

Optimal Number of Paths with Multipath Routing in Hybrid Networks (Technical Report)

Sébastien Henri
EPFL, Switzerland

Patrick Thiran
EPFL, Switzerland

Abstract—In recent years, multipath routing, i.e., employing several paths simultaneously, has emerged as an efficient way to provide significant throughput gains in local networks. This has been observed both with technologies that are not subject to interference, such as Ethernet, and with technologies that are, such as WiFi, power-line communications (PLC) and LTE. With technologies that are subject to interference, adding more paths is not always beneficial. We investigate the number of simultaneous paths necessary to reach maximal throughput when using multipath routing in multi-hop mesh networks with several self-interfering technologies. We show analytically, numerically and experimentally that the optimal number of paths M^{opt} is tightly linked with the number of technologies K . For certain classes of networks (in particular, for typical home networks), we prove analytically that $M^{\text{opt}} = K$, and our analytical findings are verified both with simulations and with experiments on a testbed composed of PLC and two orthogonal WiFi channels. In general networks, our numerical and experimental results show that the throughput loss caused by using at most K simultaneous paths is very small: The relative loss is smaller than 0.05 in 97% of the networks and smaller than 0.1 in 99% of the networks.

I. INTRODUCTION

The demand of users for high-throughput connectivity is increasing very rapidly, and new high-throughput standards have been recently published: For example, LTE for cellular networks, 802.11n and 802.11ac for WiFi, and IEEE 1901 for power-line communications (PLC). However, it is becoming increasingly difficult to significantly improve throughput by modifying the technology standards. For this reason, other directions of improvement have been pursued: (i) It is possible to use multiple technologies (e.g., WiFi, PLC, LTE), as illustrated by the standardization of hybrid networks by the IEEE 1905 working group [2]. In this paper, two orthogonal WiFi channels are considered as two different technologies. (ii) Mesh networking is also gaining momentum and an increasing number of commercial solutions are proposed [4]; mesh networking effectively improves performance, in particular coverage, but increases the complexity when compared to the infrastructure mode, because of multi-hop paths. (iii) It is also possible to simultaneously use several paths, as illustrated by the development of multipath routing solutions such as multipath TCP (MPTCP), in particular with WiFi and LTE [21]. When employed with technologies that are not subject to interference, such as Ethernet, the gains provided by multipath are very important [20]; in principle, it is always beneficial, in terms of throughput, to add as many paths as possible (obviously, adding more paths is not necessarily possible

and gives rise to other issues, such as scalability and power consumption). In hybrid local networks with *self-interfering* technologies, i.e., shared-medium technologies where two distinct links that use the same technology are subject to interference (such as WiFi, PLC, and LTE), multipath routing also provides significant throughput gains [5], [11]. However, with self-interfering technologies, the optimal number of paths M^{opt} (i.e., the minimal number of paths to employ simultaneously necessary to reach maximal throughput) is not obvious to find. Adding more paths does not always improve throughput and can even degrade it [14], [18]. Yet, to the best of our knowledge, no work has investigated analytically the number of paths necessary to reach maximal throughput when using multipath routing in multi-hop hybrid mesh networks with self-interfering technologies. Here, we show analytically and experimentally that, in multi-hop mesh networks with K technologies that self-interfere but do not interfere with each other, the optimal number of paths M^{opt} is tightly linked with K . For certain classes of networks (in particular, for typical home networks), we prove analytically that $M^{\text{opt}} = K$, and our analytical findings are verified both with simulations and with experiments on a testbed composed of PLC and two orthogonal WiFi channels ($K = 3$). In general networks with K distinct self-interfering technologies, our numerical and experimental results show that the throughput loss caused by employing only multipaths composed of at most K paths is very small. Knowing the limit on the number of paths to employ simultaneously with multipath routing has a practical interest: It means that it is possible to limit the size of the multipaths returned by a multipath-routing protocol without harming the performance of the protocol.

The rest of the paper is organized as follows. We present the related work in Section II. In Section III, we describe our model. In Section IV, we present our analytical findings valid for certain classes of networks; these findings are verified with simulations and experiments in Section V, where we also present numerical and experimental results for general networks. We conclude in Section VI.

II. RELATED WORK

Multipath routing has been widely studied in several contexts: mobile ad-hoc networks (MANETs) [23], wireless sensor networks (WSNs) [3], mesh networks [25] and traffic engineering [13]. In MANETs and WSNs, multipath-routing protocols have been shown to have several advantages, such

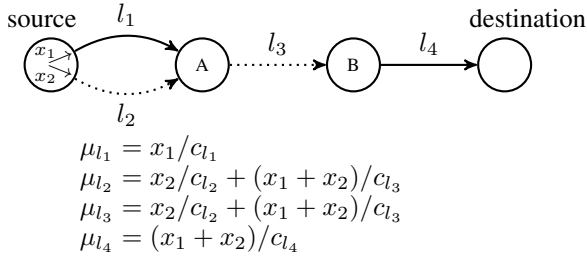


Fig. 1. Illustration of a multigraph. Dotted lines represent WiFi, plain lines PLC. The source sends traffic on the two paths P_1 and P_2 at respective rates x_1 and x_2 ($\Lambda_{P_1} = \{l_1, l_3, l_4\}$ and $\Lambda_{P_2} = \{l_2, l_3, l_4\}$).

as reduced delays and better reliability and throughput [10], [16], [24]. These protocols use heuristics to build the paths and they are consequently not guaranteed to be optimal. Optimal multipath routing and scheduling have also been studied in several works, mostly at a theoretical level [15], [17], [26]. These papers do not study the optimal number of paths; rather, they find the optimal rate provided by doing joint routing and scheduling. Multipath routing has recently received renewed attention, in particular with the development of multipath TCP (MPTCP) [7], [8]. For practical implementations of multipath, e.g., with MPTCP, the set of paths is chosen in advance and congestion control is then carried on these chosen paths. There are many works that aim at finding the best multipath in hybrid networks by using heuristics, either by explicitly trying to maximize throughput [11] or by looking for maximally disjoint paths [6], [9], [24], but they do not guarantee optimality. To the best of our knowledge, no work has yet addressed the question of finding the optimal number of paths when using multipath routing in hybrid mesh networks with self-interfering technologies.

III. MODEL

We consider a multi-hop mesh network with K different self-interfering technologies that do not interfere with each other (e.g., PLC, WiFi, LTE). Two orthogonal WiFi channels are considered as two different technologies. The network is modelled by a *multigraph* $G(\mathcal{V}, \mathcal{E})$, with \mathcal{V} the set of nodes and \mathcal{E} the set of links. \mathcal{E} is partitioned into K sets \mathcal{E}_k , $k \in \{1, \dots, K\}$, the sets of links available with each technology. A link is present whenever its two endpoints can communicate with each other with a non-zero rate on the corresponding technology. Figure 1 presents an example of a multigraph with $K = 2$ technologies, e.g., PLC and WiFi (here, $\mathcal{E}_1 = \mathcal{E}_{\text{PLC}} = \{l_1, l_4\}$ and $\mathcal{E}_2 = \mathcal{E}_{\text{WiFi}} = \{l_2, l_3\}$). For a link $l \in \mathcal{E}$, c_l is the capacity of l , i.e., the maximum rate achievable on l . For a link $l \in \mathcal{E}_k$, $\mathcal{I}_l \subset \mathcal{E}_k$ is the *interference domain* of l , defined as the set that contains l and all links that cannot transmit simultaneously with l (because doing so would cause a collision at one of the links). For example, in Figure 1, WiFi links l_2 and l_3 interfere, i.e., $\mathcal{I}_{l_2} = \mathcal{I}_{l_3} = \{l_2, l_3\}$, but PLC links l_1 and l_4 do not, i.e., $\mathcal{I}_{l_1} = \{l_1\}$ and $\mathcal{I}_{l_4} = \{l_4\}$.

If a node transmits data to another node, we call the source-destination pair a *flow*. A *path* is a self-avoiding path of

the multigraph G that connects two nodes. The source of a flow can use M paths P_1, \dots, P_M simultaneously; the set $\mathcal{P} = (P_1, \dots, P_M)$ is called a *multipath*. When $M = 1$, the multipath is a single path. The set of links belonging to any path P_i is denoted by Λ_{P_i} , with $\Lambda_{P_i} \subseteq \mathcal{E}$; for a multipath $\mathcal{P} = (P_1, \dots, P_M)$, we write $\Lambda_{\mathcal{P}} = \bigcup_{i=1}^M \Lambda_{P_i}$, and $L_{\mathcal{P}} = |\Lambda_{\mathcal{P}}|$ for the total number of links in the multipath. For example, in Figure 1, there are $M = 2$ paths, P_1 with $\Lambda_{P_1} = \{l_1, l_3, l_4\}$ and P_2 with $\Lambda_{P_2} = \{l_2, l_3, l_4\}$.

We define the *busy time* μ_l of a link l as the *fraction of time* during which no node can initiate a transmission on l , because either (i) a node in the interference domain \mathcal{I}_l is already transmitting, or (ii) the channel is idle, but the node cannot transmit because it needs to wait for the expiration of an inter-frame space, or because it is in backoff stage.

Assumption 1: When a node sends traffic at rate x_l on a single link l with no other link transmitting, we assume that if the link is not saturated (i.e., $x_l \leq c_l$), then it will obtain a busy time μ_l that is proportional to x_l , $\mu_l = \frac{x_l}{c_l}$.

Figure 1 illustrates the busy time with interfering links.

The source of a flow sends data at rate x_i on each path P_i for $i \in \{1, \dots, M\}$, and we denote by $\mathbf{x}_{\mathcal{P}}$ the rate vector $[x_i]_{i \in \{1, \dots, M\}}$. If $x_i = 0$, path P_i is not used. For each link $l \in \Lambda_{\mathcal{P}}$, the total busy time (accounting for interference) follows, if links are not saturated, from the equation of the busy time, and is given by

$$\mu_{l, \mathbf{x}_{\mathcal{P}}} = \sum_{l' \in \mathcal{I}_l} \mu_{l'} = \sum_{i=1}^M x_i \sum_{l' \in \mathcal{I}_l \cap \Lambda_{P_i}} \frac{1}{c_{l'}} = \sum_{i=1}^M x_i \alpha_{P_i, l}, \quad (1)$$

where we define $\alpha_{P_i, l} \doteq \sum_{l' \in \mathcal{I}_l \cap \Lambda_{P_i}} \frac{1}{c_{l'}}$. We say that a rate vector $\mathbf{x}_{\mathcal{P}}$ is *admissible* if for all $l \in \Lambda_{\mathcal{P}}$, $\mu_{l, \mathbf{x}_{\mathcal{P}}} \leq 1$ (the busy-time never exceeds 100%).

Writing $\boldsymbol{\alpha}_{\mathcal{P}, l} \in \mathbb{R}^M$ for the vector $[\alpha_{P_i, l}]_{i \in \{1, \dots, M\}}$, Equation (1) can be recast as $\mu_{l, \mathbf{x}_{\mathcal{P}}} = \boldsymbol{\alpha}_{\mathcal{P}, l}^T \cdot \mathbf{x}_{\mathcal{P}}$ with T denoting transposition. $\boldsymbol{\alpha}_{\mathcal{P}, l}$ is called the *multipath-impact vector* of \mathcal{P} on l ; it depends only on the network topology (i.e., the link capacities and interference domains) and on the paths, and *not* on the rate vector $\mathbf{x}_{\mathcal{P}}$. We denote by $\boldsymbol{\mu}_{\mathbf{x}_{\mathcal{P}}} \in \mathbb{R}^{L_{\mathcal{P}}}$ the vector $\boldsymbol{\mu}_{\mathbf{x}_{\mathcal{P}}} = [\mu_{l, \mathbf{x}_{\mathcal{P}}}]_{l \in \Lambda_{\mathcal{P}}}$ and by $\mathbf{A}_{\mathcal{P}} \in \mathbb{R}^{L_{\mathcal{P}} \times M}$ the matrix

$$\mathbf{A}_{\mathcal{P}} = [\boldsymbol{\alpha}_{\mathcal{P}, l}^T]_{l \in \Lambda_{\mathcal{P}}}. \quad (2)$$

With the example of Figure 1 and $\mathcal{P} = (P_1, P_2)$, we have

$$\mathbf{A}_{\mathcal{P}} = \begin{bmatrix} \boldsymbol{\alpha}_{\mathcal{P}, l_1}^T \\ \boldsymbol{\alpha}_{\mathcal{P}, l_2}^T \\ \boldsymbol{\alpha}_{\mathcal{P}, l_3}^T \\ \boldsymbol{\alpha}_{\mathcal{P}, l_4}^T \end{bmatrix} = \begin{bmatrix} 1/c_{l_1} & 0 \\ 1/c_{l_3} & 1/c_{l_2} + 1/c_{l_3} \\ 1/c_{l_3} & 1/c_{l_2} + 1/c_{l_3} \\ 1/c_{l_4} & 1/c_{l_4} \end{bmatrix}.$$

The optimal rate vector on multipath \mathcal{P} , denoted by $\mathbf{x}_{\mathcal{P}}^{\text{opt}}$, is the admissible rate vector that maximizes the 1-norm. Because $\mathbf{A}_{\mathcal{P}} \cdot \mathbf{x}_{\mathcal{P}} = \boldsymbol{\mu}_{\mathbf{x}_{\mathcal{P}}}$, $\mathbf{x}_{\mathcal{P}}^{\text{opt}}$ is a solution of the following system:

$$\begin{aligned} & \max_{\mathbf{x}} \mathbf{1}^T \cdot \mathbf{x} \\ & \text{subject to } \mathbf{A}_{\mathcal{P}} \cdot \mathbf{x} \preceq \mathbf{1} \\ & \mathbf{x} \succeq \mathbf{0}. \end{aligned} \quad (3)$$

where \succeq and \preceq denote component-wise inequalities.

For a given flow, the optimal rate or optimal throughput (the two terms are used interchangeably in this paper) is

$$x^{\text{opt}} \doteq \max_{\mathcal{P} \in \Pi} \|\mathbf{x}_{\mathcal{P}}^{\text{opt}}\|_1, \quad (4)$$

where Π denotes the set of all possible multipaths for the flow. We define the optimal number of paths M^{opt} as the minimal number of paths in a multipath \mathcal{P}^{opt} reaching the optimal rate x^{opt} . In particular, all the M^{opt} paths of \mathcal{P}^{opt} are used. Because a path has no loop, there is a finite number of paths between the source and the destination, and in theory, it is possible to find x^{opt} by computing $\mathbf{x}_{\mathcal{P}^{\text{all}}}^{\text{opt}}$ where \mathcal{P}^{all} is the multipath containing all the possible paths for a given flow. However, the number of paths in \mathcal{P}^{all} grows exponentially with the number of nodes and technologies, which makes this method impractical. In fact, there is no better solution for finding x^{opt} and \mathcal{P}^{opt} : It has been shown that in a network with interference, finding an optimal multipath is NP-hard [12].

Even if we can find the optimal rate x^{opt} , finding M^{opt} is still very challenging: It has been shown that computing the minimal-rank solution of a linear problem is also NP-hard [22]. This means that finding M^{opt} by computing the minimal-rank solution of System (3) is not practical: The number of possible multipaths grows exponentially with the number of paths, that grows itself exponentially with the number of nodes and technologies. In our analysis of Section IV, we study M^{opt} without searching for an optimal multipath of minimum rank. Finding M^{opt} without knowing the corresponding optimal multipath \mathcal{P}^{opt} remains of practical interest, because it makes it possible to limit the size of the multipaths returned by a multipath-routing protocol without harming the performance of the protocol.

IV. ANALYSIS OF THE OPTIMAL NUMBER OF PATHS

We investigate analytically M^{opt} when using multipath routing in hybrid networks with technologies that self-interfere but do not interfere with each other. The key result is that in a hybrid network with self-interfering technologies, M^{opt} is tightly linked with the number K of technologies; in specific networks, we prove that $M^{\text{opt}} = K$. This shows that with self-interfering technologies, multipath routing provides throughput gains mainly in hybrid networks (i.e., when $K \geq 2$) and is not useful in small single-technology networks.

We define the following terms.

Definition 1: The network is *multi-complete* if for every technology k , every link $l \in \mathcal{E}_k$ of the network interferes with every other link $l' \in \mathcal{E}_k$ (i.e., the interference graph for each technology is complete).

Note that this does *not* mean that the graph $(\mathcal{V}, \mathcal{E}_k)$ is complete (i.e., that every node is directly connected with every other node): For example, in Figure 2 that represents a typical network for a five-room home, all WiFi links (dotted lines) interfere with each other and all PLC links (plain lines) interfere with each other, i.e., the network is multi-complete; but there is no direct WiFi and PLC link for example between Node A and Node C and between Node D and Node E.

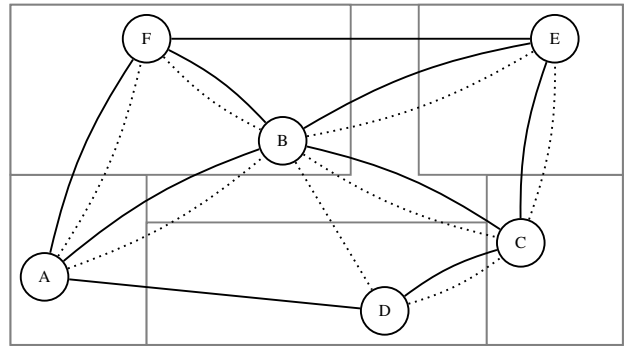


Fig. 2. Example of a typical network for a five-room home with 6 nodes and $K = 2$ technologies: WiFi (dotted lines) and PLC (plain lines). This network is both multi-complete and multi-connected.

Definition 2: The network is *multi-connected* if all the K sub-networks $(\mathcal{V}, \mathcal{E}_1), \dots, (\mathcal{V}, \mathcal{E}_K)$ are connected: For each technology k , each node in \mathcal{V} can communicate with each other node in \mathcal{V} , possibly with multiple hops, by using only links of \mathcal{E}_k .

The network represented in Figure 2 is multi-connected: every node can reach every other node by using only WiFi links and only PLC links, possibly with multi-hop paths.

Theorem 1: In a multi-complete network, $M^{\text{opt}} \leq K$.

The proofs of Theorem 1 and of all the results of this section are presented in Appendix.

Because small networks (e.g., typical home networks) are likely to be multi-complete, Theorem 1 shows in particular that in small networks with a single technology ($K = 1$), multipath routing is likely to be useless in terms of throughput.

The next analytical results are valid under the following assumption. In the numerical and experimental results presented in Section V, this assumption is not made, and we show that the results remain true in the vast majority of the cases.

Assumption 2: A property that depends on the link capacities is true if and only if there is $\epsilon_0 > 0$ such that for all $0 < \epsilon < \epsilon_0$ and all links $l_0 \in \mathcal{E}$, the property remains true if the capacity of l_0 is modified as $c_{l_0}(1 \pm \epsilon)$ whereas the capacities of the other links remain unchanged.

This means that to be true, a property must be robust against small variations of the link capacities. For example, Assumption 2 yields that for two links $l, k \in \mathcal{E}$, $c_l = c_k$ if and only if $l = k$: If we assume $c_l = c_k$ and $l \neq k$, then adding any small $\epsilon > 0$ to one of the two link capacities invalidates the equality, because $c_l + \epsilon \neq c_k$. More generally, Assumption 2 yields that

$$\sum_{l \in S_1} \frac{1}{c_l} = \sum_{l \in S_2} \frac{1}{c_l} \text{ if and only if } S_1 = S_2. \quad (5)$$

Lemma 1: With $M \leq K$, an admissible rate vector $\mathbf{x}_{\mathcal{P}}$ sent on a multipath \mathcal{P} with M paths saturates at most M technologies, i.e., at most M technologies have a link whose busy time is exactly 1.

Theorem 2: In a multi-connected network, $M^{\text{opt}} \geq K$.

Corollary 1: In a multi-complete and multi-connected network, $M^{\text{opt}} = K$.

For example, in the typical home network presented in Figure 2 that is both multi-complete and multi-connected, Corollary 1 shows that the optimal number of paths M^{opt} is equal to the number of technologies K .

V. NUMERICAL AND EXPERIMENTAL RESULTS

In this section, we verify with simulations and testbed experiments the findings of our analysis for multi-complete and multi-connected networks. We also present numerical and experimental results for more general networks that are not necessarily multi-complete and multi-connected.

A. Benchmarking Methods for the Optimal Number of Paths

1) *Brute-force method:* As explained in Section III, finding an optimal multipath in a network with interference is NP-hard [12]. To the best of our knowledge, all practical multipath-routing protocols use heuristics and do not guarantee the optimality of their results. We can find an optimal multipath by using brute-force, i.e., by computing the set \mathcal{P}^{all} of all possible paths and by solving (3) with $\mathcal{P} = \mathcal{P}^{\text{all}}$. However, the number of possible paths is exponential in the number of nodes and technologies, which makes this method highly computation-intensive. In fact, in the simulations of Section V-B, we have to limit the number of paths in order to be able to solve (3). We do so, in such a way that this is very unlikely to change the final result, by limiting the number of hops in the paths. We set the maximum number of hops to three times the minimum number of hops between the source and the destination. For example, if there exists at least one two-hop path between the source and the destination but no single-hop path, we limit the paths between the source and the destination to those with six or fewer hops. If the source and the destination are two-hop away, it is very unlikely that paths of seven or more hops are required to reach the optimal rate x^{opt} .

2) *Backpressure method:* It is also possible to find a multipath arbitrarily close to an optimal one with the method described by Neely et al. [17]. This method employs a backpressure scheme: The source initially floods the network by sending traffic in all directions; packets that arrive at the destination are removed from the network, whereas other packets stay in the queues of the nodes. Gradually, traffic is sent only to nodes that have small queues, which indicates that they are in the “right” direction. This scheme is shown to converge arbitrarily close to the maximal achievable rate x^{opt} (given by Equation (4)), namely to $(1 - 1/V) \cdot x^{\text{opt}}$ for some constant V . In the simulations and experiments, we choose $V = 1000$, such that the difference between the optimal rate and the rate that is found by this method is at most $0.001 \cdot x^{\text{opt}}$. By considering only the links that are used once the scheme has converged (i.e., the links on which traffic is sent at rate above the threshold $0.001 \cdot x^{\text{opt}}$), we compute a multipath that yields a rate arbitrarily close to optimum, and we assume that it uses the same number of paths as an optimal multipath.

These two methods give us one multipath reaching the optimal rate (or arbitrarily close to it), but they are not guaranteed to return the optimal multipath with a minimal number of paths, consequently we can only compute an upper bound of M^{opt} . As explained at the end of Section III, computing the minimal-rank solution of a linear problem is NP-hard; here, with up to several millions of possible paths, the number of possible multipaths is far too large to enable us to find the exact value of M^{opt} . However, we believe that in practice, there is a single optimal multipath in most of the cases, and that the upper bound is therefore tight in most of the cases.

We have only a benchmarking goal when we experimentally evaluate M^{opt} , and the efficiency of the benchmarking schemes is not the subject of this paper. In fact, both these benchmarking schemes are impractical. The brute-force scheme requires solving a system whose size is exponential in the number of nodes and technologies; most of our simulations take several hours to find the result. The back-pressure scheme would also be difficult to use in a real-world application for several reasons: (i) It requires knowing the interference domain of each link in advance, which is typically challenging or impractical [19]. (ii) It requires a centralized coordinator that decides at each time slot which links are to be used. (iii) It initially floods the entire network.

B. Simulation Results

We present results obtained with a Matlab simulator.¹ Each node has $K = 3$ technologies, with random ranges between 20 m and 40 m, and random maximum rates between 20 Mb/s and 180 Mb/s. Each link capacity is distributed according to a linear function that decreases with the distance, to which is added a zero-mean normally-distributed noise, with parameters chosen such that the capacities are close to the ones observed on our WiFi-PLC testbed (see our previous work [11] for a more detailed description). One technology uses the parameters found for PLC, the other two use the parameters found for WiFi (i.e., we simulate networks that have PLC and two orthogonal WiFi channels). The technologies do not interfere with each other, but are self-interfering. Two links l and l' of a same technology interfere if one node of l and one node of l' are within range of each other. We compute an upper bound on M^{opt} with the two different benchmarking methods described in Section V-A, and we keep the minimum of the two results (in the following, we slightly abuse notation and write M^{opt} for this upper bound).

We first simulate a multi-connected and multi-complete network in order to compare our analysis with the simulation results. The network, denoted by Network 1, is a 40×40 m square with 10 nodes randomly placed. We simulate 1000 different random instances of Network 1: For each instance, the placement of the nodes is made uniformly at random, the choice of the link capacities is made randomly according to the distribution described in the previous paragraph, and

¹We built on a simulator written for previous work [11]. The code of the simulator and the code for the experiments presented in Section V-C are available at <https://c4science.ch/diffusion/6360>.

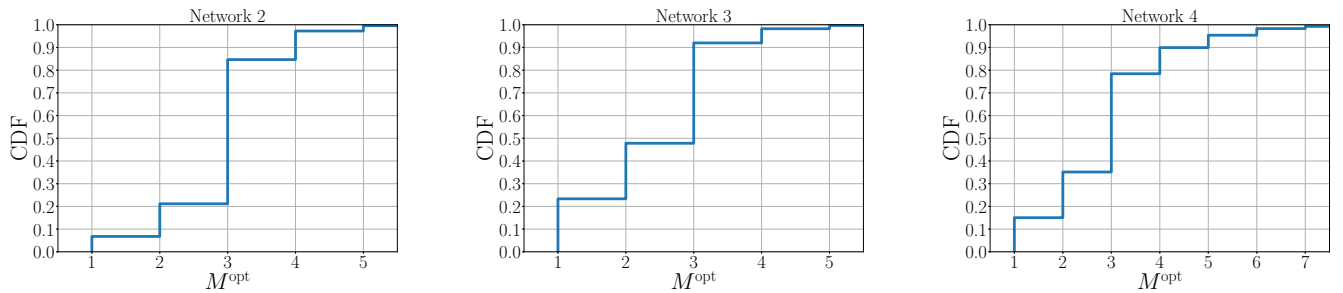


Fig. 3. Cumulative distribution function of M^{opt} for Network 2 (left), a general 100×100 m network with 15 nodes, Network 3 (center), a general 200×150 m network with 20 nodes, and Network 4 (right), a general 200×150 m network with 30 nodes (simulations).

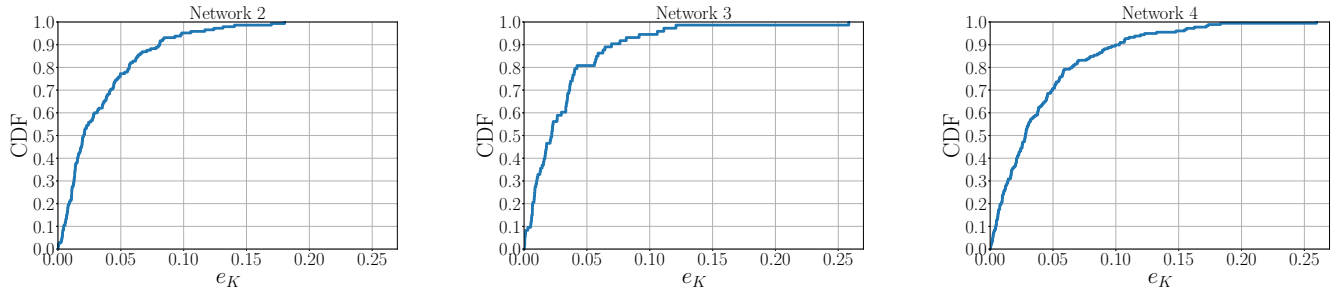


Fig. 4. Upper-bound e_K on the error made by using at most K paths when $M^{\text{opt}} > K$ for Network 2 (left), a general 100×100 m network with 15 nodes, Network 3 (center), a general 200×150 m network with 20 nodes, and Network 4 (right), a general 200×150 m network with 30 nodes (simulations).

the choice of the source and destination is made uniformly at random. If the network instance is not a multi-connected and multi-complete network, we remove the experiment (this occurs in 3% of the 1000 experiments). In 99.4% of the cases, the optimal number of paths is 3, which shows that in a multi-complete and multi-connected network, $M^{\text{opt}} = K$, as proven by Corollary 1. In a very few instances (0.2%), we find $M^{\text{opt}} = 4 > K$; these cases appear when the optimal multipath with minimal number of paths is not found. In a very few instances (0.4%), we find $M^{\text{opt}} = 2 < K$; these cases appear when the rate for a third path is below the threshold described in Section V-A, equal to $0.001 \cdot x^{\text{opt}}$.

We then study through simulations whether the analytical results presented in Section IV can be extended to more general results. We simulate three larger networks that are not necessarily multi-connected and multi-complete: Network 2, a 100×100 m square with 15 nodes; Network 3, a 200×150 m rectangle with 20 nodes; and Network 4, a 200×150 m rectangle with 30 nodes. In these larger networks, we simulate the fact that a PLC link exists only when two nodes are connected to the same *central coordinator* [1] (in particular, when two nodes are on the same electrical panel) by dividing the square in two equal parts, and by considering that, for one of the three technologies (the technology that simulates PLC), a link exists between two nodes only if the two nodes are in the same part. In particular, this means that the networks are never multi-connected. We simulate 1000 random instances of Network 2 and Network 4, and 1500 random instances of Network 3; there are more instances for Network 3 because the instances where there is no connectivity between the source

and destination nodes are more frequent. In total, there are 5% of instances with no connectivity between the source and destination nodes for Network 2, 39% for Network 3 and 17% for Network 4, and these instances are not included in the results. Figure 3 shows the cumulative distributive function of M^{opt} for respectively Network 2 (left), Network 3 (center) and Network 4 (right). Even if no theoretical result has been proven for this network, we see that the optimal number of hops M^{opt} remains tightly linked with the number of technologies K : In a large majority of the instances (respectively 85%, 92% and 79%), we have $M^{\text{opt}} \leq K$.

We finally study, for the instances where $M^{\text{opt}} > K$, the rate loss caused by using at most K paths. We do so by comparing the optimal rate x^{opt} obtained on an optimal multipath \mathcal{P}^{opt} , with the rate x_K obtained by computing the optimal rate on the K best paths in \mathcal{P}^{opt} . Note that x_K is only a lower bound on the optimal rate x_K^{opt} achieved with multipaths of K paths, as there is no guarantee that the optimal multipath with K paths contains only paths that belong to the optimal multipath \mathcal{P}^{opt} . As explained at the end of Section III, computing the actual optimal rate with multipaths of $M = K$ paths is NP-hard, and it cannot be computed in practice with up to several millions of possible paths, hence more than 10^{18} possible multipaths of $M = 3$ paths when $K = 3$. Using the K best paths of \mathcal{P}^{opt} is simple and enables us to find an upper bound $e_K \doteq \frac{x^{\text{opt}} - x_K}{x^{\text{opt}}}$ on the minimum relative error $e_K^{\text{opt}} \doteq \frac{x^{\text{opt}} - x_K^{\text{opt}}}{x^{\text{opt}}}$.

Figure 4 shows e_K for Network 2 (left), Network 3 (center) and Network 4 (right), in the (respectively) 17%, 9% and 21% of the instances where $M^{\text{opt}} > K = 3$. We see that the error made by using only K paths is very small: In (respectively)

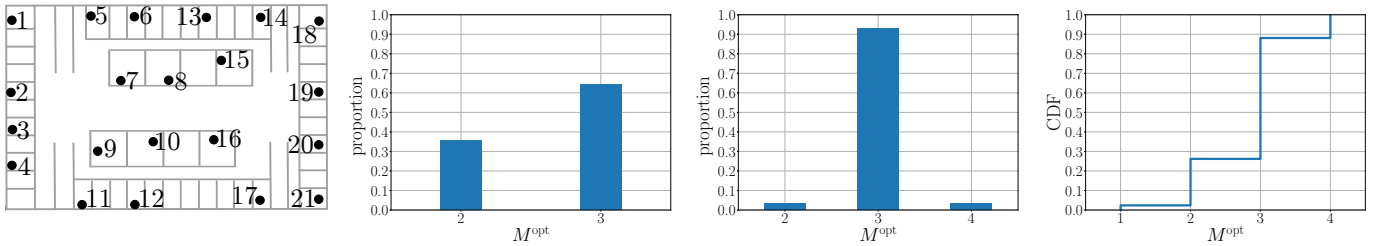


Fig. 5. Testbed of 65×40 m with PLC and two orthogonal WiFi channels (left). Proportion of each value of M^{opt} in a multi-complete but not multi-connected network (center left) and in a multi-complete and multi-connected network (center right). Cumulative distribution function of M^{opt} in a general 65×40 m network (right). Testbed experiments.

95%, 95% and 90% of the instances where $M^{\text{opt}} > K$, the relative error made by using only K paths is smaller than 0.1. Over all instances for each network, the relative error made by using only K paths is smaller than 0.1 in (respectively) 99.2%, 99.6% and 97.8% of the instances; and the relative error made by using only K paths is smaller than 0.05 in (respectively) 96.5%, 98.5% and 93.6% of the instances. Over all instances of Networks 1 to 4, the relative error made by using only K paths is smaller than 0.05 in 97.3% of the instances and smaller than 0.1 in 99.2% of the instances.

C. Experimental Results

We now present results obtained on a testbed of 21 nodes spread over an entire floor of an office building of 65×40 m (see Figure 5, left). All the nodes have two WiFi interfaces (Atheros AR9280), and a HomePlug AV PLC interface (QCA 7420) connected to the electrical network of the building. The nodes are APU1D boards running an OpenWrt Linux distribution with the open-source ath9k wireless drivers. The first WiFi channel is connected to a channel in the 2.4 GHz band, the second to a channel in the 5 GHz, consequently, they do not interfere. The PLC interface uses a Realtek Ethernet driver. We run our experiments at night to avoid external interference from the WiFi network of our university that operates in the 2.4 GHz band. To compute the interference domain \mathcal{I}_l of each link $l \in \mathcal{E}$, we run saturated traffic simultaneously on l and l' for each link $l' \neq l$, and we say that $l' \in \mathcal{I}_l$ if we observe a throughput degradation compared with the throughput when traffic is sent only on l .

We first carry experiments with Nodes 5 to 12 only. This network is multi-complete (for each technology, all links interfere with each other). Again, this does not mean that the network itself is complete (e.g., Node 5 cannot communicate directly with Node 11 with any technology). We start with a scenario where the PLC network is not multi-connected. This is achieved by setting logically two PLC networks with two different network management keys [1], one for Nodes 5 to 8, one for Nodes 9 to 12. Links in the two different PLC networks still interfere with each other. We choose randomly 28 different flows (i.e., source-destination pairs) and run the optimal back-pressure algorithm described in Section V-A. The measurements of the link capacities and interference domains take several hours, and the algorithm converges in about 20 minutes on average. Because the networks of each technology

are multi-complete, we expect that $M^{\text{opt}} \leq 3$ (Theorem 1); and because the two WiFi networks are multi-connected, we expect that $M^{\text{opt}} \geq 2$ (Theorem 2). Figure 5 (center left) shows that this is indeed the case.

Next, we connect Nodes 5 to 12 to the same logical PLC network, i.e., the network with $K = 3$ technologies is multi-connected. We choose randomly 32 different flows and run the optimal back-pressure algorithm. The proportion for each value of M^{opt} is shown in Figure 5 (center right). In more than 90% of the cases, $M^{\text{opt}} = 3$, as expected. In one case, we find $M^{\text{opt}} = 4$; this is because link capacities vary slightly, and the algorithm alternates between different paths that yield very close rates. In one case, $M^{\text{opt}} = 2$ because the capacities of two links of two different technologies are too close for a third path to exist.

Finally, we perform an experiment with the whole testbed (Nodes 1 to 21) that is neither multi-complete nor multi-connected: Nodes 1 to 12 and Nodes 13 to 21 are on two different electrical panels, i.e., on two different PLC networks, and the two PLC networks do not interfere with each other; also, WiFi links from one side (e.g., between Node 1 and Node 2) do not interfere with WiFi links from the other side (e.g., between Node 18 and Node 19). We choose randomly 42 flows. The cumulative distribution function of M^{opt} is shown in Figure 5 (right). Similarly to the simulations of Section V-B, we see that in most of the cases (about 90%), $M^{\text{opt}} \leq K$. In the remaining cases where $M^{\text{opt}} > K$, the relative error e_K made by using only K paths, as defined in Section V-B, is always below 0.1 (the maximum relative error is 0.08), and it is below 0.05 in 95.2% of the cases (all cases but two).

VI. CONCLUSION

We have presented analytical results that, for certain classes of mesh networks that include typical home networks, give bounds on the optimal number of paths when using multipath routing in hybrid networks with self-interfering technologies. They show that the optimal number of paths M^{opt} is tightly linked with the number K of non-interfering technologies. We have verified these analytical results with simulations and experiments on a three-technology testbed. We have also presented numerical and experimental results for more general networks. These results show that for general networks, the optimal number of paths M^{opt} remains close to the number of technologies K , and that the rate loss incurred by using at most

K paths is very small. This finding has a practical consequence of importance: It means that in home or enterprise networks with K distinct self-interfering technologies (e.g., PLC, WiFi with 2.4 GHz and 5 GHz, LTE), limiting a multipath-routing protocol to multipaths of at most K paths does not harm significantly the performance of the protocol.

REFERENCES

- [1] IEEE Standard for Broadband over Power Line Networks: MAC and PHY Layer Specifications. *IEEE Std 1901-2010*, 2010.
- [2] IEEE 1905.1-2013 Standard for a Convergent Digital Home Network for Heterogeneous Technologies. 2013.
- [3] J. N. Al-Karaki and A. E. Kamal. Routing Techniques in Wireless Sensor Networks: a Survey. *IEEE Wireless Communications*, 2004.
- [4] Brian X. Chen. Why Your Next Wi-Fi Setup Should Be a Mesh Network. *The New York Times*, April 2017.
- [5] Y.-C. Chen, Y.-s. Lim, R. J. Gibbens, E. M. Nahum, R. Khalili, and D. Towsley. A Measurement-based Study of MultiPath TCP Performance over Wireless Networks. In *ACM IMC*, 2013.
- [6] P. Dong, H. Qian, K. Zhou, W. Lu, and S. Lan. A Maximally Radio-disjoint Geographic Multipath Routing Protocol for MANET. *Annals of Telecommunications*, 2015.
- [7] A. Ford, C. Raiciu, M. Handley, S. Barré, and J. Iyengar. Architectural Guidelines for Multipath TCP Development. *RFC 6182*, 2011.
- [8] A. Ford, C. Raiciu, M. Handley, and O. Bonaventure. TCP Extensions for Multipath Operation with Multiple Addresses. *RFC 6824*, 2013.
- [9] J. J. Gálvez, P. M. Ruiz, and A. F. Skarmeta. Multipath Routing with Spatial Separation in Wireless Multi-hop Networks Without Location Information. *ACM Computer Networks*, 2011.
- [10] D. Ganesan, R. Govindan, S. Shenker, and D. Estrin. Highly-Resilient, Energy-Efficient Multipath routing in Wireless Sensor Networks. *ACM SIGMOBILE*, 2001.
- [11] S. Henri, C. Vlachou, J. Herzen, and P. Thiran. EMPoWER Hybrid Networks: Exploiting Multiple Paths over Wireless and ElectRical Mediums. In *ACM CoNEXT*, 2016.
- [12] K. Jain, J. Padhye, V. N. Padmanabhan, and L. Qiu. Impact of Interference on Multi-hop Wireless Network Performance. *Wireless Networks*, 2005.
- [13] G. Lee and J. Choi. A Survey of Multipath Routing for Traffic Engineering. *Information and Communications University, Korea*, 2002.
- [14] H. Lim, K. Xu, and M. Gerla. TCP Performance over Multipath Routing in Mobile Ad Hoc Networks. In *IEEE ICC*, 2003.
- [15] X. Lin and N. B. Shroff. Utility Maximization for Communication Networks with Multipath Routing. *IEEE Transactions on Automatic Control*, 2006.
- [16] W. Lou, W. Liu, and Y. Zhang. Performance Optimization using Multipath Routing in Mobile Ad Hoc and Wireless Sensor Networks. *Combinatorial Optimization in Communication Networks*, pages 117–146, 2006.
- [17] M. Neely, E. Modiano, and C.-P. Li. Fairness and Optimal Stochastic Control for Heterogeneous Networks. *IEEE/ACM Transactions on Networking*, 2008.
- [18] S. C. Nguyen, X. Zhang, T. M. T. Nguyen, and G. Pujolle. Evaluation of Throughput Optimization and Load Sharing of Multipath TCP in Heterogeneous Networks. In *IEEE WOCN*, 2011.
- [19] P. Pathak and R. Dutta. A Survey of Network Design Problems and Joint Design Approaches in Wireless Mesh Networks. *IEEE Communications Surveys & Tutorials*, 2011.
- [20] C. Raiciu, S. Barré, C. Pluntke, A. Greenhalgh, D. Wischik, and M. Handley. Improving Datascenter Performance and Robustness with Multipath TCP. In *ACM SIGCOMM*, 2011.
- [21] C. Raiciu, C. Paasch, S. Barré, A. Ford, M. Honda, F. Duchene, O. Bonaventure, and M. Handley. How Hard Can It Be? Designing and Implementing a Deployable Multipath TCP. In *USENIX NSDI*, 2012.
- [22] B. Recht, M. Fazel, and P. A. Parrilo. Guaranteed Minimum-rank Solutions of Linear Matrix Equations via Nuclear Norm Minimization. *SIAM review*, 2010.
- [23] M. Tarique, K. E. Tepe, S. Adibi, and S. Erfani. Survey of Multipath Routing Protocols for Mobile Ad-hoc Networks. *Journal of Network and Computer Applications*, 2009.

- [24] J.-Y. Teo, Y. Ha, and C.-K. Tham. Interference-Minimized Multipath Routing with Congestion Control in Wireless Sensor Network for High-Rate Streaming. *IEEE Transactions on Mobile Computing*, 2008.
- [25] J. Tsai and T. Moors. A review of Multipath Routing Protocols: From Wireless Ad Hoc to Mesh Networks. In *ACoRN ECR Workshop on Multihop Wireless Networks*, volume 30, 2006.
- [26] L. Zhou, X. Wang, W. Tu, G.-M. Muntean, and B. Geller. Distributed Scheduling Scheme for Video Streaming over Multi-channel Multi-radio Multi-hop Wireless Networks. *IEEE Journal on Selected Areas in Communications*, 2010.

APPENDIX

Before proving the theorems of Section IV, we define the following term.

Definition 3: Given a multipath $\mathcal{P} = (P_1, \dots, P_M)$, two multipath-impact vectors $\alpha_{\mathcal{P}, l_1} \in \mathbb{R}^M$ and $\alpha_{\mathcal{P}, l_2} \in \mathbb{R}^M$ are *link-independent* if there is no link of the multipath \mathcal{P} that interferes with both l_1 and l_2 , i.e., if $\mathcal{I}_{l_1} \cap \mathcal{I}_{l_2} \cap \Lambda_{\mathcal{P}} = \emptyset$.

Proof of Theorem 1: Given a multipath \mathcal{P} with M paths, if two links l and l' have same technology, then in a multi-complete network, $\alpha_{\mathcal{P}, l} = \alpha_{\mathcal{P}, l'}$ (all links of the technology interfere with each other). Consequently, the $L_{\mathcal{P}} \times M$ matrix $\mathbf{A}_{\mathcal{P}}$ given by (2) can be reduced to a $K \times M$ matrix $\tilde{\mathbf{A}}_{\mathcal{P}}$ without changing the solution of (3) ($\tilde{\mathbf{A}}_{\mathcal{P}}$ has one row per technology). If $M > K$, we next show that for each rate vector $\mathbf{x}_{\mathcal{P}} \in \mathbb{R}^M$, it is possible to build a rate vector $\mathbf{x}'_{\mathcal{P}} \in \mathbb{R}^M$ that uses $M - 1$ paths (i.e., there is an index i such that $x'_i = 0$) and such that $\mathbf{1}^T \cdot \mathbf{x}_{\mathcal{P}} \leq \mathbf{1}^T \cdot \mathbf{x}'_{\mathcal{P}}$, which proves the claim.

If there is an index i such that $x_i = 0$, then the result is trivially proven with $\mathbf{x}'_{\mathcal{P}} = \mathbf{x}_{\mathcal{P}}$. Let us now assume that $x_i > 0$ for all $1 \leq i \leq M$, and let $\mathbf{c}_1, \dots, \mathbf{c}_M \in \mathbb{R}^K$ be the columns of $\tilde{\mathbf{A}}_{\mathcal{P}}$. Then, (1) can be written as

$$\sum_{i=1}^M x_i \mathbf{c}_i = \boldsymbol{\mu}_{\mathbf{x}_{\mathcal{P}}}. \quad (6)$$

Because $M > K$ and $\mathbf{c}_i \in \mathbb{R}^K$ for $1 \leq i \leq M$, there is $(\gamma_1, \dots, \gamma_M) \neq (0, \dots, 0)$ such that

$$\sum_{i=1}^M \gamma_i \mathbf{c}_i = \mathbf{0}. \quad (7)$$

Because the \mathbf{c}_i only take positive values, some γ_i are positive and other are negative. Because for all $1 \leq i \leq M$, $x_i > 0$, $\epsilon^+ \doteq \min_{i \text{ s.t. } \gamma_i > 0} \frac{x_i}{\gamma_i} > 0$ and $\epsilon^- \doteq \max_{i \text{ s.t. } \gamma_i < 0} \frac{x_i}{\gamma_i} < 0$; let i^+ and i^- be the indices where the extremum is reached. For all $\epsilon \in [\epsilon^-, \epsilon^+]$ and all $1 \leq i \leq M$, $-x_i \leq \epsilon \gamma_i \leq x_i$; also, $x_{i^+} - \epsilon^+ \gamma_{i^+} = 0$ and $x_{i^-} - \epsilon^- \gamma_{i^-} = 0$. Clearly, due to (6) and (7), if $\mathbf{x}_{\mathcal{P}} = (x_1, \dots, x_M)$ is an admissible rate vector, then so is $\mathbf{x}'_{\mathcal{P}}(\epsilon) = (x_1 + \epsilon \gamma_1, \dots, x_M + \epsilon \gamma_M)$. Let $\sigma_{\gamma} = \gamma_1 + \dots + \gamma_M$. If $\sigma_{\gamma} \geq 0$, then $\mathbf{1}^T \cdot \mathbf{x}_{\mathcal{P}} \leq \mathbf{1}^T \cdot \mathbf{x}'_{\mathcal{P}}(\epsilon^+)$, whereas if $\sigma_{\gamma} \leq 0$, then $\mathbf{1}^T \cdot \mathbf{x}_{\mathcal{P}} \leq \mathbf{1}^T \cdot \mathbf{x}'_{\mathcal{P}}(\epsilon^-)$. Because $x'_{i^+}(\epsilon^+) = x'_{i^-}(\epsilon^-) = 0$, this proves the result with $\mathbf{x}'_{\mathcal{P}} = \mathbf{x}'_{\mathcal{P}}(\epsilon^+)$ if $\sigma_{\gamma} \geq 0$ and $\mathbf{x}'_{\mathcal{P}} = \mathbf{x}'_{\mathcal{P}}(\epsilon^-)$ if $\sigma_{\gamma} \leq 0$. ■

Proof of Lemma 1: To give an intuition, we begin by the proof for the particular case $M = 1$ (i.e., when the multipath \mathcal{P} is a single-path, denoted by P). If traffic is sent on P at an admissible rate x that saturates a technology on a link l_1 , then $\sum_{l' \in \mathcal{I}_{l_1} \cap \Lambda_P} x/c_{l'} = 1$. If l_2 is a link using

another technology, then $\mathcal{I}_{l_1} \cap \Lambda_P \neq \mathcal{I}_{l_2} \cap \Lambda_P$, i.e., from (5), $\sum_{l' \in \mathcal{I}_{l_2} \cap \Lambda_P} \frac{x}{c_{l'}} \neq \sum_{l' \in \mathcal{I}_{l_1} \cap \Lambda_P} \frac{x}{c_{l'}} = 1$. Because the rate x is admissible, $\sum_{l' \in \mathcal{I}_{l_2} \cap \Lambda_P} \frac{x}{c_{l'}} < 1$. This means that l_2 is not saturated.

Lemma 1 is trivial if $M = K$, hence we assume $M < K$. We move to $M \geq 2$. We first show the following two results.

Lemma 2: If S is a d -dimension affine subspace with $1 \leq d < M$, spanned by $d + 1$ link-independent multipath-impact vectors $\alpha_{\mathcal{P},l_1}, \dots, \alpha_{\mathcal{P},l_{d+1}}$, then a multipath-impact vector $\alpha_{\mathcal{P},l_0}$, link-independent with all the others, cannot belong to S , unless all vectors in S have the same $M - d$ entries that are zero, i.e., there are (at least) $M - d$ paths P_1, \dots, P_{M-d} such that for all i and j with $0 \leq i \leq d + 1$ and $1 \leq j \leq M - d$, $\alpha_{P_j, l_i} = 0$.

Proof: For $0 \leq i \leq d$, we denote by $\tilde{\alpha}_i \in \mathbb{R}^M$ the vector $\tilde{\alpha}_i = \alpha_{\mathcal{P},l_i} - \alpha_{\mathcal{P},l_{d+1}}$. The M elements of $\tilde{\alpha}_i$ are denoted by $\tilde{\alpha}_{ij}$ for $1 \leq j \leq M$. Because S is of dimension d , we know that the $\tilde{\alpha}_i$'s for $1 \leq i \leq d$ are linearly independent. The vector $\alpha_{\mathcal{P},l_0}$ belongs to S if and only if $\tilde{\alpha}_0$ is a linear combination of the $\tilde{\alpha}_i$'s for $1 \leq i \leq d$, i.e., if and only if there is $(\gamma_1, \dots, \gamma_d)$ such that

$$\tilde{\alpha}_0 = \sum_{i=1}^d \gamma_i \tilde{\alpha}_i. \quad (8)$$

With $\gamma = (\gamma_i)_{i \in \{1, \dots, d\}}$ and $\tilde{\mathbf{A}} \in \mathbb{R}^{d \times M}$ the matrix defined by $\tilde{\mathbf{A}} = [\tilde{\alpha}_1 \ \dots \ \tilde{\alpha}_d]^T$, this equation is equivalent to $\tilde{\mathbf{A}}^T \cdot \gamma = \tilde{\alpha}_0$. Let $k \in \{0, \dots, M\}$ be the number of paths such that for all $0 \leq i \leq d + 1$ and all $1 \leq j \leq k$, $\alpha_{P_j, l_i} = 0$ or, equivalently, $\tilde{\alpha}_{ij} = 0$. Without loss of generality, we assume that these k paths are the first k paths of \mathcal{P} . With respectively $\tilde{\mathbf{A}}_{M-k} \in \mathbb{R}^{d \times M-k}$ and $\tilde{\alpha}_{0, M-k} \in \mathbb{R}^{M-k}$ the restrictions of respectively $\tilde{\mathbf{A}}$ and $\tilde{\alpha}_0$ to their last $M - k$ columns (i.e., the multipath-impact vectors $\alpha_{\mathcal{P},l_i}$ are restricted to the last $M - k$ paths of \mathcal{P}), then (8) is equivalent to

$$\tilde{\mathbf{A}}_{M-k}^T \cdot \gamma = \tilde{\alpha}_{0, M-k}. \quad (9)$$

Let us assume that $k < M - d$, i.e., $M - k > d$. Then, the rank of $\tilde{\mathbf{A}}_{M-k}^T$ is d (because $\tilde{\alpha}_i$ for $1 \leq i \leq d$ are linearly independent), and $(\gamma_1, \dots, \gamma_d)$ is uniquely defined by d rows of $\tilde{\mathbf{A}}_{M-k}^T$. Without loss of generality, we assume that these d rows are the first d rows of $\tilde{\mathbf{A}}_{M-k}^T$. For (8) to be true, the equality (9) also needs to be verified for the last row of $\tilde{\mathbf{A}}_{M-k}^T$, i.e.,

$$\tilde{\alpha}_{0M} = \sum_{i=1}^d \gamma_i \tilde{\alpha}_{iM}. \quad (10)$$

But we know that there is $i \in \{0, \dots, d\}$ such that $\tilde{\alpha}_{iM} \neq 0$, i.e., $\mathcal{I}_{l_i} \cap \Lambda_{P_M} \neq \emptyset$. Modifying the capacity of one link of $\mathcal{I}_{l_i} \cap \Lambda_{P_M}$ by any small $\epsilon > 0$ changes the value of $\tilde{\alpha}_{iM}$ without changing the capacity of any other $\tilde{\alpha}_{ij}$ (because of the link-independence). In particular, it does not change the γ_i 's. This means that under Assumption 2, (10) cannot be true. Consequently, for $k < M - d$, there is no $(\gamma_1, \dots, \gamma_d)$ such that (8) is verified, i.e., $\alpha_{\mathcal{P},l_0}$ does not belong to S , which shows the claim. ■

Lemma 3: Let \mathcal{P} be a multipath with M paths, and $\mathbf{x}_{\mathcal{P}}$ an admissible rate vector. If $k < K$ links l_1, \dots, l_k that use k different technologies are such that the corresponding multipath-impact vectors $\alpha_{\mathcal{P},l_i}$ for $1 \leq i \leq k$ are linearly independent and verify $\alpha_{\mathcal{P},l_i} \cdot \mathbf{x}_{\mathcal{P}} = 1$ for $1 \leq i \leq k$ (i.e., the k links are saturated), then the multipath-impact vector $\alpha_{\mathcal{P},l_{k+1}}$ of link l_{k+1} using a technology different of the k others is either linearly independent with $\alpha_{\mathcal{P},l_1}, \dots, \alpha_{\mathcal{P},l_k}$, or it verifies $\alpha_{\mathcal{P},l_{k+1}} \cdot \mathbf{x}_{\mathcal{P}} \neq 1$.

Proof: Let us assume that neither of the two is verified, i.e., $\alpha_{\mathcal{P},l_{k+1}} \cdot \mathbf{x}_{\mathcal{P}} = 1$ and there are $(\gamma_1, \dots, \gamma_k)$ such that $\alpha_{\mathcal{P},l_{k+1}} = \sum_{i=1}^k \gamma_i \alpha_{\mathcal{P},l_i}$. Because for all $i \in \{1, \dots, k + 1\}$, $\alpha_{\mathcal{P},l_i} \cdot \mathbf{x}_{\mathcal{P}} = 1$, we have that $\sum_{i=1}^k \gamma_i = 1$, i.e., $\alpha_{\mathcal{P},l_{k+1}}$ belongs to the affine subspace of dimension $d = k - 1$ spanned by the multipath-impact vectors $\alpha_{\mathcal{P},l_i}$ for $1 \leq i \leq k$. But because they use different technologies, all the multipath-impact vectors $\alpha_{\mathcal{P},l_1}, \dots, \alpha_{\mathcal{P},l_{k+1}}$ are necessarily link-independent which, using Lemma 2, means that all the vectors $\alpha_{\mathcal{P},l_1}, \dots, \alpha_{\mathcal{P},l_k}$ must have $M - k + 1$ common zero-components, i.e., $\dim \ker(\alpha_{\mathcal{P},l_1}, \dots, \alpha_{\mathcal{P},l_k}) \geq M - k + 1$. Because the dimension of the space is M , the inequality is in contradiction with the fact that the vectors $\alpha_{\mathcal{P},l_1}, \dots, \alpha_{\mathcal{P},l_k}$ are linearly independent, i.e., $\text{rank}(\alpha_{\mathcal{P},l_1}, \dots, \alpha_{\mathcal{P},l_k}) = k$, which shows the result. ■

In particular, Lemma 3 shows that if we have $k = M$ links l_1, \dots, l_M of different technologies (the $\alpha_{\mathcal{P},l_i}$ are thus link-independent) that all are saturated, i.e., if $\alpha_{\mathcal{P},l_i} \cdot \mathbf{x}_{\mathcal{P}} = 1$ for $1 \leq i \leq M$, then they are necessarily linearly independent, which means that they form a basis of \mathbb{R}^M . Consequently, for any other link l_{M+1} using another technology ($\alpha_{\mathcal{P},l_{M+1}}$ is link-independent with the other multipath-impact vectors $\alpha_{\mathcal{P},l_i}$), $\alpha_{\mathcal{P},l_{M+1}}$ cannot be linearly independent with $\alpha_{\mathcal{P},l_1}, \dots, \alpha_{\mathcal{P},l_M}$, thus it must verify $\alpha_{\mathcal{P},l_{M+1}} \cdot \mathbf{x}_{\mathcal{P}} \neq 1$. This proves Lemma 1. ■

Proof of Theorem 2: Let \mathcal{P} be a multipath with $M < K$ paths, i.e., $M = K - k$ for some $k \in \{1, \dots, K - 1\}$. From Lemma 1, any rate vector $\mathbf{x}_{\mathcal{P}}$ saturates at most $K - k$ technologies. We therefore construct a K -path multipath that is strictly better by adding k paths using each only one of the (at least) k technologies that are not saturated (these paths exist because the network is multi-connected). Therefore, $M^{\text{opt}} \geq K$. ■