Exploiting Energy-Saving Potential in Heterogeneous Networks

Lefteris Mamatas and Vassilis Tsaoussidis
(v1.0 released October 2006)

We investigate the energy-saving potential of transport protocols. We seek an answer to strategic issues of maximizing energy and bandwidth exploitation, without damaging the dynamics of multiple-flow equilibrium. We claim that (i) an energy-saving strategy of the transport level needs to be associated with some energy potential index which, unlike energy expenditure, is not device-specific and (ii) system-wise an energy-efficient system of flows is not always a better choice: we show that a less energy-efficient system may be more reliable in terms of packet multiplexing and, in turn, may reduce the probability that some flows may expend their energy with zero gain. We perform experiments using a real testbed and ns-2 based simulations.

Keywords: Energy Efficiency, Extra Energy Efficiency, Risk Index, UAR, Energy Potential

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.

We investigate energy efficiency from two perspectives:

(i) The energy-saving potential of the communication mechanism.
(ii) The risk of a flow to expend its energy for minor gains due to the multiplexing limitations. In particular, we investigate whether increased energy-saving capabilities may result in further unfair behavior. Since we associate energy expenditure not only with data transmission but also with time, unfair behavior translates into energy expenditure with minor performance gains.

Wireless network interface cards usually have four basic states of operation

Corresponding author. Email: emamatas@ee.duth.gr

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 *apriori* 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 collaborate ideally. Otherwise, if higher-layer protocol operation is suspended but the power module does not adjust, the protocol potential cannot translate into energy efficiency.

Since the network interface is a significant consumer of power, considerable research has been devoted to energy efficient design of the entire network protocol stack of wireless networks [1]. Several attempts have been made to measure the energy efficiency of transport protocols, (e.g. [2], [3], [4], [5], [6]), as well as their potential for energy efficiency [7]. 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.

Furthermore, we noticed at this stage of our investigation some interesting results. While protocol *Goodput* is an important factor for energy efficiency (as we have also shown in [7]), protocol fairness is another key factor for usability, which in turn determines the amount of flows that receive bad or zero service. In this context fairness also associates with energy: bad or zero service does not translate into minor or zero energy expenditure.

Consider a scenario where a system exhibits unfair behavior. Practically, some flows are favored while some others are not. We show experimentally that a system with increased energy efficiency does not guarantee better results for its users, but instead, the potential risk for a flow to receive bad or zero service is increased. We introduce an experimental metric, named *Risk Index* (RI), which captures this behavior.

The structure of this paper is the following: In section II we discuss protocol strategies. In section III we choose metrics for experimental analysis. Additionally, we introduce and discuss the *Energy Potential* and *Risk Index.* In section IV we detail our experimental methodology and evaluation plan. Finally, in sections V and VI we present our experimental results and we conclude the paper.

2 On Protocol Strategies

Energy cost due to communication relates with:

(i) The effort that the protocol expends (in terms of data transmission rate).
(ii) The amount of time required for the completion of communication.

In general, energy-consumption is the outcome of the transmission strategy that a transport protocol implements. An aggressive protocol, for example, may generate more overhead and hence expends some extra energy due to that overhead. By the same token, a conservative protocol may expend more energy due to unexploited opportunities for successful transmission. Clearly, a sophisticated (energy-wise) protocol should alternate aggressive and conservative strategies that minimize overhead and maximize efficiency. Such sophistication requires enhanced mechanisms for detecting network dynamics.

Additive Increase Multiplicative Decrease (AIMD) [8] allows for blind congestion control. According to AIMD, all senders keep increasing their transmission rate additively (i.e. the congestion window W increases by α packets per round-trip time (RTT)), until a packet loss. When congestion is taking place (i.e. there is a packet loss), a multiplicative decrease ratio is used to avoid a congestive collapse. So, the congestion window W decreases to βW upon congestion. The standard TCP uses the values $\alpha = 1$ and $\beta = 0.5$. TCP-friendly TCP(α, β) protocols parameterize the congestion window increase value α and decrease ratio β in order to trade responsiveness for smoothness. This tradeoff guarantees friendliness to traditional TCP.

Authors of [9] introduced a simple relationship for α and β:

$$\alpha = \frac{4(1 - \beta^2)}{3} \quad (1)$$

Based on experiments, they propose $\beta = 7/8$ as the appropriate value for the reduced the window (i.e. Less rapidly than TCP does). For $\beta = 7/8$, equation 1 gives an increase value $\alpha = 0.31$.

At a first glance, one may think that conservativeness and aggressiveness of the window adjustment strategy can be regulated by the increase/decrease parameters α and β. However, the adjustment of parameters α, β cannot really regulate some conservative or aggressive behavior. For example, a protocol with an increased α parameter is not always more aggressive than one with a smaller α value. An aggressive sender may trigger the timeout mechanism more times. If bursts of packets are being lost, the RTO mechanism can suspend transmission, which indicates a conservative behavior. We investigate when a protocol should be aggressive as well as the cost of this behavior in terms of energy-efficiency and fairness. Since the timeout may be a conflicting factor for scheduling an aggressive behavior[1], our adjustments of α and β are coupled with a small fixed timeout value. Practically, the trading of α for β parameter regulate the level of smoothness / responsiveness. Smoothness and responsiveness constitute a tradeoff [10]. Authors in [11] discuss the dynamics of this behavior.

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 [11]. Consider packet drops at the end of a congestion epoch; the window decreases by a factor of $(1 - \beta)$. 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 for the window size and the *ssthresh* to 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 the rate at which the sender reduces its transmission, but rather by its capability to 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.

3 Metrics for Evaluating Energy Performance

Energy dynamics in association with protocol strategy cannot be characterized accurately based only on traditional metrics. For example *Goodput* captures protocol performance but not protocol effort. *Goodput* is defined as:

$$Goodput = \frac{Original\_Data}{Connection\_Time} \tag{2}$$

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 corresponding data delivery.

[1]that is, an aggressive transmission may result in long periods of suspension

Therefore we complement this metric with the *EEE* metric. *Extra Energy Expenditure (EEE)* [2] attempts to capture the extra energy expended due to protocol operation - not just the expended energy. That is, a protocol may transmit when there are 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. *EEE* attempts 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. The index *EEE* 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 *EEE* takes values from 0 to 1, attempting to capture both distances.

$$EEE = a\frac{Throughput - Goodput}{Throughput_{max}} + b\frac{Throughput_{max} - Throughput}{Throughtput_{max}} \quad (3)$$

The first term of the *EEE* metric represents the overhead of network communication, normalized by resource availability (i.e., $Throughput_{max}$). Protocol overhead has a different impact on energy consumption depending each time on the particular device. Consequently, for every network card a different a value should be assigned. More precisely, the coefficient a is a function of the network card P_{tran} (transmission power) value and can be estimated experimentally.

The second term of the *EEE* metric captures the amount of available resources that have been exploited. When the available resources are exhausted, *Throughput* reaches $Throughput_{max}$. This term reflects energy consumption due to unexploited resources (e.g., time passes without any transmission). The b coefficient is a function of the network card P_{idle} (idle power) value. This term is bounded by the maximum energy consumption due to protocol inactivity. Consequently, the b coefficient is a function of P_{idle} (idle power) and not P_{sleep} (sleep power).

To summarize, the a and b parameters follow the behavior of a specific network device. In many cases, a sophisticated energy efficient protocol consumes more energy than it is designed to, due to lack of sophistication of the network device. However, the energy potential of a network protocol is not device dependant.

The *ideal EEE*, is the *EEE* produced by an ideal device. We assume that an ideal network device is energy efficient and sophisticated in the sense that its states correspond always to the states of the transport protocol (i.e. when the protocol suspend transmission the device remains on an idle state). Therefore, this device allows the transport protocol to operate on it's maximum energy efficiency. According our assumption, such a network card has a $\frac{P_{idle}}{P_{tran}}$ ratio of 0.3[1] and consumes the 30% of its energy in the idle state. Note that we did not find any network card with lower ratio. For example, according to [12], the Wavelan 2.4 GHz wireless network card have a $\frac{P_{idle}}{P_{tran}}$ ratio of 0.78. In this context, the *EEE* metric normalized with the parameters a=1 and b=0.3 behaves almost ideally.

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).

In the same context, *Fairness* derived from the formula given in [8] and defined as:

$$F(x) = \frac{\sum_0^{n-1}(Throughput_i)^2}{n\sum_0^{n-1}(Throughput_i^2)} \tag{4}$$

where $Throughput_i$ is the *Throughput* of the $i_t h$ flow and n the flow number.

Fairness captures overall multiplexing capabilities but does not indicate clearly whether flows exist that expend significant energy for zero return. Therefore, we complement this metric with the *Risk Index* defined as:

$$RiskIndex = \frac{NumberOfUnfavoredFlows}{TotalNumberOfFlows} \tag{5}$$

We regard as unfavored flows, the flows that have less *Goodput* than a specific threshold. In our case, the threshold is the 50% of the *average Goodput*.

Energy Potential can be defined as:

$$EP = 1 - EEE_{ideal} = 1-$$

[1]This assumption is subject of further work and may be explored theoretically.

$$\frac{Throughput - Goodput}{Throughput_{max}} + 0.3\frac{Throughput_{max} - Throughput}{Throughtput_{max}} \qquad (6)$$

An ideal energy efficient protocol should have *Energy Potential* 1 (which means zero extra energy expenditure).

For the sake of our analysis, and in particular, in order to be able to classify the cause of energy loss we specifically introduce the *UAR* index, defined as:

$$UAR = 1 - [k\frac{Throughput}{Throughput_{max}} + l\frac{Goodput}{Throughput}] \qquad (7)$$

where, typically, k=0.5 and l=0.5 (the k, l parameters may be adjusted according to a specific hardware). The *UAR* index ranges also from 0 to 1, expressing also a negative performance aspect.

Unexploited Available Resources (*UAR*) [2] captures how well did the protocol 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 *EEE* metric) but a major penalty.

UAR metric captures the behavior of the protocol in terms of available resources exploitation. A smooth protocol, which has a small α value, cannot exploit available bandwidth very fast. So, it has a high *UAR* value in the beginning. After some time, the protocol (due to the increased β value) is more aggressive. Consequently, the protocol may exploit available bandwidth efficiently further on.

The choice of metrics is very important for the experimental analysis. Each metric captures a different view of the protocol behavior. Additionally, the specific application's type calls for specific metrics. Table 1 summarizes the metrics used to highlight the different aspects of system performance.

4 Experimental Methodology

4.1 *Evaluation Plan*

We developed a real testbed in order to perform measurements and support our claims. Our testbed consists of a laptop, a desktop PC and a switch. We used *ACPI* (Advanced Configuration Performances Interface) to sample current voltage level, current drawn and available energy (in mAh) from

Goodput	Captures protocol performance.
Extra Energy Expenditure	Captures extra energy expended due to protocol operation.
Unexploited Available Resources	Captures how well the protocol exploits the windows of opportunities for successful transmission.
Fairness	Captures the how fair is the system to each participating flow.
Risk Index	Indicates whether flows exist that expend significant energy for minor gains.
Energy Potential	Indicates the energy potential of a protocol.

Table 1. Metrics for Evaluating Energy Performance

the laptop battery. *ACPI* is integrated in the Linux kernel and maps to the proc filesystem. *ACPI* takes measurements directly from the battery when an application accesses the corresponding file of the proc filesystem (*/proc/acpi/battery/BAT1/state*). Authors in [13] use similar methodology to measure energy consumption of "basic" application-level tasks, such as processing, input/output (disk, display, etc.) and communication (transmission and reception over the network).

We used an Acer Aspire 1692WLMI with Debian Linux OS, equipped with a Sanyo 65W Li-Ion battery, an Intel PRO/Wireless 2200BG 802.11b/g network card and a Broadcom BCM5700 network card for wired network.

We developed a tool for analyzing protocol performance which is focused on energy consumption. Our tool is based on Almost Tcp over UDP (*atou*) [14], an application-level implementation of TCP. We integrated our protocols and performance metrics into *atou* and evaluated the impact of different transport mechanisms on the energy consumption. Every experiment started with a full battery. We repeated our experiments several times in order to have statistically accurate results. Each experiment lasted 600 seconds, a time-period deemed appropriate to allow all protocols to demonstrate their potential. We used standard New-Reno TCP(1, 0.5), an extreme aggressive TCP(1.2, 1) with a small fixed timeout (50 ms) and a conservative TCP (0.3, 0.2) with a large fixed timeout (1 sec), in order to explore the limits of the energy consumption due to the network communication and to adjust our metrics. We used the adjusted metrics to evaluate 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 used the same testbed for the two-node scenario.

For a more extensive experimental analysis, we complemented our results by using ns-2 [15] based simulations. We used the same protocols and performance metrics. The network topology used is the typical single-bottleneck *dumbbell*,

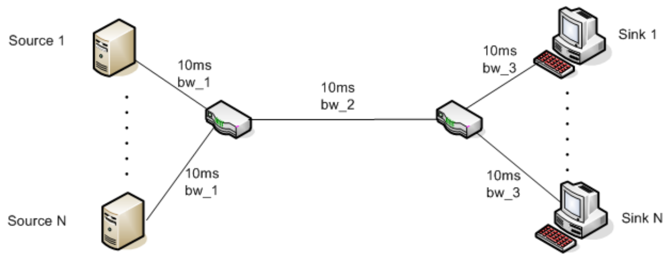

Figure 1. Dumbbell Topology

as shown in Figure 1. The bw_1 link is 10Mbps, the bw_2 link is 10Mbps and the bw_3 is 1Mbps. 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 packet-level errors with configurable packet error rate (PER). The simulation time was fixed at 120 seconds. Due to the deterministic nature of the experiments, statistical validity is not an issue.

5 Results and Discussion

5.1 *Energy efficiency results & adjustments of metrics*

In Figure 2(a) we observe the energy that three different transport mechanisms expend. The Idle curve depicts the energy consumption of our laptop battery when no communication takes place. When the TCP connection is on an Idle state (it does not actually transmit or receive any packet), the energy consumption slightly increases (Idle TCP curve). The actual communication-related energy consumption of a mechanism is therefore represented by the area between the corresponding energy-consumption curve and the Idle TCP one.

We assume that the aggressiveness of the transport mechanism ranges from the Idle TCP (which is zero) to the aggressive TCP and we adjust the *EEE* metric accordingly. We also assume that the *Throughput* of the aggressive TCP approaches the *maximum Throughput* that can be achieved under the specific network conditions. So, in the case of the aggressive TCP, the value of *EEE* should be close to 1 and the value of *UAR* approaches 0. In contrast, for the extremely conservative TCP, the *UAR* index should approach 1 and the *EEE* should approach the value $1.34/1.86 = 0.72$, where 1.34 is the average Idle TCP's power and 1.86 is the maximum power in the Figure 2(a).

Based on the equation 3 and on the results depicted from figures 2(a), 3(a), we get:

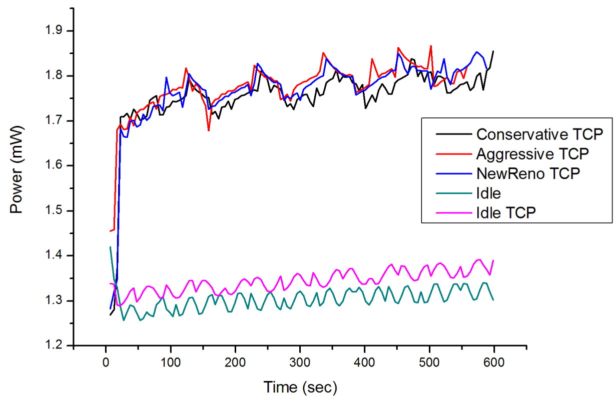

(a) Energy Consumption of Different Transport Mechanisms

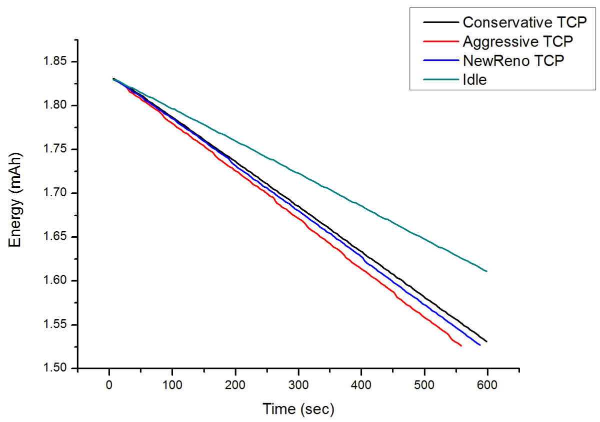

(b) Available Energy of the System

Figure 2. Energy-wise behavior of Different Transport Mechanisms

$$a = \frac{Throughput_{max}}{Throughput_{max} - Goodput_{aggressive}} = 5 \tag{8}$$

$$b = 0.72 \tag{9}$$

In Figure 2(b) we show the impact of the different transport mechanisms on the available energy of the system. In the case of aggressive TCP the battery is drained faster. The conservative TCP is more energy efficient than TCP NewReno. The aggressive TCP consumes 4 mAh more energy than the conservative TCP and the NewReno 3mAh.

The effort/gain dynamics of the system can be observed by Figure 3(a). The conservative TCP has less overhead, less *Throughput* but more *Goodput* than NewReno. Although it expends less effort, it achieves more gains. Consequently, NewReno expends more effort in this specific scenario. Similarly, the aggressive protocol expends significant effort (26% more) for only 8% gain.

In Figure 3(b) we plot the behavior of the three protocols in terms of *Extra Energy Expenditure* and *Unexploited Available Resources*. The EEE_1 curve represents the *ideal EEE* while the EEE_2 the actual one (normalized to the particular network device). We can observe that the aggressive protocol consumes more energy and instead the conservative protocol is the most energy efficient. We can also claim, based on the same figure, that the space for improvement is significant for all protocols.

The three protocols transmit data for about 600 seconds. The conservative TCP transmits 6.3GB with 174.2MB overhead. The aggressive TCP transmits 8.4GB with 1.5GB overhead and the NewReno TCP 6.1GB with 543MB. The system consumes about 300, 304 and 303 mAh energy, respectively.

In contrast to the conservative version of the protocol, the aggressive version expends extra effort for 1.4GB and consumes 4 mAh more energy in order to transmit 2.1GB of useful data. However, the conservative version would have required an extra minute of communication in order to transmit the same amount of useful data (2.1GB); the specific parameters of our experiments, would have caused more energy consumption than the 304 mAh of the aggressive version. However, we note that this conclusion may have been reverse had the network card idle state consumption been different (i.e. more conservative).

The NewReno TCP appears less energy efficient and is outperformed by conservative TCP in terms of *Goodput*. The additional effort expended by NewReno is not invested in performance gains. This result is quite interesting: 302MB less effort, which also corresponds to 3mAh less energy consumption achieves 66MB more *Goodput*.

5.2 *Evaluation of different transport mechanisms using testbed*

We evaluate three different versions of TCP: NewReno TCP, Responsive TCP and Smooth TCP. The Responsive TCP is the TCP(1.24, 0.25) and the Smooth TCP is the TCP(0.31, 0.875). We repeat the experiment 10 times in order to investigate the statistical accuracy of the results. In the following Figures (4(a) - 5(d)) we plot the average values for the 10 experiments. We

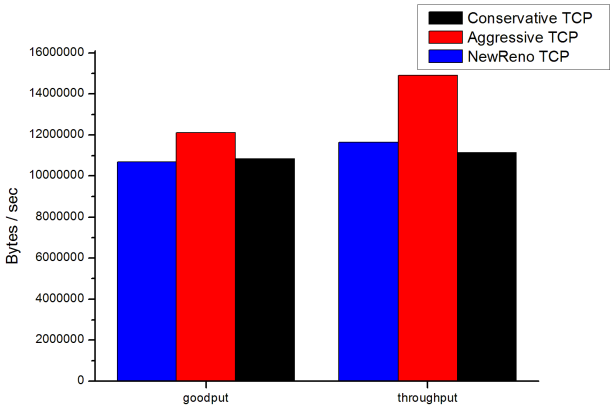

(a) Performance of Different Transport Mechanisms

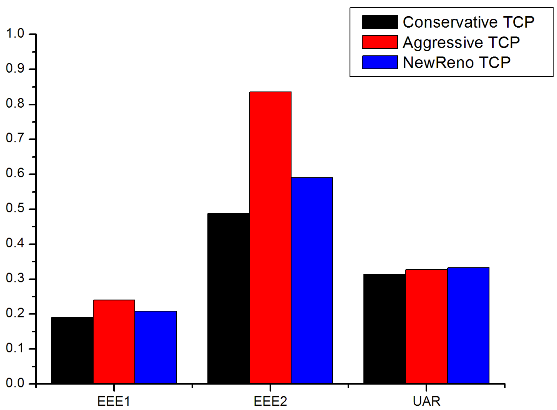

(b) EEE & UAR of Different Transport Mechanisms

Figure 3. Behavior of Different Transport Mechanisms

didn't observe significant deviation between the 10 experiments. For example, in the case of *Fairness*, the maximum deviation was 0.08, the minimum deviation was 0 and the average was 0.00899 (1.18 %).

According to Figures 4(a), 4(b) the aggressive behavior of NewReno TCP is not translated into increased *Goodput*. Compared with the Responsive TCP, Smooth TCP expends slightly more effort (Figure 4(a)) for a very significant return in *Goodput* (Figure 4(b)). This behavior is also captured by the *UAR* curve (Figure 5(c)). However, this extra effort is not distributed uniformly

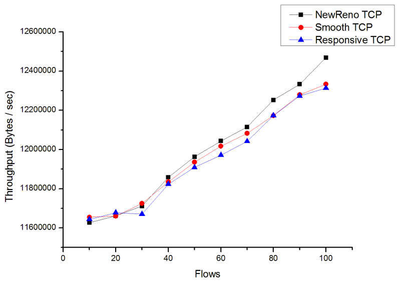

(a) Throughput of Different Transport Mechanisms

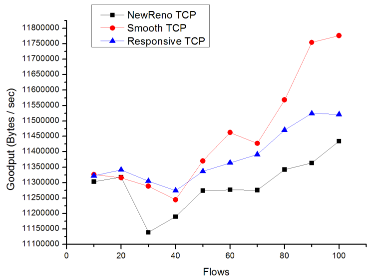

(b) Goodput of Different Transport Mechanisms

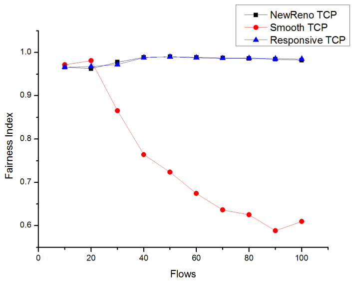

(c) Fairness of Different Transport Mechanisms

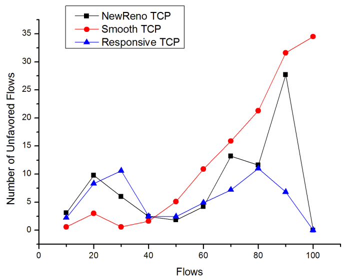

(d) Number of unfavored flows of Different Transport Mechanisms

Figure 4. Behavior of Different Transport Mechanisms

among participants (Figure 4(c)). In Figure 4(d) we plot the amount of flows that receive bad service due to the unfair system behavior. We defined as bad service the situation where a flow does not achieve at least 50% of the *average Goodput*. While NewReno and Responsive TCP exhibit similar behavior in terms of *Fairness*, the Smooth TCP is not fair (Figures 4(c), 4(d)).

According to the *Risk Index* (Figure 5(a)), the Smooth TCP appears unfair indeed. It causes several flows to receive bad service, which in turn causes great uncertainty to users of such system, especially when contention is high. There, the probability to expend significant energy for minor return is higher, even if the system is in general, more energy efficient.

Furthermore, Smooth TCP appears more energy efficient (Figure 5(b)). The situation uncovers a very interesting tradeoff. At least occasionally, in order to achieve better energy efficiency system-wise, we may let the *Risk Index* grow. In Figure 5(b) we show the *ideal EEE* curve. In Figure 5(d) we plot the behavior of the three protocols in terms of *Energy Potential*. We can see that, independently of the network device, the Smooth TCP has the best

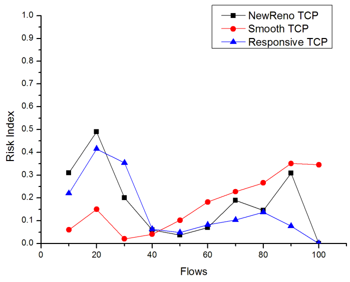

(a) Risk Index of Different Transport Mechanisms

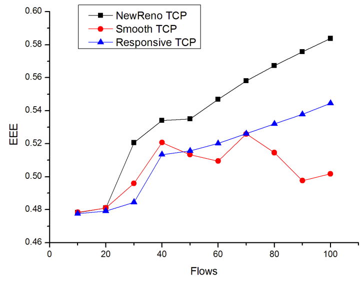

(b) EEE of Different Transport Mechanisms

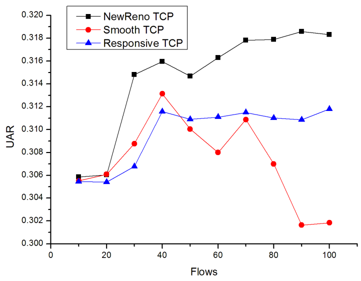

(c) UAR of Different Transport Mechanisms

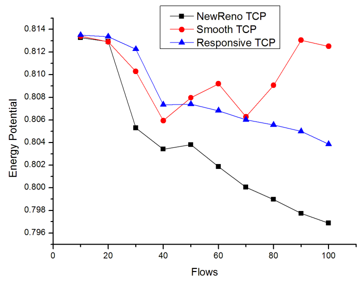

(d) EP of Different Transport Mechanisms

Figure 5. Behavior of Different Transport Mechanisms

Energy Potential in this particular case. Additionally, the NewReno TCP is outperformed by Responsive TCP in terms of energy efficiency (Figure 5(d)).

5.3 *Evaluation of different transport mechanisms using Simulations*

5.3.1 ***Low Error-Rate Favors Responsive Protocols.*** The first scenario simulates a heterogeneous environment with random transient errors increasing from 0.01 to 0.1 PER. We used 30 flows and a 10Mbps bottleneck, a relatively low-contention environment. The responsive protocol outperforms the smooth one in terms of energy efficiency (Figure 6(a)) and performance in terms of *Goodput* (Figure 6(d)) because it exploits resources better (Figure 6(b)). In this case, the responsive protocol deals with the transient errors sooner due to the setting of parameter α, without any negative impact on the system's fairness (Figure 6(c)).

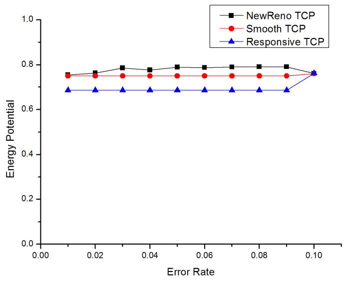

(a) EP & Low Error-Rate

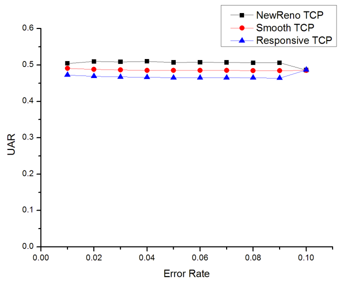

(b) UAR & Low Error-Rate

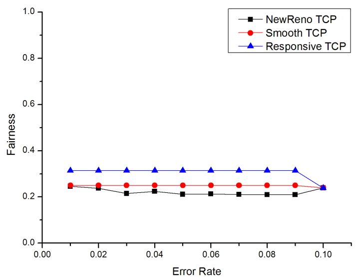

(c) Fairness & Low Error-Rate

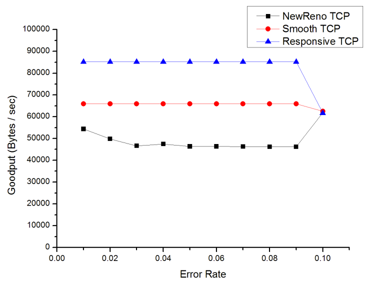

(d) Goodput & Low Error-Rate

Figure 6. Low Error-Rate Favors Responsive Protocols

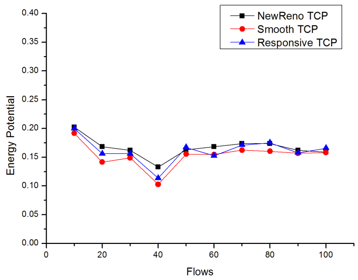

(a) EP & Effort/Gain Dynamics

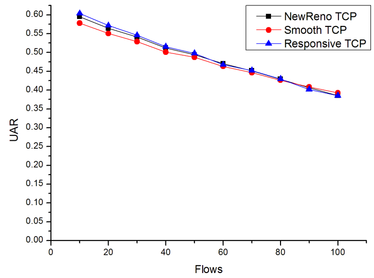

(b) UAR & Effort/Gain Dynamics

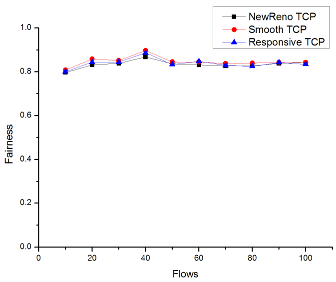

(c) Fairness & Effort/Gain Dynamics

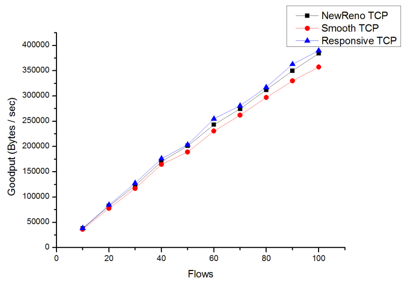

(d) Goodput & Effort/Gain Dynamics

Figure 7. A Macroscopic View of the Effort/Gain Dynamics

5.3.2 ***A Macroscopic View of the Effort/Gain Dynamics.*** In this scenario we used handoffs with duration 0.2 seconds in a 10Mbps bottleneck. We measured performance in terms of *Goodput*, ranging the number of flows from 10 to 100. We can observe that, better resource (Figure 7(d) and energy exploitation (Figure 7(a)) may have a positive impact on protocol goodput, although the reverse is also possible. See, for example the contrasting outcome with less and more effort, in figures 7(a), 7(b), 7(d) and 6(a), 6(b), 6(d), respectively. Although smooth TCP appears fair (Figure 7(c)), it is less energy efficient (Figure 7(a)) due to worse resources exploitation (Figures 7(d), 7(b)). A sophisticated protocol should have gains in terms of energy consumption and performance but without being unfair. Otherwise, some flows may drain their resources for minor data transmission.

5.3.3 ***Observations with Contention Decrease.*** The next scenario presented here intends to provide a framework for characterizing protocol behavior when bandwidth becomes available rapidly in heterogeneous networks. We measure *Energy Potential (EP)* - Figure 8(a), *Unexploited Available Resources*

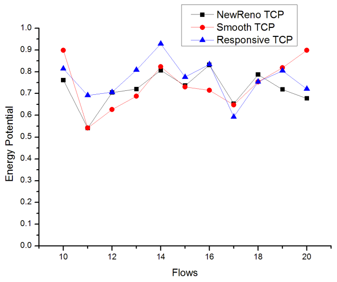

(a) EP & Contention Decrease

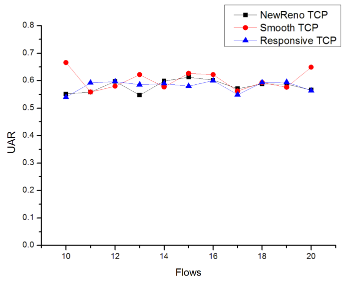

(b) UAR & Contention Decrease

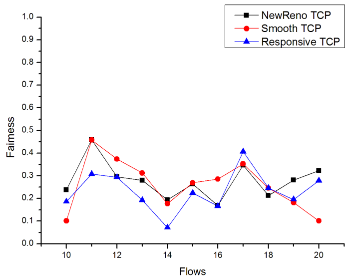

(c) Fairness & Contention Decrease

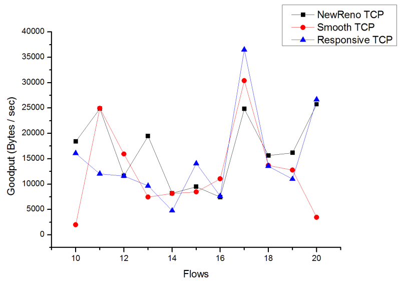

(d) Goodput & Contention Decrease

Figure 8. Observations with Contention Decrease

Index (UAR) - Figure 8(b), *Fairness* - Figure 8(c) and *Goodput* - Figure 8(d) for a range of flows from 10 to 20. We used a 0.2 PER. All flows are entering in the system within the first two seconds. For the rest 118 seconds we have a graduated contention decrease, starting from 10 flows and repeating the experiment for 11 to 20 flows. At each stage we reduce the number of flows to half every Decrease_Step seconds, where Decrease_Step, is the step needed, in order for the last flow to exit at the 120th second.

The small value of parameter α of Smooth TCP leads to slow resource exploitation (Figure 8(b)) without any gains in terms of energy efficiency (Figure 8(a)). On the other hand, Responsive TCP consumes less energy (Figure 8(a)) but exploits resources (Figure 8(d) in an unfair manner (Figure 8(c)).

5.3.4 ***Error-Rate Increase Cancels Responsive TCP's Advantages.*** In the following scenario, we used 30 flows, a 10Mbps bottleneck and a variable error-rate from 0.01 to 0.4 PER. During small error rates the responsive protocol has better return for its effort, however, when error-rate exceeds 0.1, these advantages are canceled (see Figures 9(a), 9(b), 9(c)). In Figures 9(a),

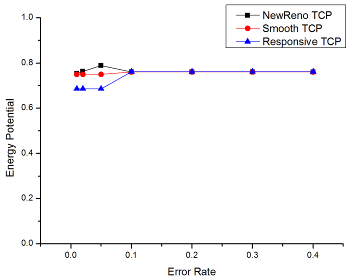

(a) EP & Error-Rate

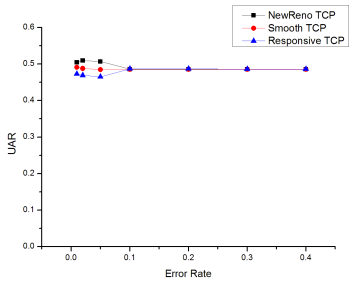

(b) UAR & Error-Rate

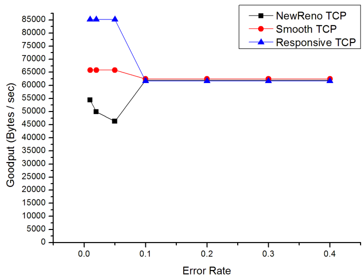

(c) Goodput & Error-Rate

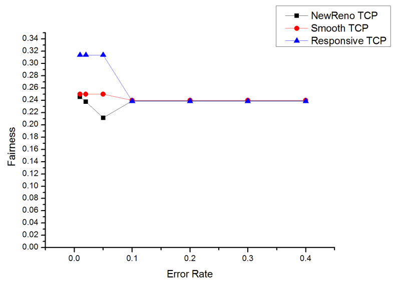

(d) Fairness & Error-Rate

Figure 9. Error-Rate Increase Cancels Responsive TCP's Advantages

9(b), 9(c), 9(d), we summarize the difference in *Energy Potential*, *Unexploited Available Resources Index*, *Goodput* and *Fairness*.

We can see that the responsive protocol is favored at the beginning. After a certain point, which is relevant to the specifics of the experiment (which in our case is 0.1), the smooth protocol may even become more efficient (in goodput) and fair, while it expends less extra energy.

6 Conclusions

We explored the energy-saving potential of different transport protocols using a real testbed. We introduced two new metrics, the *Energy Potential* and *Risk Index*. *Energy Potential* is a device-independent metric which captures the energy-saving potential of a protocol. *Risk Index* refers to a system's behavior and captures the potential risk for a flow to expend its energy for minor *Goodput* due to the multiplexing limitations.

We confirm experimentally that, in general, smoothness and responsiveness constitute a tradeoff; however, we show that this tradeoff does not correspond to a conservative/aggressive behavior. Energy-wise, existing protocol tactics cannot always be justified; our results suggest that an adaptive congestion control algorithm is needed to integrate the dynamics of heterogeneous networks into protocol behavior.

Acknowledgment

This work was supported by the project PENED 03, funded by European Commission and the G.S.R.T. of Greek Ministry of Development.

References

[1] Jones, C., Sivalingam, K., Agrawal, P. and Chen, J., July 2001, A Survey of Energy Efficient Network Protocols for Wireless Networks. *Wireless Networks,* **7**(4), 343-358
[2] Mamatas, L., Tsaoussidis, V., 2004, Protocol Behavior: More Effort, More Gains? *Proceedings of PIMRC 2004,* Barcelona, Spain
[3] Singh, H., Singh, S., 2002, Energy Consumption of TCP Reno, Newreno, and SACK in Multi-Hop Networks. *Proceedings of ACM SIGMETRICS 2002,* Marina del Rey, California, USA
[4] Wang, B., Singh, S., March 2004, Computational Energy Cost of TCP. *Proceedings of the INFOCOM 2004,* Hong Kong
[5] Agrawal, S., Singh, S., 2001, An Experimental Study of TCPs Energy Consumption over a Wireless Link. *Proceedings of the 4th European Personal Mobile Communications Conference)*
[6] Anastasi, G., Conti, M., Gregori, E. and Passarella, A., 2002, A Power Saving Architecture for Web Access from Mobile Computers. *Proceedings of the 2nd International IFIP Networking Conference (Networking 2002),* **2345**, 240-251
[7] Tsaoussidis, V., Badr, H., Ge, X. and Pentikousis, K., July 2000, Energy / Throughput Tradeoffs of TCP Error Control Strategies. *Proceedings of the 5th IEEE Symposium on Computers and Communications (ISCC),* Antibes-Juan les Pins, France
[8] Chiu, D., Jain, R., June 1989, Analysis of the Increase/Decrease Algorithms for Congestion Avoidance in Computer Networks. *Journal of Computer Networks and ISDN,* **17**(1), 1-14
[9] Yang, Y., Lam, S., November 2000, General AIMD Congestion Control. *Proceedings of the IEEE International Conference on Network Protocols,* Osaka, Japan
[10] Zhang, C., Tsaoussidis, V., July 2002, The Interrelation of TCP Responsiveness and Smoothness. *Proceedings of the 7th IEEE ISCC*
[11] Tsaoussidis, V., Zhang, C., 2005, The Dynamics of Responsiveness and Smoothness in Heterogeneous Networks. *IEEE Journal on Selected Areas in Communications,* **23**(6), 1178-1189
[12] Yong, W., Chandra, S. and Bhandarkar, S., 2004, A Statistical Prediction-based Scheme for Energy-aware Multimedia Data Streaming. *Wireless Communications and Networking Conference (IEEE WCNC 2004),* 2053-2057
[13] Margi, C.B., Obraczka, K., Manduchi, R., June 2005, Characterizing system level energy consumption in mobile computing platforms. *Proceedings of the International Conference on Wireless Networks, Communications and Mobile Computing,* **2**, 1142-1147
[14] Dunigan, T., Almost TCP over UDP (ATOU). *Available: http://www.csm.ornl.gov/ dunigan/netperf/atou.html*
[15] The Network Simulator - NS-2. *Web Page: http://www.isi.edu/nsnam/ns/*