Strategic Evaluation of Performance-Cost Trade-offs in a Multipath TCP Multihoming Context

Sinh-Chung Nguyen, Thi-Mai-Trang Nguyen, Guy Pujolle, Stefano Secci
LIP6 - Pierre and Marie Curie University (UPMC), Paris, France. Emails: {firstname.lastname}@lip6.fr

***Abstract*—Today's mobile terminals have several access network interfaces. In practice, the use of different access technologies is subject to different interconnection costs, and mobile users have preferences on interfaces jointly depending on performance and cost factors. There is therefore an interest in defining "light" yet rational multipath communication policies less expensive than greedy ones such as with basic Multipath TCP (MP-TCP). We analyze the performance-cost trade-off of multi-homed end-to-end communications from a strategic standpoint. We model the communication between multi-homed terminals as a multi-criteria non-cooperative game so as to achieve performance-cost decision frontiers. The resulting potential game always allows to select multiple equilibria, which correspond to a strategic load-balancing distribution over the available interfaces, possibly constraining their use with respect to basic MP-TCP. We specify how the resulting model may be in practice implemented by users willing to jointly control the interconnection cost and the performance, based on user Quality of Experience (QoE) assessments. By simulation of a realistic 3-interface scenario, we show how the achievable performance is bound by the interconnection cost; we show that we can halve the interconnection cost with respect to basic (greedy) MP-TCP under a reasonable trade-off, while offering double throughputs with respect to single-path TCP.**

I. Introduction

In the recent few years, mobile terminals have been equipped with several network interfaces. 3G terminals have integrated Wi-Fi and bluetooth antennas. Laptops often have Wi-Fi, Ethernet and 3G access. 4G terminals will also have LTE-A and WiMax interfaces. Indeed, multihoming for mobile terminals becomes a desirable feature because it can provide users with ubiquitous access, enhanced Quality of Experience (QoE) and application performances.

At the Internet Engineering Task Force (IETF), novel protocols such as Stream Control Transmission Protocol (SCTP) [1], Level 3 Multihoming Shim Protocol for IPv6 (SHIM6) [2], Host Indentity Protocol (HIP) [3], multiple Care-of Addresses registration in Mobile IPv6 (mCoA) [4] and Multipath Transmission Control Protocol (MP-TCP) [5][6] have been proposed as possible solutions. Among them, SCTP can be considered as the first transport protocol supporting multihoming. Many studies based on SCTP, especially the proposition of Concurrent Multipath Transfer (CMT) [7], have been carried out to enable the simultaneous data transmission over multiple end-to-end paths. SHIM6, HIP and mCoA are multihoming IP-level protocols for end-hosts. More recently, MP-TCP has been proposed as an extension of TCP to support multihoming, interoperable with the legacy Internet.

The concurrent use of multiple interfaces as allowed by MP-TCP can obviously provide users with a better throughput. However, a greedy use of different access technologies may be costly under common billing schemes. For instance, 3G and 4G accesses are typically more expensive than Wi-Fi or Ethernet ones, due to the use of licensed bands. Certainly, the majority of multi-homed mobile users prefer to use inexpensive technologies as much as possible while maintaining an acceptable performance. The trade-off between performance and cost is therefore subjective, and it seems therefore quite interesting to offer users tools to control it. The specification of such tools is certainly out of scope of the IETF. In fact, the MP-TCP specification and current implementations fully use the available interfaces, which can produce, for example, fast file transfers and better-quality real-time communications. However, in practice, the majority of the users is not willing to greedily use all interfaces concurrently because of performance-cost trade-off preferences.

Actually, research on MP-TCP essentially concentrates on multipath transmission performance improvement, for instance, on joint congestion control of the multiple subflows as in [8], or on reordering avoidance in heterogeneous environments as in [9]. However, there is no work as of our knowledge that investigates on the performance-cost trade-off, and that proposes strategic multi-homed communication mechanisms to constrain the basic greedy mode of MP-TCP.

In this paper, we adopt a game-theoretic approach to model and control the load-sharing over multiple paths. We model the communication between multi-homed endpoints as a bi-criteria non-cooperative game considering both cost and delay factors; a game modeling is appropriate for these situations because each terminal's utility is not only affected by its outgoing interface decision, but also by the other endpoint's decision on its incoming interface decision. To evaluate different trade-off strategies, we extend the existing MP-TCP implementation. By simulation of a realistic 3-interface scenario we show how our strategic load-sharing framework can control the trade-off, highlighting the price to pay (to allow strategic interactions among endpoints) in terms of throughput, and the related savings in terms of interconnection cost, with respect to greedy MP-TCP. In particular, we can halve the interconnection cost while doubling the throughput with respect to basic TCP. Our model is valid for situations in which the achievable throughput with greedy MP-TCP is more than really needed, with a user modestly requiring a moderate increase of throughput with respect to basic TCP, at a reasonable interconnection cost.

The paper is organized as follows. Section II presents our load sharing game framework. In Section III, we evaluate different cost-performance trade-off strategies by simulations. Finally, Section IV concludes the paper.

II. Framework

In this section, we present how to model the communication between multi-homed endpoints with non-cooperative game

TABLE I
GAME WITH NO ENDPOINT COORDINATION (INTERCONNECTION COST)

I \ II	I1	I2
II1	**3,11**	**8,11**
II2	**3,5**	**8,5**

TABLE II
GAME WITH INTERCONNECTION COSTS UNDER JOINT FLOW ROUTING

I \ II	S_1D_1	S_1D_2	S_2D_1	S_2D_2
S_1D_1	6,22	11,22	**6,16**	**11,16**
S_1D_2	6,16	11,16	***6,10***	**11,10**
S_2D_1	11,22	16,22	11,16	16,16
S_2D_2	11,16	16,16	11,10	16,10

TABLE III
GAME WITH UNIDIRECTIONAL INTERCONNECTION COSTS ONLY

I \ II	S_1D_1	S_1D_2	S_2D_1	S_2D_2
S_1D_1	$(9,31)^9$	$(14,31)^9$	$(9,23)^1$	$(14,23)^1$
S_1D_2	$(9,25)^9$	$(14,25)^9$	$(9,17)^1$	$(14,17)^1$
S_2D_1	$(8,31)^8$	$(13,31)^8$	**(8,23)**^{0}	**(13,23)**^{0}
S_2D_2	$(8,25)^8$	$(13,25)^8$	***(8,17)***^{0}	**(13,17)**^{0}

theory, so as to select coordinated load-balancing strategies. We start with a simple game setting, dealing with interconnection costs only, and then we gradually develop the model.

A. Modeling scenario

Let us consider the case where two multi-homed MP-TCP endpoints exchange an equivalent amount of data via multiple available paths. These paths use different interfaces, such as Ethernet, Wifi, 3G, 4G, bluetooth, etc, which have various characteristics in terms of connection cost, bandwidth, and delay. Aiming to improve their connection performance while considering the user interests, endpoints can announce to each other their respective interface preferences. For example, an endpoint may prefer Ethernet to 4G because the Ethernet interface is faster and less expensive.

As a first step, let us model the interaction between the two endpoints as if they did not coordinate the interface path decision. In this case, an endpoint autonomously decides on the destination endpoint's incoming interface, impacting an incoming interconnection cost to the destination endpoint for the incoming flow. For the moment, we do not consider the outgoing interface selection so as to emulate, therefore, the basic MP-TCP behavior, which fully uses the outgoing interfaces without considering their possible interconnection cost. Let us consider a simple example of two endpoints, I and II, with two interfaces each, with associated bidirectional interconnection costs of 3 and 8 for interface I1 and I2, and 11 and 5 for interface II1 and II2, respectively. Taking into account the interconnection cost impacted by the other endpoint decision, we have the strategic game of Table I: endpoints I and II have as strategies the endpoint II's and I's incoming interface, respectively; each cell, corresponding to a strategy profile, indicates the costs for players I and II for that strategy profile, on the left and on the right, respectively. It is easy to notice that all the profiles in Table I are (pure-strategy) Nash equilibria, i.e., for each player there is no preference over the available strategies [10]. Indeed, the game is a dummy game, which highlights that unilaterally selecting the destination's incoming interface without an unilateral performance improvement is a decision rationally not motivated. Therefore, it is necessary to define coordination mechanisms so as to benefit strategically and not greedily from the multihoming capabilities.

The two endpoints can agree in jointly routing their flows following implicit coordination equilibria of the multihoming game. This means accounting not only for the (incoming) cost that the other player decision impacts on its own network, but also for the (outgoing) cost of its own decision. For the moment, let us suppose that for each interface incoming and outgoing interconnection costs are the same.

In Table II, the strategies have now the notation S_iD_j, where i and j indicate the source outgoing interface and the destination incoming interface, i.e., a MP-TCP subflow. In fact, now the decision is not simply on the destination interface where to send the traffic, but also on the source outgoing interface; in MP-TCP, subflows are natively identified and therefore this strategy set seems appropriate to the technology context. Table II indicates in bold the four Nash equilibria of the corresponding balancing game. For example, (S_1D_1, S_2D_2) is a Nash equilibrium but the equal-cost (S_2D_2,S_1D_1) strategy profile is not; indeed, for (S_1D_1, S_2D_2), both the players have no incentive to change their strategies, while for (S_2D_2,S_1D_1) player II has incentives to change to a strategy with a lower unilateral cost such as (S_2D_2,S_2D_1). In addition, among the four (pure-strategy) equilibria of Table II, the italic one (S_1D_2,S_2D_1) is the efficient one (more precisely, Pareto-superior to the others).

An assumption made above is that the incoming cost is equal to the outgoing cost for a given interface. In practice, they may not be the same for a number of cases, as commonly in access networks you have asymmetric service levels (e.g., different upstream and downstream bandwidths). Therefore, a more generic game setting has different incoming and outgoing costs. For instance, in Fig. 1a, for each interface, the incoming cost is close to endpoint while the outgoing cost is near the interface; we obtain the new strategic form of Table III. Also for this case we have four Nash equilibria, with one Pareto-superior to the others. The meaning of the exponent in Table III, as well as the presentation of the resulting game properties need a preliminary mathematical formalization.

B. Notations

The resulting multihoming game can be described as $G_{cost} = (X, Y; f, g) = G_s + G_d$, sum of a selfish game and a dummy game, respectively; let f and g be the cost functions, and X and Y the strategy sets, of endpoint I and endpoint II, respectively. Each strategy $x \in X$ or $y \in Y$ indicates the source and destination interfaces. The strategy set cardinality is equal to the number of source interfaces $\times$ the number of destination interfaces. G_s considers the outgoing cost only, while G_d considers incoming cost only impacted by the other endpoint's interface selection decision.

$G_s = (X, Y; f_s, g_s)$, is a purely endogenous game, where $f_s, g_s : X \times Y \to \mathcal{N}$ are the cost functions for endpoints I and II, respectively. In particular, $f_s(x, y) = \Phi_s(x)$, where $\Phi_s : X \to \mathcal{N}$ and $g_s(x, y) = \Psi_s(y)$, where $\Psi_s : Y \to \mathcal{N}$.

$G_d = (X, Y; f_d, g_d)$, is a game of pure externality, where $f_d, g_d : X \times Y \to \mathcal{N}$. $f_d(x, y) = \Phi_d(y)$, where Φ_d :

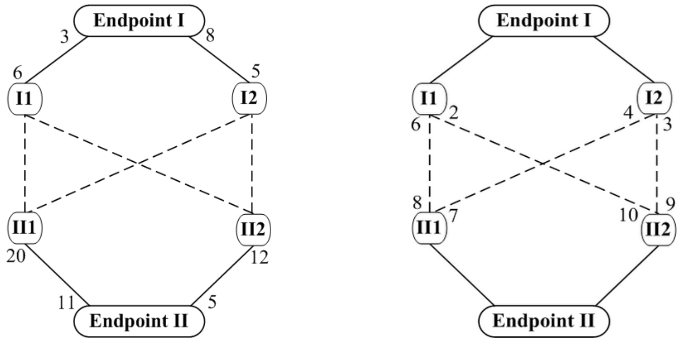

(a) Interconnection costs (b) Path delay costs

Fig. 1. Multihoming example with 8 subflows

$Y \rightarrow \mathcal{N}$ and $g_d(x,y) = \Psi_d(x)$, where $\Psi_d : X \rightarrow \mathcal{N}$. For example, to calculate the cost of strategy (S_2D_1, S_1D_2):
$f_s(S_2D_1, S_1D_2) = \Phi_s(S_2D_1) = 5$
$g_s(S_2D_1, S_1D_2) = \Psi_s(S_1D_2) = 20$
$f_d(S_2D_1, S_1D_2) = \Phi_d(S_1D_2) = 8$
$g_d(S_2D_1, S_1D_2) = \Psi_d(S_2D_1) = 11$.

G_s is a cardinal potential game [11], i.e., the incentive to change players' strategy can be expressed with a single potential function, $P : X \times Y \rightarrow \mathcal{N}$, for all players, and the difference in individual costs by an individual strategy move has the same value as the potential difference. G_d can be seen as a potential game too, but with null potential, so that $G = G_s + G_d$ is a potential game, as sum of potential games. In Table III, and following tables, the exponents to the strategy profiles indicates the corresponding potential values.

Generally, in non-cooperative games, the Nash equilibrium existence is not guaranteed. As property of potential games, the P minimum corresponds to a (pure-strategy) Nash equilibrium and always exists. The inverse is not necessarily true, but it is easy to prove that for G it is true due to the endogenous nature of G_s. To explicate P in calculus an arbitrary starting potential has to be chosen; e.g., in Table III we set to 0 the potential of social welfare profiles, i.e., $P(x_0, y_0) = 0 \ \forall (x_0, y_0) \in X \times Y$ such that $f(x_0, y_0) + g(x_0, y_0) = \min\{f(x,y) + g(x,y)\}$.

C. Accounting for one-way end-to-end delay cost components

In transport level end-to-end communications, several factors can affect the connection performance such as the one-way delay, the round trip time, or for TCP communications the congestion window. In particular, it is well known that in TCP communications the throughput is inversely proportional to the round-trip time. In our multihoming decision context, the end-to-end paths may be asymmetric (the path between the interfaces depends on the Border Gateway Protocol, which implements various routing policies), and a strategy profile indicates a single direction (an MP-TCP subflow) from a source endpoint interface towards a destination endpoint interface. For these reasons, for performance improvement, the simplest yet most appropriate factor one shall include in the game, as an additional cost component, is the one-way delay (obviously, the round-trip time is the sum of the one-way delays along the two subflows in opposite directions). Such an information is nowadays retrievable using Internet monitoring platforms, and is commonly used by many Internet applications (e.g., in overlay services). In Fig. 1b, for each path between the endpoints, a subflow delay cost component is placed next to the outgoing interface. Moreover, it is easy to show that the delay game is symmetric and is a potential game too.

TABLE IV
MULTIHOMING GAME WITH INTERCONNECTION AND DELAY COSTS

I \ II	S_1D_1	S_1D_2	S_2D_1	S_2D_2
S_1D_1	$(23,45)^{10}$	$(27,44)^{9}$	$(25,39)^{4}$	$(29,38)^{3}$
S_1D_2	$(19,35)^{6}$	$(23,34)^{5}$	$(21,29)^{0}$	$\mathbf{(25,28)}^{-1}$
S_2D_1	$(20,43)^{7}$	$(24,42)^{6}$	$(22,37)^{1}$	$(26,36)^{0}$
S_2D_2	$(19,36)^{6}$	$(23,35)^{5}$	$(21,30)^{0}$	$\mathbf{(25,29)}^{-1}$

In order to jointly take both interconnection and delay cost criteria into account for the multihoming coordination, we can integrate the two games into a single one. The objective is to use a multihoming game that takes into account monetary interconnection costs, and a performance criterion that directly affects MP-TCP performances. In order to explore the cost-performance trade-off in the strategic situation, we can extend G as $G = G_{cost} + \beta G_{delay}$, where β is the trade-off coefficient (with $\beta = 0$ just the interconnection cost is taken into account, while as β increases more importance is given to the performance). All the properties previously discussed are maintained for the resulting bi-criteria G game: it is still a potential game as sum of potential games. Table IV shows the resulting strategic form of the example, with $\beta = 1$ (again, the equilibria, corresponding to the potential minima, are highlighted in bold).

So far, the multihoming game example has showed only a limited number of equilibria. The multihoming solution corresponding to the two equilibria of Table IV implies that the endpoint I uses the two MP-TCP subflows $S_1 \rightarrow D_2$ and $S_2 \rightarrow D_2$, evenly distributing the load on them, and that the endpoint II uses the subflow $S_2 \rightarrow D_2$. However, with the objective to further enlarge the equilibrium set, and therefore the number of used subflows, while allowing for arbitrary load-balancing on the selected subflows, we can exploit the potential value as described in the next section.

D. Strategic load-balancing distribution computation

In potential games, the potential value qualifies the profile propensity to reaching equilibrium and predicts the behavior of the potential game: the lower it is, the finer the profile is. In fact, potential value can help in extending the equilibrium set including also those profiles that are not pure-strategy equilibria, but that have a possibility of becoming equilibria if minor changes occur. With the aim of increasing the diversity of the load-balancing decision, we can thus elevate those profiles that are not Nash equilibria, but that have a very low potential, to the equilibrium status and include them in the load-balancing decision. This corresponds to selecting as equilibrium all the strategy profiles that have a potential value equal or below a threshold.

Therefore, we can exploit the potential as a means to increase the path diversity of the multihoming game solution. Increasing the potential threshold, the equilibrium set is larger and the set of used interfaces is larger, while guaranteeing that they are rationally selected. Since the trade-off coefficient β can already be used to enhance performance by weighting the importance of the one-way delay in the multihoming decision, the way the potential threshold is computed shall depend on β. An acceptable simple way to compute the potential threshold (τ) as a function of β is to set it linearly with β between the minimum (P_{min}) and the maximum (P_{max}) potential:

$$\tau(\beta) = (P_{max} - P_{min}) \cdot \beta / \beta_{max} + P_{min} \qquad (1)$$

Let $S \in X \times Y$ be the set of strategy profiles with a potential below the potential threshold τ (hence kept as solution equilibria), i.e., $\forall(x, y) \in S \; P(x, y) < \tau$. A still open problem is therefore to compute the load-distribution among the interfaces corresponding to the selected equilibria in S. It cannot be an even load-balancing, because a subflow load should instead be an arbitrary distribution computed as a function of the potential values of the equilibria for the subflow. Let b_x and b_y be the load-balancing ratio for strategy $x \in X$ and $y \in Y$, for endpoints I and II, respectively. The load-balancing ratios can be computed as the proportional weight, with respect to the distance from the potential threshold, of the unilateral strategy over all the available strategy profiles (dually $\forall y' \in Y$):

$$b_{x'} = \frac{\sum_{(x,y)\in S}^{x=x'}[1 + \tau - P(x, y)]}{\sum_{(x,y)\in S}[1 + \tau - P(x, y)]}, \; \forall x' \in X \qquad (2)$$

E. A user QoE feedback policy

In practice, a specific policy can be conceived around the tuning of the trade-off coefficient β and therefore the path diversity induced by τ. The principle is that the end-user, aware of the interconnection cost and perceiving the performance, can decide if increasing the Quality of Experience (QoE) at the expense of a higher interconnection cost. For instance, a QoE-feedback policy could be implemented by an application installed in mobile terminals. The user would simply tune the performance-cost coefficient β, enabling the usage of more subflows while loading the least-expensive interfaces (with an effect on both endpoints). The following simulation results can allow understanding which values of β would be likely chosen by the user during the exploration of the trade-off frontier.

III. Simulation results

We present simulation results to assess the cost-performance tradeoff of our approach, highlighting the differences with the basic MP-TCP implementation. We extended the NS-2 MP-TCP implementation [12]. It is important to mention that, for the sake of simplicity, in the simulations we use the same β value for both players, while in practice the different values for the two players are uncorrelated and can be different.

We emulated a case with three interfaces at each endpoint, with 10 Mbps links. The interface connection cost and path delay are randomly chosen as shown in Fig. 2. We generated persistent FTP traffic in both directions for 60s, with a trade-off coefficient β ranging in the interval [0.01, 4] (above 4 there is no relevant change for the given example). In the given example, for $\beta = 4$ all strategy profiles are used (however, this does not correspond to greedy MP-TCP since a strategic load-balancing is still enforced), and for $\beta = 0.01$ only one Nash equilibrium appears (which in fact corresponds to single-path TCP over the least cost interface).

Fig. 3a shows the throughput plot, with one curve for each endpoint, the global throughput and the throughput

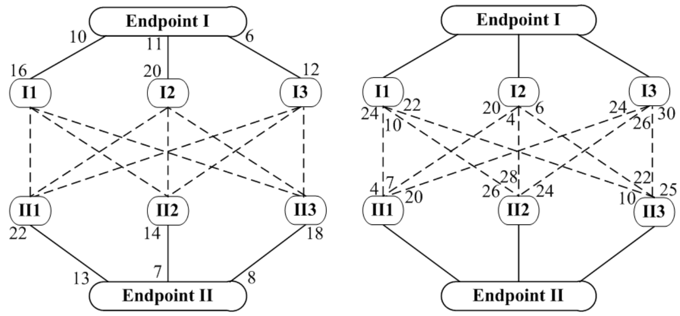

(a) Interconnection costs

(b) Path delay cost

Fig. 2. Multihoming example with 18 subflows

with greedy MP-TCP. As expected, the throughput generally increases with β, and so does the path diversity, since more importance is given to the one-way delay and more subflows are selected (see Fig. 4 with the load-balancing distribution). However, this is not a continuous throughput increase, sudden variations in single-player throughput are in fact due to subflow changes, while smooth variations are induced by equilibrium set modification without subflow changes. At the highest values of β, all subflows are used at both sides (see Fig. 4). However, the throughput of greedy MP-TCP is not reached as we keep computing the load-balancing distribution assigning higher weight to the equilibria with lower potential. The gap with greedy MP-TCP, about 35% less, is the price to pay to maintain a strategic load-distribution and ensure a rationally acceptable coordination between endpoints. On the other hand, even with low values of β (e.g., $\beta < 1$) we obtain a throughput up to four times the single-path TCP throughput (for $\beta = 0.01$). Indeed, our distribution follows the requirements of those users willing to access at the least possible cost multiple paths in a coordinated way. The coordination granted by the game-theoretic modeling of our approach manifests in Fig. 3 with quite close throughput and interconnection cost for the two players.

Fig. 3b shows the interconnection cost results: its increase as a function of β is quite similar to the throughput increase behavior, since a performance improvement always comes with an interconnection cost increase. With low values of β, that is, with more importance given to the interconnection cost than to the performance, one can save about 50% in interconnection cost with respect to high values of β. Moreover, our strategic MP-TCP scheme grants more than 50% saving with respect to greedy MP-TCP, for low trade-off values (e.g., $\beta < 1$).

Coupling the analysis of plots in Fig. 3, one can deduce that with a user QoE feedback policy, the trade-off coefficient tuning would end with a value of β that corresponds to local minima of the interconnection cost. The tuning of β would consist in moving from a local minima to a next one. For example, in Fig. 3, it is easy to identify six values of β corresponding to local minima, indicated by vertical lines. In particular, the most likely chosen values will be the one with the longest distance to the next minima, in our case $\beta = 0.65$. Such trade-off equilibrium points can be the result of an autonomous learning, or of an initiation learning phase between MP-TCP speakers.

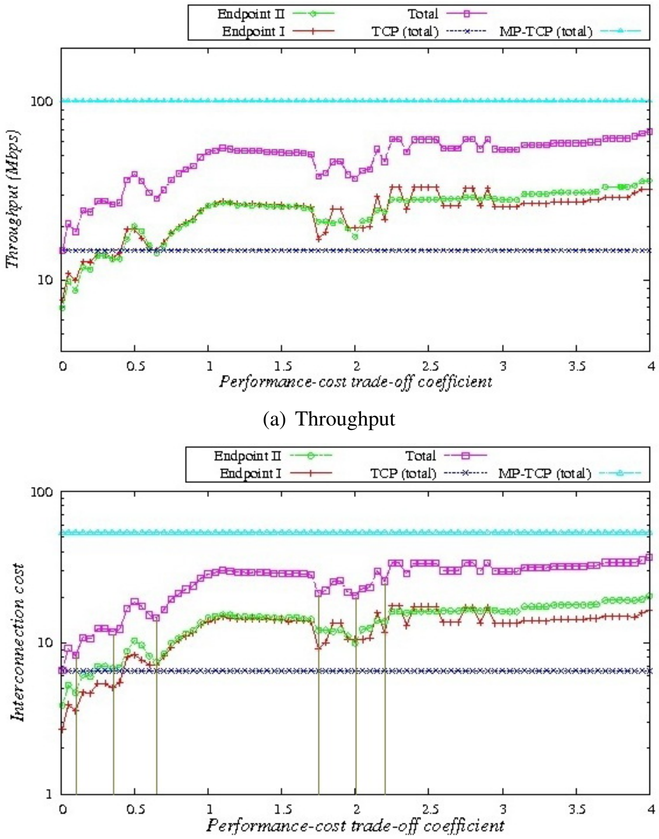

(a) Throughput

(b) Interconnection cost

Fig. 3. Results as a function of the trade-off coefficient

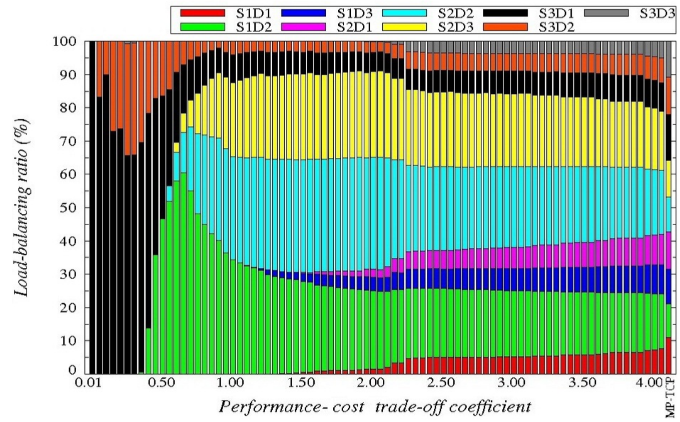

Fig. 4. Load-balancing distribution

As evidenced by Fig. 4, reporting the load-balancing distribution for the first endpoint (similar for the second), the $\beta = 0.65$ point corresponds to a solution with 4-5 subflows and 2-3 interfaces concurrently used for endpoint I. We can observe how giving higher importance to performance (increasing β) the path diversity (number of subflows and interfaces) increases until reaching the maximum number of subflows (18), as done with greedy MP-TCP (plotted in the last column for comparison). Indeed, basic MP-TCP greedily uses all the available subflows and interfaces (filling the corresponding buffer in a round-robin fashion). Nevertheless, even for high values of β, our distribution differs from greedy MP-TCP because we differently weight the load using the potential value of the corresponding strategy profiles.

IV. Conclusion

With the extremely rapid pace at which mobile Internet usages increase, novel solutions have been proposed to increase the performance of multihomed devices with many Internet access interfaces. The most recent and interoperable one seems MultiPath TCP (MP-TCP), which in its current form fully uses the available interfaces while performing multiple end-to-end subflow control. Nevertheless, the basic specification of MP-TCP does not cover practical issues related to the different costs of access technologies. The objective of this paper is to precisely study this topic, assessing the importance of the performance-cost trade-off and proposing a strategic MP-TCP load-balancing scheme mixing performance and interconnection cost factors.

We modeled the interaction among distant multihomed devices as a bi-criteria non-cooperative game so as to allow a rational coordination towards multihoming equilibria. In particular, a trade-off coefficient allows users to weight the MP-TCP subflow load-balancing according to their propensity to pay more for the interconnection, hence to get a better quality of experience. Simulation results show that our approach can grant roughly 50% saving with respect to greedy MP-TCP, under an acceptable trade-off, with a roughly double throughput with respect to single-path TCP. Moreover, it is possible to identify isolated values of the trade-off coefficient following local minima of the interconnection cost behavior; we described a rational yet light coordination scheme among MP-TCP endpoints to set up arbitrary load-balancing distribution on the available subflows. More generally, our analysis allows understanding the rather unexplored aspect of performance-cost tradeoff in access multihoming.

References

[1] R. Stewart, "Stream Control Transmission Protocol". RFC 4960,. 2007.
[2] E. Nordmark, M. Bagnulo, "SHIM6: Level 3 Multihoming Shim Protocol for IPv6". RFC 5533, 2009.
[3] P. Nikander, T. Henderson, C. Vogt, J. Arkko "End-Host Mobility and Multihoming with the Host Identity Protocol". RFC 5206, 2008.
[4] R. Wakikawa et al., "Multiple Care-of Addresses Registration". RFC 5648, Oct. 2009.
[5] A. Ford et al, "Architectural Guidelines for Multipath TCP Development". RFC 6182, 2011.
[6] A. Ford, C. Raiciu, M. Handley, O. Bonaventure, "TCP Extensions for Multipath Operation with Multiple Addresses". drat-ietf-MP-TCP-multiaddressed-05, 2012.
[7] J.R. Iyengar, P.D. Amer, R. Stewart, "Concurrent Multipath Transfer Using SCTP Multihoming Over Independent End-to-End Paths". *IEEE/ACM Transactions on Networking*, Vol.14, No. 5, Oct. 2006.
[8] C. Raiciu, M. Handley, D. Wischik, "Coupled Congestion Control for Multipath Transport Protocols". RFC 6356, 2011.
[9] X. Zhang, T.M.T Nguyen, G. Pujolle, "A Cross-layer Approach to Optimize the Performance of Concurrent Multipath Transfer in Wireless Transmission". in *Proc. of Wireless days 2009*.
[10] R.B. Myerson, *Game Theory: Analysis of Conflict*, Harvard Univ. Press.
[11] D. Monderer, L.S. Shapley, "Potential Games", *Games and Economic Behavior*, Vol. 14, No. 1, May 1996.
[12] NS-2 MP-TCP working group: http://code.google.com/p/multipath-tcp