

Measuring Transport Protocol Potential for Energy Efficiency

S. Kontogiannis, L. Mamatras, I. Psaras, and V. Tsaoussidis

Democritus University of Thrace,
Dept. Of Electrical & Computer Engineering, Xanthi, Greece
{skontog, emamatras, ipsaras, vtsaousi}@ee.duth.gr

Abstract. We investigate the energy-saving potential of transport protocols. We focus on the system-related aspect of energy. Do we have to damage or enhance system fairness in order to provide energy efficiency? We depart from defining protocol potential; we compare different transmission strategies and protocol mechanisms; and we report our results on the impact of each mechanism on system energy. We highlight our conclusion that protocol fairness appears to be a key factor for system energy efficiency.

1 Introduction

Energy consumption is becoming a crucial factor for wireless, ad-hoc and sensor networks, which affects system connectivity and lifetime. Standard TCP, originally designed for wired network infrastructure, does not cope with wireless conditions such as fading channels, shadowing effects and handoffs, which influence energy consumption.

Wireless network interface cards usually have four basic states of operation and each of these states has different power requirements. The most power-demanding states are the active states where transmission and reception of data take place. The standby/listen state, is the state where a network interface card is simply waiting. The extended period of idle state may lead to a sleep state, which is the least power-demanding state, where the radio subsystem of the wireless interface is turned off. Note that the transition mechanism itself is also energy consuming. Regardless of the states, their number and the frequency of transition, energy consumption is itself device-specific.

Due to the complexity of energy management and the fact that the state transition is device specific, each transmission or reception attempt by a higher-layer protocol does not necessarily correspond to a similar power transition. That is, we cannot accept *a priori* that the measured energy expenditure reflects the ability of a protocol to administer energy resources. Therefore, we distinguish protocol energy potential from actual device expenditure. The former approaches the latter when the sophistication of devices increases in a manner that all network layers operate in parallel states. Otherwise, if higher-layer protocol operation is suspended but the power module does not adjust, the protocol potential cannot translate into energy efficiency.

Several attempts have been made to measure the energy efficiency of transport protocols, (e.g. [10], [12]) as well as their potential for energy efficiency [14]. Energy efficiency is clearly device-specific while energy potential is not clearly defined. We attempt to define the latter, by introducing a corresponding index; we also attempt to measure actual expenditure, using specific device characteristics. We used Goodput in order to characterize protocol potential and an experimental extra energy expenditure index in order to characterize protocol energy performance.

Furthermore, we go beyond measuring energy potential within the confines of a single flow operation. We also investigate the system behavior of protocols attempting to address the question: “What are the design characteristics of transport protocols that impact system rather than single-flow energy efficiency”? In other words, what is the behavior of energy-efficient protocols within a multi-flow system? We noticed at this early stage of our investigation, some interesting results. While protocol Goodput is an important factor for energy efficiency (as we have also shown in [14]), protocol fairness seems to be another key factor for *system* energy efficiency.

The structure of this paper is the following: In section 2 we present the congestion control mechanisms that affect energy performance, according to distinct wireless conditions. In section 3 we present our proposed energy expenditure and energy potential metrics. In section 4 we present our scenario and evaluation plan and in section 5 we discuss the results.

2 Transmission Strategies and Network Conditions

The basic factor that determines the transmission strategies of the transport protocols is the window adjustments made by the congestion control algorithms. Different protocols employ distinct algorithms to control congestion. We focus on two basic categories of such algorithms. The first one considers the network as a black box and hence follow a blind procedure; the second one measures network conditions and adjust accordingly.

In the first category, in which most standard TCP versions belong, there are four widely available versions: Tahoe, Reno, New Reno and Sack. Tahoe is the most conservative version which includes Slow Start and Fast Retransmit [5], [8]. Reno is somewhat more aggressive due to its Fast Recovery mechanism. New Reno is even more aggressive when multiple drops occur within a single window of data, while Sack [9], the newest TCP version, is the most aggressive due to its selective acknowledgment strategy and its associated selective repeat mechanism.

The second category is represented by various standard (e.g. Vegas [2]) or experimental (e.g. Westwood [3], [11], Real [17], Jersey [7]) TCP protocols. We selected Vegas and Westwood for our experiments. TCP Vegas [2] congestion control is based on sample RTT measurements. The sender calculates throughput rate every RTT. This rate is compared to an expected rate, which is calculated based on what is measured as best RTT. TCP Westwood computes a sample

of Bandwidth by measuring and low pass filtering the rate of returning ACKs. TCP Westwood departs from the AIMD paradigm by proposing the additive increase adaptive decrease (AIAD) paradigm. No theoretical proof is given that AIAD converges to fairness.

In the context of transport protocol energy potential, we cannot isolate transmission strategy apart from distinctive error characteristics. We consider two major categories of errors, which are further classified into four different types. Each one of them calls for distinctive transmission tactics. We note that these types by no means traverse in detail the whole spectrum of distinct errors but are rather abstract. The first category, *congestion losses*, is separated into two types: *burst congestion losses* and *transient congestion losses*. During burst errors several consecutive transmitted packets are lost due to buffer overflow. By the term transient congestion errors, we characterize a situation where a small number of flows coexist in the same channel, causing in that way buffer overflowing sparsely, (e.g due to TCP synchronization). It is clear that both types of this category are associated with system's queuing delay. Under such conditions, we expect that the timeout mechanisms of the transport protocols have to be adjusted to accommodate the extra queuing delay. Furthermore, in case of burst congestion errors, the congestion window have to be drastically reduced, while transient errors may require smooth window adjustments.

The second category, *non-congestion losses*, includes the last two types of errors: *burst non-congestion errors* and *transient/random non-congestion errors*. Non-congestion losses, appear mostly in wireless/heterogeneous networks. Burst errors in the wireless portion of the network include handoffs, shadowing events, errors due to low SNR, etc. Under such conditions, data transmission would better be suspended until the communication channel recovers. This idea is implemented in TCP-Probing [13] where a probing mechanism gets aware of the situation and suspends data transmission for as long as the error persists.

Current TCP versions including these in our experiments, cannot distinguish those categories but mainly differentiate their mechanisms towards congestion losses. In other words, current TCP protocols are not suited for the distinct characteristics of wireless networks and thus an ideal protocol that can distinguish between those characteristics, could be much more energy efficient.

3 Measuring Energy Performance

In order to evaluate TCP performance over wireless networks and present useful directions in the context of energy consumption, we used traditional metrics, such as system *Goodput* and *Fairness Index* along with: *Extra Energy Expenditure* [10]. System Goodput is used to measure the overall system efficiency in bandwidth utilization and defined by (1). Fairness is measured by Fairness Index, derived from the formula given in (2).

$$Goodput = \frac{OriginalData}{ConnectionTime} \quad (1)$$

$$\mathcal{F.I.} = \frac{(\sum_{i=0}^n \|Throughput_i\|)^2}{n(\sum_{i=0}^n \|Throughput_i\|^2)} \quad (2)$$

The energy efficiency of a protocol is defined as the average number of successful transmissions per energy unit, which can also be computed as the average number of successes per transmission attempt as pointed out by Jones et al [6]. Energy expenditure or energy efficiency is a very important factor that has a major impact on wireless, battery-powered devices. However, apart from the overhead metric, there is no other metric in the literature that monitors the potential of a protocol for energy saving. Departing from that point and in order to capture the amount of *extra* energy expended, we introduce a new metric that was first presented in [16]. We call this new metric, *Extra Energy Expenditure (3E)*. The 3E metric, quantifies the extra effort expended without return in Goodput as well as the energy loss due to insufficient effort when aggressive transmission could have resulted in high Goodput. Three variables take place in this new metric. These are Throughput_{max}, Throughput and Goodput. The idea behind Throughput_{max} is that it captures the best possible data transmission that can be achieved under the given network conditions. The other two variables are the Throughput and Goodput metrics that monitor protocol performance. The 3E metric is given by the following formula:

$$\mathcal{E}\mathcal{E}\mathcal{E} = \alpha \frac{Thr - Goodput}{Thr_{max}} + b \frac{Thr_{max} - Thr}{Thr_{max}} \quad (3)$$

In order to explore the extra energy expenditure of a system of flows, we introduce system's 3E. System's 3E is equal to the sum of all competing flows extra energy expenditure:

$$\mathcal{E}\mathcal{E}\mathcal{E}_s = \alpha \frac{\sum_{i=1}^n (Thr_i - G_i)}{Thr_{max}} + b \frac{Thr_{max} - \sum_{i=1}^n Thr_i}{Thr_{max}} \quad (4)$$

It is clear that in all cases, Throughput_{max} ≥ Throughput ≥ Goodput. Extra Energy Expenditure (3E) takes into account the difference of achieved Throughput from maximum Throughput (Throughput_{max}) for the given channel conditions, as well as 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}. All available energy is consumed into efficient transmissions only when $Thr - Goodput = Overhead$ and $Thr = Thr_{max}$. That is, for an ideal TCP protocol that has an overhead of 40 Bytes in a 1024 Bytes TCP segment, EEE should be:

$$\mathcal{E}\mathcal{E}\mathcal{E} = \alpha \frac{0.04}{Thr_{max}} \quad (5)$$

In order for the 3E index to estimate the device specific extra energy expenditure, the value of α must be linked with the device transmission power: $\alpha = P_{Tx}(W)$ and the value of b must be linked with the device idle power: $b = P_{Idle}(W)$. In our experiments we normalized our α , and b parameters according to the the Lucent OriNOCO wireless device. We used the values of $\alpha = 1$ and $b = 0.45$.

4 Experimental Methodology

We have implemented a scenario, with two wireless nodes: The sender (node 0) and the receiver (node 1). The simulator used was the ns-2 network simulator and the topology an area 100x100 meters with a stable 100 meter distance between transmitter and receiver. The wireless link capacity is 2 Mbit. We used ns-2 energy model to simulate a specific device energy expenditure. The power values that were used for transmit, receive and idle states, where those of the Lucent OriNOCO wireless card. In our experiments we used a two-state error model for the process of packet errors, combined with the Bernoulli geometric distribution, to simulate the probability of packet drops. This is also known as the Gilbert channel model [4].

Our evaluation plan is consisted of two stages. At the first stage we modified the error-rate for a single flow scenario. We used different transport protocols in order to confirm the impact of different congestion control strategies energy potential and energy expenditure, for the one-flow system. At the second stage of this plan we modified the number of the flows for distinct error rates. Points of interest for us were those ones with similar Goodput performance but different fairness performance, utilizing different energy potential; or, those with worse Goodput performance which however were counterbalanced by fairness performance, resulting in better energy performance.

5 Results and Discussion

5.1 One-Flow Scenario Results

Energy expenditure or energy efficiency is a very important factor that has a major impact on wireless battery-powered devices. It is known already that a communication channel with low error rates should be utilized aggressively; when the error rate increases, a more conservative behavior yields better results.

Figure 1(a) compares the three standard TCP versions. TCP Reno seems to be more energy consuming when the packet error rate is greater than 15%. This is probably happening because Reno does not back-off to its initial congestion window (like Tahoe does) and neither does it recover with Fast Recovery (like New Reno).

Similarly, in figure 1(b) Vegas does not waste much energy when the error rate is low, while for higher error rates, Vegas behaves aggressively and under-achieves in terms of energy potential. More precisely, Vegas algorithm estimates accurately the available bandwidth at low error rates and thus presents better energy potential. However, when the error rate increases, Vegas estimator seems to estimate the available bandwidth without taking into consideration the persistent error conditions of the network. Under these conditions, Vegas false estimations are clearly outperformed by Tahoe's conservative strategy. Based on the above analysis, we confirm that a more aggressive behavior (Vegas) performs better under low error rate conditions, while the opposite might happen when

the error rate increases. Furthermore, Goodput proves one more time to be the most significant factor for TCP Energy Efficiency.

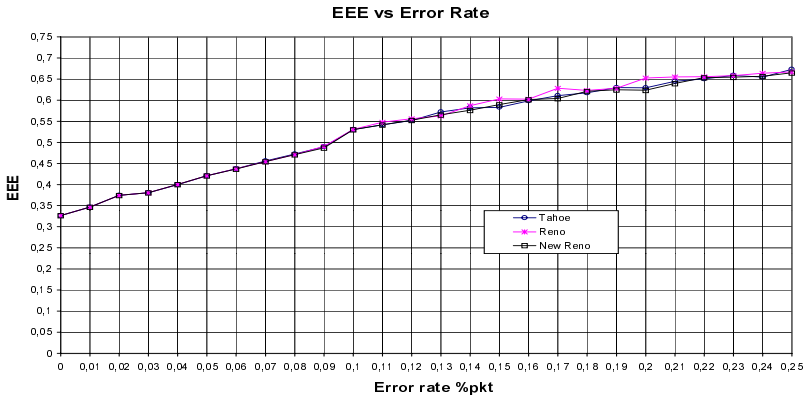
In the same scenario, Sack protocol neither appears energy efficient (figure 1(b)), nor does it achieve satisfactory Goodput performance. As Singh and Singh [12] stated, Sack “energy” performance suffers from extended timeouts and computational burden. In multi-drop situations where New Reno would timeout, Sack aggressively continues to retransmit packets. The aggressive re-transmissions, along with the computational burden and the extended timeouts are translated into extra energy expenditure.

TCP Westwood occasionally fails to adjust to the level of the available bandwidth, mainly at burst errors. Also it utilizes an adaptive policy appropriate for congestive losses and not for wireless errors. That is why its performance cannot overcome the performance of conservative TCP Tahoe both at random and burst error rates. However, as shown in figure 1(c), Westwood estimates available bandwidth more accurately at low error rates. For Westwood, when Goodput increases also energy potential increases.

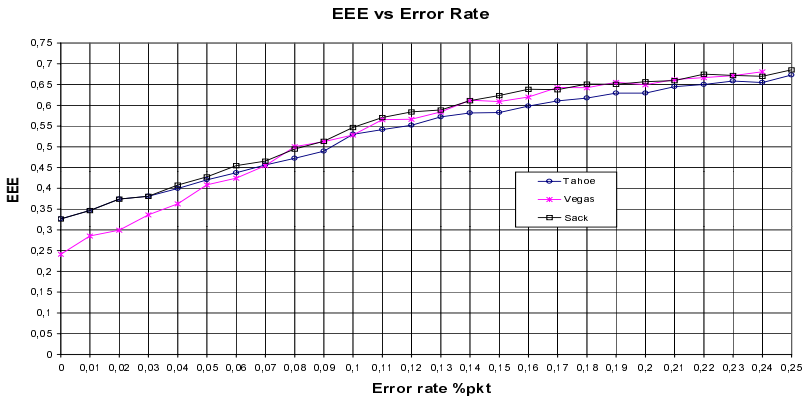
5.2 Multiple-Flow Scenario Results

We confirmed from previous one-flow scenario results that as Goodput performance increases, energy performance increases as well. The aforementioned conclusion is not quite accurate for a multi-flow system. In that case both Goodput and Fairness affect energy performance. In order to confirm the latter, we compare the behavior of two systems of flows. The first system utilizes TCP Vegas flows, while the second system utilizes TCP Tahoe flows. We focus on finding the points where both systems have the same amounts of Goodput but different values of Fairness index. According to figure 2(b), For a system of 5, 8 and 25 flows, Tahoe’s Goodput is equal or more than Vegas. On the other hand, Vegas is more fair than Tahoe for the 5, 8 and 25 flows systems. The impact of such behavior on energy performance is depicted in figure 2(c). Vegas increases its energy performance towards Tahoe, even if Tahoe performs equal or even better than Vegas. That is Tahoe shows increased amount of Goodput compared to Vegas. This confirms further our assertion that fairness does contribute to the system’s energy potential and energy performance. For a system of flows both Fairness and Goodput should be increased in order to improve protocol energy potential.

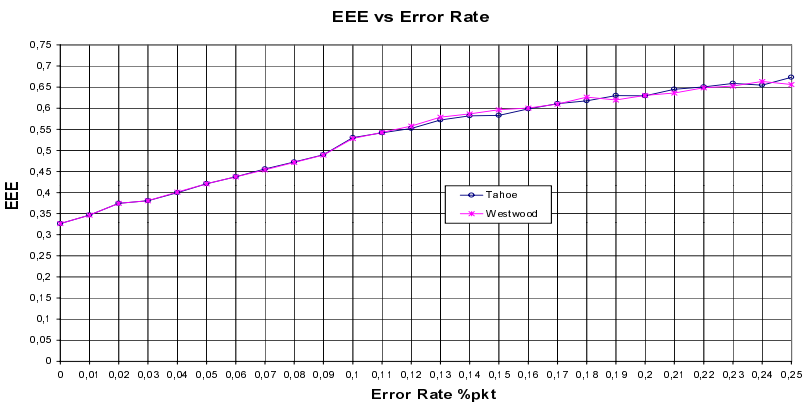
How far is fairness a dominant factor for energy efficiency? As we can see in figures 2(a) and 2(b), for a system of 3 flows, Vegas protocol is fair compared to Tahoe but performs poorly in terms of Goodput. The result for this system is that Tahoe has better energy potential. There is a point from where protocol energy performance is not affected by fairness, or, in other words, fairness impact on protocol energy performance is not the dominant factor. Moreover, as Goodput difference between two systems increases, fairness impact on protocol energy performance decreases. From a point and beyond, energy performance is mainly affected by Goodput performance. Systems Energy Expenditure accommodates the behavioral characteristics of systems energy potential. As depicted



(a) Tahoe, Reno and New Reno Extra Energy Expenditure vs Error rate.

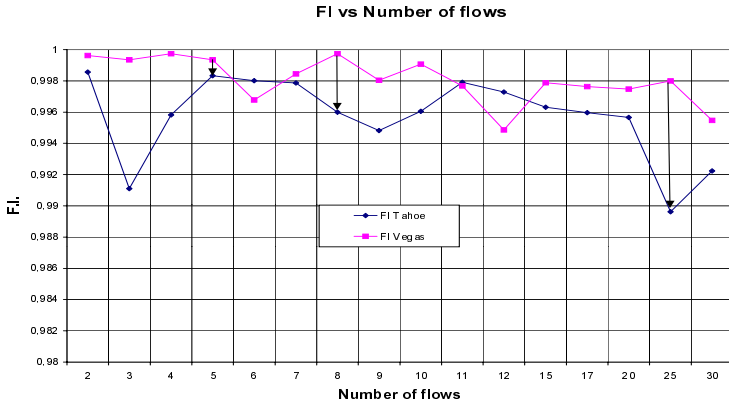


(b) Tahoe, Vegas and Sack Extra Energy Expenditure vs Error rate.

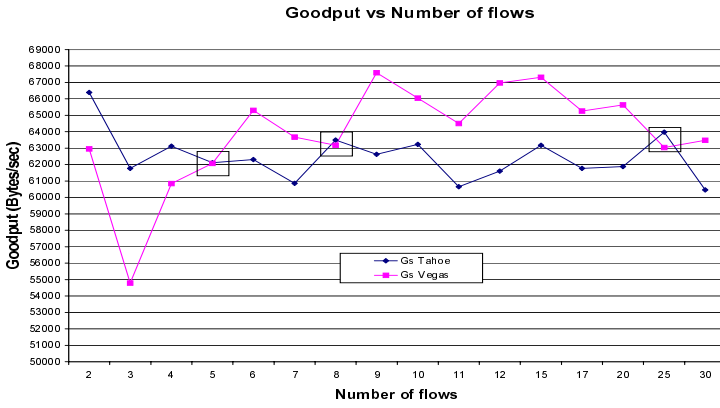


(c) Tahoe and Westwood Extra Energy Expenditure vs Error rate.

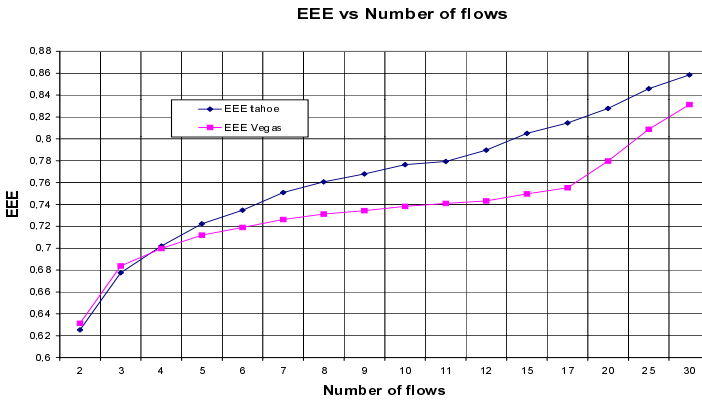
Fig. 1. TCP Protocols (a) Tahoe Reno and New Reno, (b) Tahoe Vegas and Sack (c) Tahoe and Westwood, Extra Energy Expenditure index vs Error rate



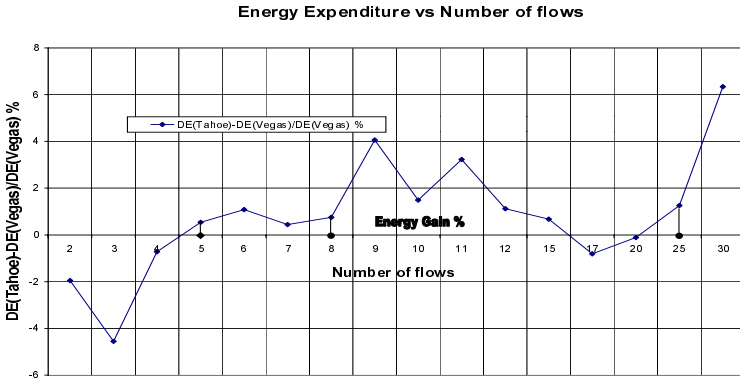
(a) Tahoe and Vegas Fairness Index vs Number of Flows - 0.1% Error rate.



(b) Tahoe and Vegas Goodput vs Number of Flows - 0.1% Error Rate.



(c) EEE vs Number of Flows - 0.1% Error Rate.



(d) Energy Gain vs Number of Flows - 0.1% Error Rate.

Fig. 2. TCP Protocols (a) F.I., (b) Goodput, (c) Extra Energy Expenditure and (d) Energy Gain

in figure 2(d), for the marked points 5, 8 and 25 of the Vegas flows system, fairness increases and Goodput decreases while system’s protocol energy potential increases. The actual energy gain of Tahoe versus Vegas due to the difference in fairness does not exceed 1% of the transmitter node total energy expenditure, while in general energy gain of Tahoe reaches 6%. However both protocols are far from reaching energy-conserving strategies. That is a new design can clearly reach much greater levels of energy efficiency.

6 Conclusions

Energy saving is not a property of one operation, layer, or protocol: Many design factors of different levels can contribute to achieve energy gains. We attempted to isolate energy gains due to transport protocol design characteristics. Since the energy-saving functionality of transport protocols may not be reflected in actual energy savings, due to device limitations, we introduced the notion of energy potential and linked it with the Extra Energy Expenditure (3E) index. We also adjusted this index to a specific device in order to establish a relation of *potential* with *real* expenditure. Using the aforementioned criteria, we evaluated the energy behavior of transport protocols. We report two important conclusions. First, we confirmed our previous assertion that high Goodput does contribute towards energy saving. Second, we observed that fairness is inherently correlated with system energy: when two systems achieve similar Goodput performance, the one that is more fair appears to be more energy-efficient as well.

References

1. I. Batsiolas and I. Nikolaidis. “Selective Idling: Experiments in Transport Layer Energy Conservation”. In: *Journal of Supercomputing, Special Issue on Design and Evaluation of Transport Services*, pages 101–114, 2001.
2. Lawrence S. Brakmo, Sean W. O’Malley, and Larry L. Peterson. TCP vegas: New techniques for congestion detection and avoidance. In *SIGCOMM*, pages 24–35, 1994.
3. C. Casetti, M. Gerla, S. Mascolo, M. Y. Sanadidi and R. Wang. “TCP Westwood: Bandwidth Estimation for Enhanced Transport over Wireless Links”. In: *Proc. of ACM Mobicom*, pages 287–297, July 2001.
4. E.N.Gilbert. “Capacity of a Burst-Noise Channel”. In: *The Bell System Technical Journal*, September 1960.
5. V. Jacobson. “Congestion avoidance and control”. *Proc. of ACM SIGCOMM’ 88*, August 1998.
6. Christine E. Jones, Krishna M. Sivalingam, Prathima Agrawal, and Jyh-Cheng Chen. “A Survey of Energy Efficient Network Protocols for Wireless Networks”. In: *Journal of Wireless Networks*, 7(4):343–358, 2001.
7. Y. Tian K. Xu and N. Ansari. “TCP-Jersey for Wireless IP Communications”. *IEEE Journal on Selected Areas in Communications*, vol. 22, No. 4:747–756, May 2004.
8. V. Paxson M. Allman and W. Stevens. “TCP Congestion Control”, RFC 2581, April 1999.
9. S. Floyd M. Mathis, J. Mahdavi and A. Romanow. “TCP Selective Acknowledgment Options”, IETF RFC 2018, 1996.
10. L. Mamas and V. Tsaoussidis. “Protocol Behavior: More Effort, More Gains?, PIMRC 2004”, September 2004, Barcelona, Spain.
11. Saverio Mascolo, Claudio Casetti, Mario Gerla, M. Y. Sanadidi, and Ren Wang. TCP Westwood: Bandwidth estimation for enhanced transport over wireless links. In *ACM Mobicom 2001*, 2001.
12. H. Singh and S. Singh. “Energy Consumption of TCP Reno, Newreno, and SACK in Multi-Hop Networks”. In *ACM SIGMETRICS 2002*, 2002.
13. V. Tsaoussidis and H. Badr. “TCP-Probing: Towards an Error Control Schema with Energy and Throughput Performance Gains”. In: *Proc. of the 8th IEEE Conference on Network Protocols ICNP*, 2000.
14. V. Tsaoussidis, H. Badr, X. Ge, and K. Pentikousis. “Energy / Throughput Trade-offs of TCP Error Control Strategies”. In: *Proc. of the 5th IEEE Symposium on Computers and Communications (ISCC)*, 2000.
15. V. Tsaoussidis and A. Lahanas. “Exploiting the adaptive properties of a probing device for TCP in heterogeneous networks”. In: *Computer Communications, Elsevier Science*, 2:40, Nov. 2002.
16. V. Tsaoussidis and I. Matta. “Open Issues on TCP for Mobile Computing”. *The Journal of Wireless Communications and Mobile Computing*, February 2002.
17. V. Tsaoussidis and C. Zhang. “The dynamics of responsiveness and smoothness in heterogeneous networks”. In: *IEEE Journal on Selected Areas in Communications (JSAC), Special issue on Mobile Computing and Networking*, March 2005.
18. M. Zorzi and R. Rao. “Energy Efficiency of TCP in a Local wireless Environment”. *Mobile Networks and Applications*, 6, Issue 3, ISSN:1383-469X:265–278, 2001.