energy expended, we introduce a new metric. Extra Energy Expenditure (3E) [10] attempts to capture the *extra* energy expended due to protocol operation – not just the expended energy. That is, a protocol may transmit when there windows of opportunities for error-free transmission, without expending extra energy, or vice versa. In contrast, it may waste opportunities for transmission expending energy (even in an idle state) and extending communication time. 3E attempt to capture extra energy expenditure as an associated result of Goodput, Throughput and maximum Throughput, each one represented as a moving point on a line. 3E takes into account the difference of achieved Throughput from maximum Throughput ($Throughput_{max}$) for the given channel conditions along with the difference of Goodput from Throughput, attempting to locate the Goodput as a point within a line that starts from 0 and ends at $Throughput_{max}$. The metric 3E takes values from 0 to 1, attempting to capture both distances.

$$EEE = a \frac{Throughput - Goodput}{Throughput_{max}} + b \frac{Throughput_{max} - Throughput}{Throughput_{max}}$$

where $a=1$ and $b=0.3$

When Goodput approaches Throughput, which approaches 0, the extra expenditure is only due to time waiting (probably in an idle state). We assume that the extra expenditure at this stage is 0.3 (the first term is 0). Instead, when $Goodput=Throughput=Throughput_{max}$ the extra expenditure is 0, since all the expended energy has been invested into efficient transmissions. Also, when $Throughput_{max}=100$, $Throughput=99$, $Goodput=1$, the extra expenditure due to unsuccessful retransmission grows to an almost maximum value (0.993)

We need to introduce another metric as well, in order for us to capture the level of *Unexploited Available Resources* (UAR) [10]. That is, how well did we exploit the windows of opportunities for successful transmissions. More precisely, holding transmission when conditions call for transmission, will perhaps result in minor energy expenditure but have a great cost on protocol goodput. Reasonably, the case of $Goodput=Throughput=0$ should not give us at this point a minor (as with the 3E metric) but a major penalty.

$$UAR = 1 - \left[a \frac{Throughput}{Throughput_{max}} + b \frac{Goodput}{Throughput} \right]$$

where $a=0.5$ and $b=0.5$ ². The UAR index ranges also from 0 to 1, expressing also a negative performance aspect.

² The a, b values in both EEE & UAR indices can be modeled on the behavior of a specific wireless network device.

The protocol efficiency can be studied from another perspective. *Overhead* is used as a metric to realize the protocol transmission effort to complete reliable data delivery.

$$Overhead = \frac{BytesSent - OriginalBytes}{BytesSent} = 1 - \frac{Goodput}{Throughput}$$

BytesSent is the total bytes transmitted by TCP senders, while OriginalBytes is the number of bytes delivered to the higher level protocol by receivers, excluding retransmitted packets and TCP header bytes. This metric captures the portion of consumed bandwidth, or the percentage of the transmission energy (a scarce resource in mobile computing), that is wasted on packet retransmissions and protocol header overhead. It differs from *EEE* in that it only captures the extra energy expenditure due to retransmissions.

4. EXPERIMENTAL METHODOLOGY

4.1 Evaluation Plan

We have implemented our testing plan on the ns-2 network simulator. The network topology used as a test-bed is the typical single-bottleneck *dumbbell*, as shown in Figure. 1. The bw_1 link is 10Mbps and the bw_3 link is 1Mbps. The link's capacity (bw_2) is 10Mbps, unless it is explicitly stated otherwise. We used equal number of source and sink nodes. We simulated a heterogeneous (wired and wireless) network with ns-2 error models, which were inserted into the access links at the sink nodes. The Bernoulli model was used to simulate link-level errors with configurable bit error rate (BER). We did not use an ARQ mechanism in the link layer. The number of flows occasionally changes for the different scenarios. The simulation time was fixed at 60 seconds, a time-period deemed appropriate to allow all protocols to demonstrate their potential.




Figure 1. Simulation topology

Due to the deterministic nature of the experiments, statistical validity is not an issue. In order to validate our statements, we selected and evaluated three protocols that satisfy the TCP-friendly equation [18]. We used standard New-Reno TCP (1, 0.5), a responsive New-Reno TCP (1.25, 0.25) and a smooth New-Reno TCP (0.31, 0.875).

In the scenarios without graduated contention decrease, ftp flows are entering the system within the first two seconds. All flows are fixed, during the rest 58 seconds. In order to evaluate how efficiently and fairly the protocols can exploit available bandwidth, we used, additionally, scenarios with graduated contention decrease.

5. RESULTS AND DISCUSSION

In this work, we comment on five different scenarios:

1. A simple wired scenario
2. A wireless scenario with low Error-Rate
3. A wireless scenario with low Error-Rate and Graduated Contention Increase
4. A scenario with long handoffs
5. A scenario with brief handoffs

We focus on protocol behavior with respect to effort expended and gains achieved. Although effort expended is a rather unified metric (i.e. packets transmitted over time), gain achieved is application-specific and cannot be expressed in a unified manner.

5.1 Simple wired scenario

In the simple wired scenario, the three TCP variations have similar gains in terms of throughput (figure 5). While the responsive TCP have better fairness index (figure 2), the smooth TCP have less extra energy expenditure (figure 3). An adaptive transport protocol could follow, in this situation, either a responsive or a smooth transmission tactic, in order to be fair or energy efficient. So, there is a tradeoff between fairness and energy efficiency. The responsive TCP does not exploit the network resources very well in case of fewer flows than 40 (figure 6). This results in an increased UAR index (figure 4). Based on the effort / gain perspective, we note that the increased effort that the responsive TCP expends (figure 3) makes the protocol more fair (figure 2). However, it has no performance gains (figure 6). In case of real-time / multimedia applications, the traditional TCP achieves more gains, as indicated by its application success index (figure 7)




Figure 2. Fairness




Figure 3. EEE




Figure 4. UAR




Figure 5. Throughput




Figure 6. Goodput




Figure 7. Application Success

5.2 Scenario with low Error-Rate (error-rate 0.02 BER)

In our second scenario, we simulated a heterogeneous environment with random transient errors (error-rate 0.02 BER). We measured performance, ranging the number of flows from 10 to 100. Based on figure 11 we note that the responsive TCP has an increased UAR index. In this situation, the responsive TCP follows a conservative strategy. So, there is not always a direct relationship between responsiveness and aggressiveness. Although the responsive TCP is not most efficient (figures 12, 13), it appears more fair (figures 8, 9). This result is interesting and calls for further research.




Figure 8. Fairness




Figure 9. Worst Case Fairness




Figure 10. EEE




Figure 11. UAR




Figure 12. Throughput




Figure 13. Goodput




Figure 14. Application Success

5.3 Scenario with low Error-Rate and Graduated Contention Increase

We simulated a heterogeneous environment with random transient errors (error-rate 0.01 BER). Additionally, we were gradually increasing the contention level. While the protocols exhibit similar behavior in terms of throughput, goodput and energy efficiency (figures 19, 20, 17), the responsive protocol is more fair (figures 15, 16) because it adjusts faster to the corresponding contention level. On the other hand, the smooth protocol has a better application success index since the smooth behavior reduces the jitter (figure 21). So, the same effort (figure 17) can result, in case of different transport mechanisms, in different gains. If the transport protocol was aware of the application's demands, it could have selected a different strategy to reach the desired gain.




Figure 15. Fairness




Figure 16. Worst Case Fairness




Figure 17. EEE




Figure 18. UAR




Figure 19. Throughput




Figure 20. Goodput




Figure 21. Application Success

5.4 Scenario with long Handoffs

In this scenario, we used handoffs with 1 second duration. We measured performance, ranging the number of flows from 10 to 100. Although the smooth TCP is less fair (figure 22), it has outperformed the other protocols in terms of Extra Energy Expenditure (figure 23) because they have increased overhead (figures 25, 26). The increased UAR index (figure 24) shows that the smooth protocol has not exploited the available resources efficiently. So, it can become more efficient using a sophisticated congestion control algorithm. The less effort expenditure of the smooth TCP (figure 23) leads to a slight difference in terms of goodput (figure 26). On the other hand, it is not suitable for real-time / multimedia applications because of the increased jitter (figure 27).




Figure 22. Fairness




Figure 23. EEE




Figure 24. UAR




Figure 25. Throughput




Figure 26. Goodput




Figure 27. Application Success

5.5 Wireless Scenario with brief handoffs

In this scenario, we used handoffs with low duration (0.1 sec). We measured performance, ranging the number of flows from 1 to 10. While all protocols have the same behavior in terms of energy efficiency, goodput and throughput (figures 30, 33, 32), smooth TCP is unfair (figures 28, 29) and achieves worse application success (figure 34). This result calls for investigating further whether a smooth behavior is always more suitable for applications, which demand low jitter. From the effort / gain perspective, which is our present focus, we note that different transport mechanisms can have different gains (either having the same effort or not). In this example, while the smooth protocol is more suitable for application, which demand fairness (like applications on grid computing), it is not appropriate for multimedia / real-time applications (because of the increased jitter).




Figure 28. Fairness




Figure 29. Worst Case Fairness




Figure 30. EEE




Figure 31. UAR

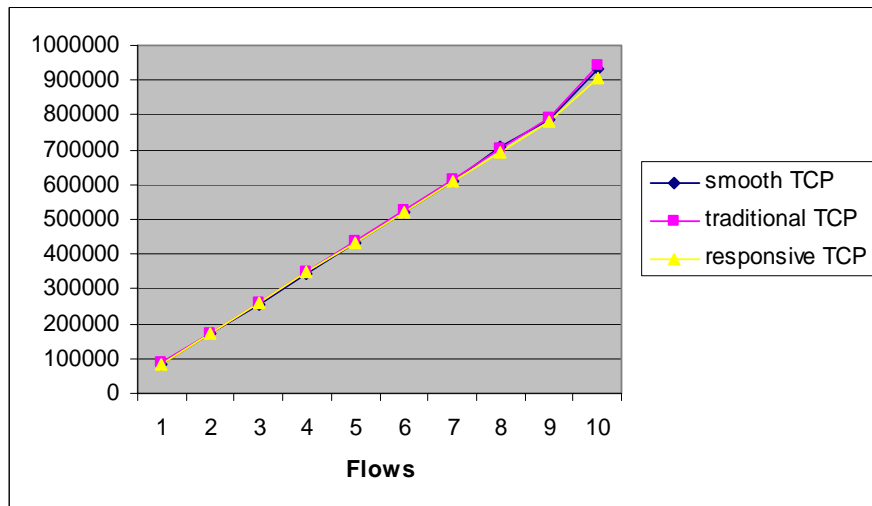


Figure 32. Throughput




Figure 33. Goodput




Figure 34. Application Success

6. CONCLUSIONS AND FUTURE WORK

In this paper, we seek an answer to the following questions, regarding the transport protocol design:

- How can we measure the effort that a protocol expends?
- What is the return of this effort?
- Are the achieved gains suitable for our application?
- Are the achieved gains worth of the expended effort?
- Can we have more gains with less effort?

If TCP's traffic could be shaped to conform to the appropriate effort/gain dynamics of the application, system performance in terms of goodput, fairness, energy efficiency and resource utilization would be handled better.

We have shown that smooth/responsive protocols do not always have a conservative/aggressive behavior respectively, as it is was traditionally believed. We have predicted through a basic analysis and confirmed experimentally a better behavior of conservative protocols in a high error-rate environment, in contrast to the aggressive ones. In case of low error-rates and sufficient availability of bandwidth, the situation is reversed.

Initially, we plan to work towards a measurement-based detection of network characteristics, such as the one presented in [15]. Departing from there, we plan to apply error recovery tactics, which integrate the adaptive strategy, in accordance with the results shown here.

7. REFERENCES

- [1] M. Allman, V. Paxson, and W. Stevens, "TCP Congestion Control", RFC2581, April 1999.
- [2] P. C. Attie, A. Lahanas and V. Tsaoussidis, "Beyond AIMD: Explicit fair-share calculation", In Proceedings of ISCC 2003, June, 2003.

- [3] L. Brakmo and L. Peterson, "TCP Vegas: End to end congestion avoidance on a global Internet", IEEE Journal on Selected Areas in communication, 13(8):1465--1480, October 1995
- [4] C. Casetti, M. Gerla, Saverio Mascolo, M.Y. Sansadidi, and Ren Wang, "TCP Westwood: End-to-End Congestion Control for Wired/Wireless Networks" In Wireless Networks Journal 8, 467-479, 2002
- [5] D.-M. Chiu and R. Jain, "Analysis of the Increase and Decrease Algorithms for Congestion Avoidance in Computer Networks", Computer Networks and ISDN Systems, 17(1):1-14, 1989.
- [6] S. Floyd, "Congestion Control Principles", RFC 2914, September 2000.
- [7] S. Floyd, M. Handley and J. Padhye, "A Comparison of Equation-based and AIMD Congestion Control", May 2000. URL: <http://www.aciri.org/tfrc/>.
- [8] S. Floyd, M. Handley, J. Padhye, and J. Widmer, "Equation-Based Congestion Control for Unicast Applications", Proceedings of ACM SIGCOMM 2000, August 2000.
- [9] Handley, M., Floyd, S., Padhye, J., and Widmer, J, "TCP Friendly Rate Control (TFRC): Protocol Specification", RFC 3448
- [10] L.Mamatas, V.Tsaoussidis: Protocol Behavior: More Effort, More Gains? The 15th IEEE International Symposium on personal, indoor and mobile radio communications, Barcelona, Spain.
- [11] L.Mamatas, V.Tsaoussidis and C.Zhang, "Approaches to Congestion Control in packets networks", Technical Report: TR-DUTH-EE-2004-10, Demokritos University of Thrace
- [12] J. Padhye, V. Firoiu, D. Towsley, and J. Kurose, "Modeling TCP Throughput: A Simple Model and its Empirical Validation", ACM SIGCOMM 1998, August 1998.
- [13] V. Tsaoussidis, H. Badr, X. Ge, K. Pentikousis, "Energy/Throughput Tradeoffs of TCP Error Control Strategies", *5th IEEE Symposium on Computers and Communications IEEE ISCC 2000*, July 2000.
- [14] V. Tsaoussidis and I. Matta, "Open issues on TCP for Mobile Computing", Journal of Wireless Communications and Mobile Computing, Wiley Academic Publishers, Issue 2, Vol. 2, February 2002.
- [15] V. Tsaoussidis and C. Zhang "TCP-Real: Receiver-oriented Congestion Control", The Journal of Computer Networks COMNET, Elsevier Science pp 477-497, Volume 40, Issue 4, November 2002.
- [16] K.Xu, Y.Tian and N.Ansari, "TCP-Jersey for wireless IP communications", IEEE JSAC, vol.22, no.4, pp.747-756, 2004/05.
- [17] Y.R. Yang, M.S. Kim and S.S. Lam, "Transient Behaviors of TCP-friendly Congestion Control Protocols", Proceedings of IEEE INFOCOM 2001, April 2001.
- [18] Y.R. Yang and S.S. Lam, "General AIMD Congestion Control", Proceedings of the 8th International Conference on Network Protocols", Osaka, Japan, November 2000.

Lefteris Mamatas is a Phd student in the Demokritos University of Thrace. He graduated in 2003 from the Dept of Electrical and Computer Engineering, Demokritos University of Thrace. His research interests are focused on transport protocols over wired/wireless networks.

Vassilis Tsaoussidis received a B.Sc in Applied Mathematics from Aristotle University, Greece; a Diploma in Statistics and Computer Science from the Hellenic Institute of Statistics; and a Ph.D in Computer Networks from Humboldt University, Berlin, Germany (1995). Vassilis held faculty positions in Rutgers University, New Brunswick, SUNY Stony Brook and Northeastern University, Boston. In May 2003, Vassilis joined the Department of Electrical and Computer Engineering of Demokritos University, Greece. His research interests lie in the area of transport/network protocols, i.e. their design aspects and performance evaluation. His COMNET (Computer Networks) Research Group includes 7 Ph.D students.

Vassilis is an editor for *IEEE Transactions in Mobile Computing* the Journal of *Computer Networks* the journal of *Wireless Communications and Mobile Computing* and the journal of *Mobile Multimedia*. He participates in several Technical Program Committees in his area of expertise, such as INFOCOM, GLOBECOM, ICCN, ISCC, EWCN, WLN, and several others. Vassilis graduated 2 Ph.D students in the United States: Adrian Lahanas, now a lecturer at the University of Cyprus; and Chi Zhang, now Asst. Professor at Florida International University.