ecMTCP: An Energy-Aware Congestion Control Algorithm for Multipath TCP

Tuan Anh Le, Choong Seon Hong, *Member, IEEE,* Md. Abdur Razzaque, *Member, IEEE,* Sungwon Lee, *Member, IEEE*, and Heeyoung Jung, *Member, IEEE*

***Abstract*—In this paper, we develop an energy-aware congestion control algorithm for multipath TCP, called ecMTCP. ecMTCP moves traffic from the most congested paths to the more lightly loaded paths, as well as from higher energy cost paths to the lower ones, thus achieving load-balancing and energy-savings. Our simulation results show that ecMTCP can achieve greater energy-savings compared to MPTCP, and preserve fairness to regular TCP.**

***Index Terms*—Multipath TCP, energy-saving, energy-aware congestion control, end-to-end energy cost measurement model.**

I. INTRODUCTION

RECENT portable devices such as smart phones and tablet PCs have multiple radio interfaces. If multipath transport protocols were used, the performance, mobility, and energy-savings of these devices would likely be improved. A multipath TCP connection can create multiple simultaneous sub-flows among the end hosts, where each sub-flow maintains the ability to send data packets over a path [1]. In such environments, devising a congestion control algorithm involving a number of issues, such as the rate control of the path, traffic sharing amongst the paths, and energy awareness, is a challenging task.

Several solutions for the multipath congestion control problem have been suggested in the literature. For example, mTCP [2] stripes data packets across multiple paths, and detects shared bottlenecks to alleviate unfair sharing. Coupled multipath-aware congestion control (MPTCP) [3] [4] views a set of resources on all paths as a single resource, as suggested by the resource pooling principle [5]. The coupled congestion control moves more data traffic load onto the less congested paths as a load-balancing mechanism. However, the problem of designing an efficient multipath congestion control algorithm that is simultaneously energy-aware and capable of handling load-balancing has not yet been addressed.

In this work, we are motivated by the fact that the design of energy-aware transmission protocols is currently an urgent need [6]. This is mainly due to the following two reasons: the increased use of the Internet by energy-critical portable devices, and the IEEE P802.3az task force [7], which has also recently standardized guidelines for saving energy on Ethernet networks. Here we develop an energy-aware congestion control algorithm for multipath TCP (ecMTCP), in which the rate control is based on a traffic sharing policy amongst the paths, and which is driven by their energy costs and traffic loads. The evaluation results show that ecMTCP can achieve energy-savings and fairness.

Manuscript received August 31, 2011. The associate editor coordinating the review of this letter and approving it for publication was H. Yomo.

This research was supported by the IT R&D program of KCA (10913-05004: Study on Architecture of Future Internet to Support Mobile Environments and Network Diversity).

T. A. Le, C. S. Hong (corresponding author), and S. Lee are with the Dept. of Computer Engineering, College of Electronics and Information, Kyung Hee University, Korea (e-mail: cshong@khu.ac.kr).

M. A. Razzaque is with the Dept. of Computer Science and Engineering, University of Dhaka, Bangladesh.

H. Y. Jung is with ETRI, Korea.

Digital Object Identifier 10.1109/LCOMM.2011.120211.111818

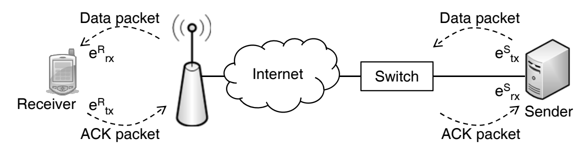

Fig. 1. The energy cost measurement model between two-end hosts.

The rest of this paper is organized as follows. Section II presents an energy cost measurement model. In Section III, we describe the details of the ecMTCP design. Section IV presents the simulation results. We conclude our work in Section V.

II. END-TO-END ENERGY COST MEASUREMENT MODEL

TCP congestion control can play a significant role in regulating the network load and in determining the network interface load. We propose an energy cost measurement model between the two-end hosts, as shown in Fig. 1. The sender measures the energy costs e^S_{tx} (μJ/byte) for transmitting a data packet, and e^S_{rx} for receiving an ACK at its network interface. Similarly, at the receiver side, the costs for the data and ACK packets are e^R_{rx} and e^R_{tx}, respectively. The energy cost at the receiver is piggybacked to the sender periodically via the ACK option. The sender calculates the sum of the end-to-end (e2e) energy cost of the successful transmission of one data packet (μJ/byte) on path s as

$$e_s = e^S_{tx} + e^R_{rx} + (e^R_{tx} + e^S_{rx})ACK_s/(n\_ACK_s DATA_s),$$

where $DATA_s$ and ACK_s are the data packet and ACK packet sizes, respectively, on path s. n_ACK_s denotes the number of data packets acknowledged by one ACK packet on that path. The short-term variations of the instantaneous e2e energy cost e_s are then smoothed out by using the exponentially weighted moving average formulae, with weight factor $\gamma = 0.125$, as follows: $e_s(t) = (1-\gamma)e_s(t-1) + \gamma e_s$.

III. ECMTCP DESIGN

Designing ecMTCP necessitates the following aims: (i) *improve throughput*, (ii) *fairness* to both regular single-path

TCP and MPTCP flows, (iii) *load-balancing* [3], and (iv) *energy-efficiency*, such that - the algorithm should exploit potential energy-savings when the e2e energy costs between paths are diverse, while balancing the loads amongst the paths.

Now, we adopt MPTCP, proposed in [4, Eqn. (3)]. MPTCP was designed in order to satisfy the first three requirements. In this work, we focus on the energy-efficiency requirement, which is satisfied by moving traffic off of the higher energy cost paths to the lower paths, while guaranteeing load-balancing. This leads to the establishment of a relationship between the linked windows [4] and the e2e energy costs, known as a *linking function* (lf_s). Finally, ecMTCP is implemented as follows: Whenever a sub-flow on path s receives a positive ACK, it increases its window (w_s) such that

$$w_s(t+1) \leftarrow w_s(t) + \min\left\{\delta lf_s, \frac{1}{w_s(t)}, \delta' \frac{w_s^{1-k}(t)}{w^{2-k}(t)}\right\}, \quad (1)$$

$$lf_s = \left(\frac{w_s^{1-k}(t)}{w^{2-k}(t)}\right)\left(N\frac{1/e_s(t)}{e(t)}\right)^{\theta+k},$$

whereas whenever duplicate ACKs are received indicating a packet loss, it decreases w_s such that

$$w_s(t+1) \leftarrow w_s(t) - \beta w_s(t), \quad (2)$$

where $w(t) = \sum_s w_s(t)$, and $e(t) = \sum_s 1/e_s(t)$. N denotes the number of paths used during a multipath TCP session. k determines the trade-off between fluctuation and load-balancing, $0 \leq k \leq 2$ [4]. The protocol's energy efficiency is adjusted by both $\theta \geq 0$ and k. $\beta = 0.5$ denotes the decrease factor. To make ecMTCP completely backward-compatible with MPTCP when the energy costs between the paths are identical or not in support of the measurement model described in Section II, the k value must be set to 1 [3] [8]. Parameters δ and δ' determine the congestion window increment recompense for multiple paths with divergent round trip times (RTTs). Without them, the first and third requirements are not satisfied since the congestion windows of the sub-flows (having longer RTT values) on their least congested paths are increased more slowly than expected. To ensure fairness to both regular TCP and MPTCP flows sharing a common bottleneck link, ecMTCP's window increment does not exceed that of either the regular TCP flow or MPTCP sub-flow due to the use of the two last arguments of $\min\{.\}$.

The first two requirements are satisfied when the total throughput of a multipath TCP flow equals the maximum throughput over all paths s that a regular TCP flow could achieve [4] [8], as

$$\sum_s \hat{w}_s/RTT_s = \max_s\left\{\hat{w}_s^{TCP}/RTT_s\right\}, \quad (3)$$

where $\hat{w}_s$ and $\hat{w}_s^{TCP}$ denote the equilibrium value of $w_s(t)$ and a regular TCP flow's congestion window on path s, respectively. The additive increase function in (1) for each sub-flow is interlinked by using the ratio of its window to the sum (the left term of lf_s) [4] [8], satisfying the third requirement. For the energy-efficiency requirement, we propose an energy-aware control mechanism, where the congestion window increase is inversely proportional to the energy cost on each path (the right term of lf_s). The path with higher energy cost will be less aggressive in increasing the congestion window than the one incurring less cost.

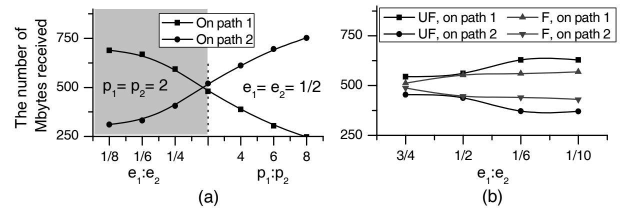

Fig. 2. (a) Trade-off between load-balancing and energy-saving; (b) trade-off between fairness and energy-saving.

By neglecting δlf_s in (1), ecMTCP becomes MPTCP. Hence, $\delta' = \hat{w}^{2-k}\max_s\left\{\hat{w}_s^k/RTT_s^2\right\}/(\sum_s \hat{w}_s/RTT_s)^2$ [4], where $\hat{w}$ denotes the equilibrium value of $w(t)$. Similarly, by neglecting $\delta' w_s^{1-k}(t)/w^{2-k}(t)$ in (1) to calculate δ, the fluid model of an ecMTCP sub-flow on path s corresponding to (1) and (2) is expressed as

$$\frac{d}{dt}w_s(t) = \min\left\{\delta lf_s \frac{w_s(t)}{RTT_s}, \frac{1}{RTT_s}\right\}(1-p_s) - \frac{\beta w_s^2(t)}{RTT_s}p_s. \quad (4)$$

We assume that the packet drop probability on path s (p_s) is small. Therefore, the fixed point equation of (4) gives

$$\min\left\{\delta\frac{\hat{w}_s^{1-k}}{\hat{w}^{2-k}}\left(\frac{N}{e_s(t)e(t)}\right)^{\theta+k}, \frac{1}{\hat{w}_s}\right\} = \beta\hat{w}_s p_s. \quad (5)$$

The packet loss rate of a regular TCP flow in the steady state using path s could be

$$p_s = 1/\beta\left(\hat{w}_s^{TCP}\right)^2. \quad (6)$$

By substituting (6) into (5), we obtain

$$\hat{w}_s^{TCP} = \max\left\{\left(\frac{\hat{w}^{2-k}\hat{w}_s^k}{\delta}\right)^{1/2}\left(\frac{e_s(t)e(t)}{N}\right)^{(\theta+k)/2}, \hat{w}_s\right\}. \quad (7)$$

δ is determined by solving equations (3) and (7) as

$$\delta = \hat{w}^{2-k}\left(\frac{e(t)}{N}\right)^{\theta+k}\frac{\max_s\left\{\hat{w}_s^k e_s^{\theta+k}(t)/RTT_s^2\right\}}{\left(\sum_s \hat{w}_s/RTT_s\right)^2}.$$

We consider a two-path ecMTCP flow. Neglecting $1/\hat{w}_s$ in (5), the ratio of the congestion windows between paths is

$$\hat{w}_1/\hat{w}_2 = \left(p_2 e_2^{\theta+k}(t)/p_1 e_1^{\theta+k}(t)\right)^{1/k}. \quad (8)$$

For $k = 1$, as mentioned above, we choose $\theta = 0$ such that the ecMTCP reaction to the e2e energy cost on each path is equivalent to the packet drop on that path.

We define the energy cost ratio as the ratio of the energy cost of path 1 to that of path 2. From (8), the congestion window sizes of the sub-flows are identical, as $p_1 : p_2 = e_2 : e_1$. Fig. 2(a) shows how ecMTCP (without the fairness consideration) trades-off between load-balancing and energy-saving by transferring a 1GB file. Simulation parameters are set as shown in Fig. 3(b). The dotted line indicates the balance of two paths at $e_1 : e_2 = 1/2$ and $p_1 : p_2 = 2$. In the shaded region, given $p_1 : p_2 = 2$, more traffic is sent on the least energy cost path (path 1), as the energy cost ratio is more diverse. Alternatively,

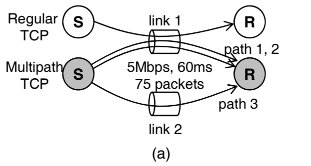

Fig. 3. (a) A three-path ecMTCP flow vs. a regular TCP flow; (b) a two-path ecMTCP/MPTCP flow.

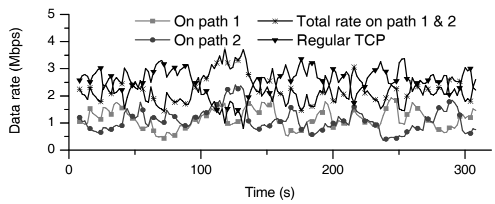

Fig. 4. Fairness with $e_1 : e_3 = e_2 : e_3 = 1/4$ and a 0.01% packet loss rate on links 1 and 2. The data rate was sampled once every 2 seconds.

in the non-shaded region, given $e_1 : e_2 = 1/2$, the traffic is dominated by the load-balancing mechanism, hence more traffic is moved to the least congested path (path 2), as $p_1 : p_2$ is inversely proportional to $e_1 : e_2$ in the previous case. In contrast to the trade-off regions, ecMTCP moves more traffic aggressively from a congested and higher energy cost path to a better path. Since traffic aggregation on a path is bounded by $\min\{.\}$ in (1) to preserve fairness, this makes another trade-off between fairness and energy-efficiency. Fig. 2(b) shows that the gap between two paths in the unfairness condition (UF) (without the third argument of (1)) is larger, as $e_1 : e_2$ is more diverse. However, the gap between paths in the fairness condition (F) remains almost parallel, as $e_1 : e_2 \leq 1/2$.

IV. Simulation Results

In this section, we evaluate ecMTCP's energy efficiency and fairness using NS-2 [9] with the SACK option. Two scenarios shown in Fig. 3 with a 1000-byte data packet and droptail are used in the evaluations. In the experiments, the energy costs on the paths are unchanged during the multipath TCP session.

We evaluate ecMTCP fairness in a typical configuration, as shown in Fig. 3(a), where two ecMTCP sub-flows (on low energy cost paths 1 and 2) compete against a regular TCP flow in link 1, and the third sub-flow runs on high energy cost path 3. Fig. 4 shows that the two ecMTCP sub-flows on paths 1 and 2 do not get more bandwidth than the regular TCP flow. Therefore, ecMTCP can fairly share with regular TCP.

The energy-efficiency is determined by comparing the energy consumption of our ecMTCP with MPTCP when transferring a 1GB file in the simulation scenario shown in Fig 3(b). The simulation setup environment is given in Table I, representing a variety of network conditions. We have taken into account energy costs for retransmitted data packets as well. Fig. 5 shows greater variation in the e2e energy cost ratio results in more energy-savings for all experiments. The energy-savings are higher when more throughput can be passed on the least energy cost path, path 1 (e.g., larger link capacity at the bottleneck link, shorter RTT, or lower packet loss rate). Therefore, the achievable energy-savings depend on

TABLE I

Simulation Setup Environment.

Exp. ID	**Link 1**	**Link 2**	**Packet loss rates (%)**
EQ	5Mbps, 60ms	5Mbps, 60ms	p_1 = 0.01, p_2 = 0.01
CL	5Mbps, 60ms	10Mbps, 60ms	p_1 = 0.01, p_2 = 0.01
CG	10Mbps, 60ms	5Mbps, 60ms	p_1 = 0.01, p_2 = 0.01
DL	5Mbps, 60ms	5Mbps, 120ms	p_1 = 0.01, p_2 = 0.01
DG	5Mbps, 120ms	5Mbps, 60ms	p_1 = 0.01, p_2 = 0.01
PL	5Mbps, 60ms	5Mbps, 60ms	p_1 = 0.01, p_2 = 0.02
PG	5Mbps, 60ms	5Mbps, 60ms	p_1 = 0.02, p_2 = 0.01
HE1	5Mbps, 60ms	5Mbps, 240ms	p_1 = 0.02, p_2 = 0.01
HE2	5Mbps, 240ms	5Mbps, 60ms	p_1 = 0.01, p_2 = 0.02

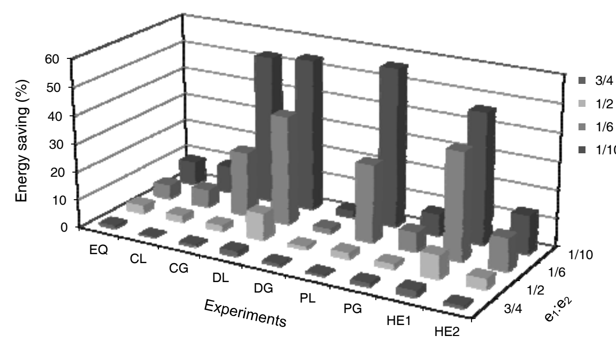

Fig. 5. Energy-savings vs. network conditions and energy cost ratios.

the throughput on the least energy cost path, which is affected by the packet loss rate, RTT, network congestion, link capacity, or any combination of these factors.

V. Conclusion

Our ecMTCP can support the exploitation of potential energy-savings between two-end hosts when using both the e2e energy cost measurement model and the energy-aware multipath congestion control algorithm. The energy-savings and fairness of ecMTCP are validated through simulations.

References

[1] A. Ford, C. Raiciu, and M. Handley, "TCP extensions for multipath operation with multiple addresses," IETF Internet Draft, July 2011.

[2] M. Zhang, J. Lai, A. Krishnamurthy, L. Peterson, and R. Wang, "A transport layer approach for improving end-to-end performance and robustness using redundant paths," in *Proc. 2004 USENIX Conference*.

[3] C. Raiciu, M. Handley, and D. Wischik, "Coupled congestion control for multipath transport protocols," IETF RFC 6356, Oct. 2011.

[4] D. Wischik, M. Handley, and C. Raiciu, "Control of multipath TCP and optimization of multipath routing in the Internet," in *NetCOOP 2009*.

[5] D. Wischik, M. Handley, and M. B. Braun, "The resource pooling principle," in *2008 ACM SIGCOMM CCR*.

[6] N. Vasic and D. Kostic, "Energy-aware traffic engineering," in *Proc. 2010 Inter. Conf. on Energy-Efficient Computing and Networking.*

[7] IEEE P802.3az energy efficient Ethernet. Available: http://www.ieee802.org/3/az/

[8] D. Wischik, C. Raiciu, A. Greenhalgh, and M. Handley, "Design, implementation and evaluation of congestion control for multipath TCP," in *Proc. 2011 USENIX NSDI11 Conference*.

[9] NS-2 network simulator. Available: http://www.isi.edu/nsnam/ns/