



HAL
open science

A stack-vector routing protocol for automatic tunneling

Mohamed Lamine Lamali, Simon Lassourreulle, Stephan Kunne, Johanne Cohen

► **To cite this version:**

Mohamed Lamine Lamali, Simon Lassourreulle, Stephan Kunne, Johanne Cohen. A stack-vector routing protocol for automatic tunneling. IEEE INFOCOM 2019 - IEEE Conference on Computer Communications, Apr 2019, Paris, France. pp.1675-1683, 10.1109/INFOCOM.2019.8737536 . hal-01987354

HAL Id: hal-01987354

<https://hal.science/hal-01987354>

Submitted on 23 Jan 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A stack-vector routing protocol for automatic tunneling

Mohamed Lamine Lamali*, Simon Lassourreille*, Stephan Kunne†, Johanne Cohen†

*LaBRI-CNRS. Université de Bordeaux. France.

†LRI-CNRS. Université Paris-Sud, Université Paris Saclay. France.

mohamed_lamine.lamali@u-bordeaux.fr johanne.cohen@lri.fr

Abstract—In a network, a tunnel is a part of a path where a protocol is encapsulated in another one. A tunnel starts with an encapsulation and ends with the corresponding decapsulation. Several tunnels can be nested at some stage, forming a protocol stack. Tunneling is very important nowadays and it is involved in several tasks: IPv4/IPv6 transition, VPNs, security (IPsec, onion routing), etc. However, tunnel establishment is mainly performed manually or by script, which present obvious scalability issues. Some works attempt to automate a part of the process (e.g., TSP, ISATAP, etc.). However, the determination of the tunnel(s) endpoints is not fully automated, especially in the case of an arbitrary number of nested tunnels. The lack of routing protocols performing automatic tunneling is due to the unavailability of path computation algorithms taking into account encapsulations and decapsulations. There is a polynomial centralized algorithm to perform the task. However, to the best of our knowledge, no fully distributed path computation algorithm is known. Here, we propose the first fully distributed algorithm for path computation with automatic tunneling, i.e., taking into account encapsulation, decapsulation and conversion of protocols. Our algorithm is a generalization of the distributed Bellman-Ford algorithm, where the distance vector is replaced by a protocol stack vector. This allows to know how to route a packet with some protocol stack. We prove that the messages size of our algorithm is polynomial, even if the shortest path can be of exponential length. We also prove that the algorithm converges after a polynomial number of steps in a synchronized setting. We adapt our algorithm into a *proto-protocol* for routing with automatic tunneling and we show its efficiency through simulations.

Keywords—Tunneling; encapsulation; path computation; routing protocol; distributed algorithms.

I. INTRODUCTION

Routing is one of the most important tasks in any network, and particularly in the Internet. Routing protocols are generally (distributed) versions of path computation algorithms in graphs. For instance, RIP relies on the distributed Bellman-Ford algorithm, while OSPF uses Dijkstra’s algorithm. These routing protocols were developed in the early years of the Internet. Thus, they work only in networks using the same *protocol*¹ (generally IP). However, Internet encompasses now several protocols. For instance, IPv4 and IPv6 coexist, and a lot of other situations involve interoperability of different protocols within the Internet. Thus, a path between a source and a destination may contain several portions using different protocols (for example, if the path crosses several domains

or Autonomous Systems). The mapping from a protocol to another one along a path is generally done in two ways: i) *Conversion*: a packet of some protocol is *converted* into a packet of another one (e.g., NAT-TP [1]); ii) *Encapsulation*: a packet of some protocol is *encapsulated* (or nested) in a packet of another one, thus being transparent to the crossed nodes until its *decapsulation* (the reverse operation). All these operations are called *adaptation functions* further in the paper.

A tunnel is a subpath beginning at an encapsulation and ending at the corresponding decapsulation. Tunneling is a ubiquitous feature in the Internet nowadays. It is involved in virtual networks, security (an encryption can be modeled as an encapsulation), interoperability, etc. Tunnels can be nested to achieve different tasks. In this context, topology connectivity is not enough to ensure communication. Even if the network is (strongly) connected, there may not exist a suitable adaptation function sequence to reach some destination from some node. If a path has such a suitable sequence, it is said to be *feasible*. However, current routing protocols are unable to handle protocol heterogeneity. Nowadays, the endpoints of tunnels are manually configured or determined from a precomputed list. Several works try to automate a part of the process (e.g., TSP [2], ISATAP [2], etc.). They are referenced as automatic tunneling. However, even if the negotiation of some parameters of the tunnel establishment are automated, such as tunnel type and DNS registration, the determination of the tunnel endpoints is not fully automated. Moreover, these approaches cannot handle an arbitrary number of nested tunnels. The main reason is that the underlying algorithms of routing protocols cannot handle adaptation functions. The authors of [3] propose a centralized polynomial algorithm to solve path computation problem in a multiple protocol context. However, their approach is based on language theory, and is unlikely to be distributed. To the best of our knowledge, no distributed algorithm for fully automated tunneling, i.e., with automatic determination of the (possibly nested) tunnel endpoints was proposed before this work.

Our goal is to design a distributed algorithm for path computation with automatic tunneling. This corresponds to the classical *All-Pairs Shortest Path* (APSP) problem, but in a network involving adaptation functions. To achieve this goal, we generalize the distributed Bellman-Ford algorithm, in order to take into account possible encapsulations, and to automatically establish (nested) tunnels. The *distance-vector* is replaced by a *stack-vector*. The route to follow for a packet is then determined by its destination and its protocol stack.

This paper will appear in the proceedings of IEEE INFOCOM 2019.

¹In this paper, *communication protocols* and *routing protocols* should not be confused. For simplicity, we refer to the first ones as *protocols*, and to the second ones explicitly as *routing protocols*.

Our contributions:

- 1) We design the first distributed algorithm for routing with automatic tunneling;
- 2) We prove nontrivial bounds on the maximum protocol stack height and on the algorithm convergence: we show that the maximum protocol stack height (number of encapsulated protocols at the same time) of the shortest path (involving tunnels) between two nodes is at most λn^2 , where n is the number of nodes and λ is the number of protocols in the network. This implies that the maximum message size is also polynomial, despite the fact that the shortest path may be of superpolynomial length in this context;
- 3) We design a *proto*-protocol for routing with automatic tunneling and we evaluate its efficiency through simulations.

The paper is organized as follows: Section II describes the problem and discusses the related work. Section III details the model used through this paper and formalizes the notion of path feasibility. Section IV describes the proposed algorithm and study its convergence and its message size. Section V describes a basic implementation of a routing protocol based on our algorithm, while Section VI presents the simulation results. Finally, Section VII concludes the paper.

II. THE PROBLEM

A. Problem definition and illustration

We aim to illustrate the path computation problem through an example. Figure 1 depicts a network with 8 nodes². We want to compute a path from node S to node D . Both of them use protocol a . However, the nodes connecting S and D use different protocols. And some nodes are able to map a protocol to another one. A packet of any protocol comprises a *header* field and a *data* field. The structure of the header field is specific to a protocol. Converting a packet of protocol a to protocol b consists in rewriting the header and converting all the information (for example, the node identifiers, the routing options, etc.). Encapsulating a packet of a in b consists in considering the whole packet (header and data) of a as the data field of b , and then adding a header of protocol b at the top of the packet. For example, node U_1 is only able to encapsulate (a packet of) protocol a in protocol b , while node U_2 is only able to receive and send protocol b without any change of the possible encapsulated protocols (which are transparent to it). The capabilities (adaptation functions) of each node are listed above it. Each step of the path is associated with a *protocol stack*, i.e., the current sequence of encapsulated protocols, the current protocol being at the top of the stack. In Figure 1a, node S emits (a packet of) protocol a , so at this stage the protocol stack contains only a . When the packet crosses node U_1 it is encapsulated in protocol b . Thus, the current protocol stack is ab (from bottom to top). At U_2 the protocol stack remains the same, since node U_2 passively retransmits packets

with current protocol b . Node U_4 then decapsulates protocol a from b , thus the current stack is again a . Node U_5 receives protocol a and it is only able to decapsulate protocol b from a ; however, there is no b encapsulated in a at this stage. This path is said to be *not feasible*. The same problem would appear if a node receives a protocol while it is only able to handle another one. In contrast, a path is *feasible* if at each stage, the current node can handle the current protocol stack, and the destination node receives a packet with no encapsulated protocols. Figure 1b illustrates a feasible path that involves a loop. This loop is necessary to collect the required encapsulations to cross the subpath $U_4 \dots U_6$. This example illustrates some specificities of feasible paths. There is not always a feasible path between two nodes, even if the graph is connected. Moreover, the shortest feasible path (if any) can involve loops. And optimal feasible paths do not exhibit optimal substructure (a subpath of a feasible path is not necessarily feasible, since the protocol stacks at the start and the end of the subpath may be different). In practice, the protocol stack is the sequence of nested headers of a packet that underwent multiple encapsulations.

The problem we address is to compute the shortest feasible path (if any) between each pair of nodes in a distributed way, i.e., the *All Pairs Shortest paths* (APSP) problem in a network involving adaptation functions. Our algorithm builds routing tables allowing a packet with some protocol stack to reach the destination through the shortest feasible path. The main motivation of this work is to propose a generic algorithm and a basic protocol for routing. Given a network with different protocols and some adaptation functions on some nodes, the goal is to compute the shortest path between any pair of nodes. Tunnels should be automatically established during the routing table construction. This can have important applications on several fields: IPv4/IPv6 interoperability, VPNs, secured tunneling, emulation of lower layer protocols over IP [4], etc.

Note that our problem is different from path computation in multi-layer networks where the number of layers is fixed [5], while our problem is to resolve the APSP in a network where the number of needed layers is unknown, the layers being induced by the node encapsulation and decapsulation capabilities. Moreover, in our setting, even two layers with the same stack height may not be comparable, since two stacks such as abb or aba cannot be routed in the same way, for example.

B. Related Work

1) *Centralized approach:* In [6], The authors prove that the path feasibility problem under bandwidth constraint is NP-hard. They propose an exponential algorithm using a breadth-first search approach to solve the problem. The authors of [7] prove that the problem is polynomial³ if the bandwidth constraint is relaxed. They model the network as a Push-Down Automaton (PDA), and use tools from language theory to find the shortest feasible path between two nodes. In [3], they widely generalize the previous algorithms, and prove that the feasible path problem under bandwidth constraint remains NP-hard even with only two protocols and on symmetric directed

²Note that, in this example, the underlying topology is not a symmetric directed graph, while in our model (see Section III) the directed graph must be symmetric. In this context, a symmetric directed graph is defined as a graph where a link (U, V) exists if and only if the reverse link (V, U) exists. However, the example on Figure 1 is simplified for readability purpose. It can be easily converted into a symmetric directed graph with the same properties.

³More precisely, the decision version, i.e., deciding if a feasible path exists between two nodes, is polynomial. However, as stated in Section IV-A, explicitly computing the path can be superpolynomial.

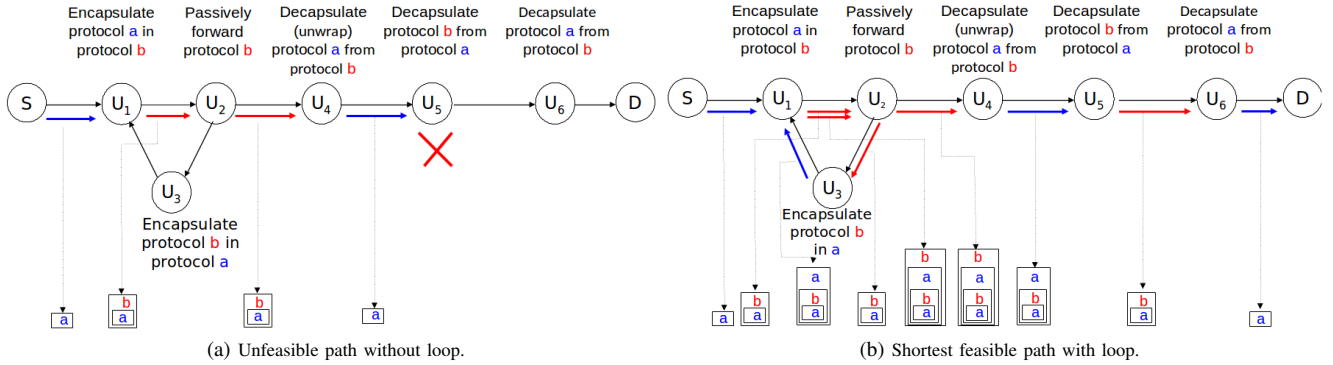


Fig. 1: A network requiring a shortest path with loop.

graphs. The authors of [8] propose a matrix based model and an algorithm to compute feasible paths. However, the algorithm takes exponential time, and can only compute loopless feasible paths. Unfortunately, the language theoretic approach seems unlikely to be distributed. Some works have been devoted to distributed PDAs (see [9] for example). However, in these works, the nodes are themselves PDAs. In contrast, our model is equivalent to a PDA where the states collaborate in a distributed way to find an accepting transition sequence.

2) *Distributed algorithm related work:* In distributed literature, adaptation functions are generally not considered. The distributed Bellman-Ford algorithm (on which relies RIP) solves the APSP problem in $O(n)$ rounds in a synchronized setting if the message size is unbounded (e.g., the LOCAL model [10]). However, the APSP problem mostly received attention in the CONGEST model, where the message size is restricted to be logarithmic [11], [12]. Unfortunately, these works cannot apply in our context since they do not take into account the adaptation functions and the interoperability issues.

3) *Networking related work:* Feasible path computation is a major challenge in networking. It underlies several technical problems: establishing tunnels, unifying control planes, etc. IPsec is a classical example. The endpoints of IPsec tunnels are often manually determined. The Tunnel Endpoints Discovery mechanism [13] allows to automatically determine IPsec tunnel endpoints. However, this can be only done through a single path. This path corresponds to the shortest path without encryption. IPv4/IPv6 interoperability is mainly achieved through conversion of protocols or by tunneling. In the first case, an IPv4 (resp. IPv6) header is converted into an IPv6 (resp. IPv4) header. This operation is costly, and it requires to know the nodes able to perform the conversion. NAT-TP [1] uses this approach. Some other mechanisms (6over4 [14], 6to4 [15], etc.) use tunneling. 6to4 is referenced as automatic tunneling. Any 6to4 encapsulation capable router has an anycast address [16]. Discovering the tunnel endpoints is done by sending a packet to the anycast address. However, the end-hosts should know that they need a tunnel. Moreover, this mechanism is specific to IPv4/IPv6 interoperability, and cannot handle nested tunnels with arbitrary protocols. While our goal is to design a generic protocol without these limitations. Some other works attempt to automate a part of the process. For example, the TSP protocol [2] proposes the (automatic) negotiation of

the tunnel parameters (e.g., keep alive duration) between the endpoints. However, no algorithm or process is provided to select the optimal (and may be nested) tunnel endpoints. The ISATAP protocol [17] proposes to select the endpoints from a predetermined list via a tunnel Broker. But, again, no algorithm is provided to automatically compute the tunnel endpoints. Moreover, these mechanisms are unable to handle an arbitrary number of nested tunnels. Fully automatic tunneling is still an open problem. The RFC 7059 [18] proposes a comparison between the different IPv4/IPv6 tunneling mechanisms.

C. Our approach

We propose to generalize the distributed Bellman-Ford algorithm. Since the path to follow by a packet depends upon the destination, but also upon the packet protocol stack, the distance-vector should be replaced by a stack-vector. Each node advises its neighbors that it can reach some destination, with some protocol stack, and at some cost. The rows of the routing table of each node should contain a next hop per destination and per protocol stack. The termination of such an algorithm is an important issue. Since feasible paths may involve loops, the same destination can be reached with different protocol stack heights. The question is to know at what stack height to stop. We resolve the problem by showing a bound on the stack height of any shortest feasible path. The convergence speed of the algorithm depends upon the length of the shortest feasible paths. Our algorithm is generic: it can handle any number of protocols, and it does not limit the number of nested tunnels to find a feasible path.

III. MODEL AND DEFINITIONS

This section describes the mathematical model used in this paper and formalizes the notion of path feasibility. We use the same notations and definitions as in [3].

A. Network model

A network is modeled as a 4-tuple $\mathcal{N} = (\mathcal{G}, \mathcal{A}, \mathcal{F}, w)$. The network topology is represented by a symmetric directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. Each node in \mathcal{V} corresponds to a router, and each pair $(U, V) \in \mathcal{E}$ is a unidirectional link. The number of nodes is denoted by $|\mathcal{V}| = n$, and the number of links is denoted by $|\mathcal{E}| = m$. The set of protocols available in the network

is denoted by $\mathcal{A} = \{a, b, \dots\}$. Its cardinality is denoted by $|\mathcal{A}| = \lambda$. Each node $U \in \mathcal{V}$ has a set $\mathcal{F}(U)$ of adaptation functions that it can perform. These functions are:

- *Conversion*: A protocol a is converted into a protocol b without any change of the possible encapsulated protocols. It is denoted by $(a \rightarrow b)$ (example: IPv4/IPv6 conversion through NAT-PT). Observe that a classical retransmission without protocol change is a special case of conversion. It is denoted by $(a \rightarrow a)$.
- *Encapsulation*: A protocol a is encapsulated in b . It is denoted by $(a \rightarrow ab)$ (e.g., IPv4/IPv6 encapsulation).
- *Decapsulation*: A protocol a is decapsulated from b . It is denoted by $(a \rightarrow \overline{ab})$.

The set of all adaptation functions available in the network is denoted by \mathcal{F} . We denote the set of protocols that a node can receive by $In(U)$, and the set of protocols that a node can send as $Out(U)$.

Finally, performing an adaptation function on a node U has a cost defined by the function $w : \mathcal{V} \times \mathcal{F} \times \mathcal{V} \rightarrow \mathbb{R}_+$. The value $w(U, f, V)$ (where $U, V \in \mathcal{V}$ and $f \in \mathcal{F}(U)$) corresponds to the cost of using link (U, V) with adaptation function f on node U . Hence, function w represents any additive metric associated with both links and adaptation functions. So the cost of a path is the sum of the costs of each triple (U, f, V) involved in it in the network. For example, if one wants to minimize the number of encapsulations in the path, the cost function should be $w(U, f, V) = 1$ if f is an encapsulation and 0 otherwise.

B. The protocol stack

A sequence of adaptation functions induces a protocol stack. For example, the sequence $(a \rightarrow a)(a \rightarrow ab)(b \rightarrow b)$ induces the stack ab (from bottom to top). For each position i in a path, H_i denotes the protocol stack at this position, i.e., the protocol stack induced by $f_0 \dots f_i$, and h_i denotes the protocol stack height. The protocol at the top of a stack H is denoted by $Top(H)$ and the protocol just below $Top(H)$ in the stack is denoted by $Top^{-1}(H)$. The height of a stack H is denoted by $h(H)$. The “forbidden” stack is denoted by \emptyset (note that it should not be confused with the empty word ϵ).

More formally, let f be an adaptation function, and let H be a stack and $H = H'.Top(H)$ ⁴, where H' is eventually empty. We will also denote by f the function taking as argument a stack and performing the adaptation function on this stack:

- if $f = (x \rightarrow y)$ and $Top(H) = x$, then $f(H) = H'.y$
- if $f = (x \rightarrow xy)$ and $Top(H) = x$, then $f(H) = H.y$
- if $f = (\overline{(x \rightarrow xy)})$ and $Top(H) = y$ and $Top^{-1}(H) = x$, then $f(H) = H'$
- $f(H) = \emptyset$ otherwise.

⁴The notation “.” stands for a simple concatenation. For example, if $H = abab$ then $H.b = ababb$.

We also denote by \bar{f} the reverse function of f , i.e., if $H' = f(H)$ and $H' \neq \emptyset$ then $\bar{f}(H') = H$. Note that if f is an encapsulation, then \bar{f} is the corresponding decapsulation. $f(H) = \emptyset$ means that the adaptation function f cannot handle the stack H (e.g., f encapsulates a in b while $Top(H) \neq a$). Note that $f(\emptyset) = \emptyset$ for any function in our context.

Thus, the protocol stack H_i induced by a sequence of adaptation functions $f_0 \dots f_i$, is recursively defined as following:

- $H_0 = x$ if $f_0 = (x \rightarrow x)$ and $H_i = f_i(H_{i-1})$

C. Path feasibility

In our context, a path should take into account the adaptation function capabilities of the network nodes. A path should contain the list of adaptation functions involved in it. Thus, a path from node S to node D in a network \mathcal{N} is a sequence of nodes and adaptation functions $Sf_0U_1f_1U_2f_2 \dots U_\ell f_\ell D$ where each U_i , $i = 1, \dots, \ell$, is a node, and each f_i is an adaptation function⁵. When needed, a directed path in the graph without taking into account the adaptation functions (i.e., just a sequence of nodes) is referred to as a *classical path*.

Definition 1: A path $\mathcal{P} = Sf_0U_1f_1U_2f_2 \dots U_\ell f_\ell D$ is *feasible* if and only if:

- 1) The sequence $SU_1U_2 \dots U_\ell D$ is a classical path in $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ and each $f_i \in \mathcal{F}(U_i)$;
- 2) $H_\ell = x$ and $x \in In(D)$.

Actually, the protocol sequences of feasible paths can be characterized as a context-free language [7]. In Figure 1a, the depicted path is:

$$S(a \rightarrow a)U_1(a \rightarrow ab)U_2(b \rightarrow b)U_4(\overline{(a \rightarrow ab)})U_5(\overline{(b \rightarrow ba)})U_6fD.$$

It cannot be feasible for any adaptation function $f \in \mathcal{F}(U_6)$ since $(a \rightarrow ab)$ appears before any encapsulation of protocol b . In contrast, Figure 1b depicts the feasible path:

$$S(a \rightarrow a)U_1(a \rightarrow ab)U_2(b \rightarrow b)U_3(b \rightarrow ba)U_1(a \rightarrow ab)U_2(b \rightarrow b)U_4(\overline{(a \rightarrow ab)})U_5(\overline{(b \rightarrow ba)})U_6(\overline{(a \rightarrow ab)})D.$$

The corresponding protocol stacks are below the links in Figure 1.

D. Problem formalization

The cost of a path $\mathcal{P} = Sf_0U_1f_1U_2f_2 \dots U_\ell f_\ell D$ from node S to node D is defined as $w(\mathcal{P}) \stackrel{def}{=} \sum_{i=0}^{\ell} w(U_i, f_i, U_{i+1})$ where $S = U_0$ and $D = U_{\ell+1}$. An optimal (i.e., shortest) feasible path \mathcal{P} between two nodes is a feasible path that minimizes $w(\mathcal{P})$. When needed, we refer to the number of hops in a path \mathcal{P} by $|\mathcal{P}|$. Note that $|\mathcal{P}|$ may be different from $w(\mathcal{P})$.

The problem we aim to solve is the following: compute the routing tables at each node, such that for each pair of nodes U and V , a packet from U to V following these routing tables follows an optimal path from U to V . This corresponds to an APSP in a network involving adaptation functions. Note that there is not always a feasible path between two nodes even if the underlying graph \mathcal{G} is strongly connected.

⁵The adaptation function f_0 is dummy. If node S emits packets of protocol a , then f_0 is denoted by $(a \rightarrow a)$ by convention.

IV. STACK-VECTOR ROUTING ALGORITHM

A. Different approaches to distributed routing

There are two main approaches in distributed routing: link-state routing and distance/path vector routing. The first one consists in spreading the whole topology then computing locally the path in a centralized way. However, this approach has a drawback: computing the whole shortest path at a node is not enough to route the packets. A node should know the next neighbor at each stage of the route. This is not possible unless the header of the packet contains the whole path to follow. The second approach is to use a Bellman-Ford algorithm (as RIP) where each node shares its current routing table with its neighbors. We cannot use this approach since the path leading to the destination depends upon the protocol stack of the packet to route. Storing the cost and the destination in the routing table is not enough.

Path vector routing consists in sharing not only the destination and the cost, but also the whole path to reach the destination. The main goal to do so is to avoid loops (if a node receives a path where it already appears, then it discards it). BGP uses this approach. This would be a possible solution for our problem, since a complete feasible path (the node sequence and the adaptation function sequence) characterizes the route to take for a packet. However, the shortest feasible path can be superpolynomial [3]. A lower bound of $\Omega((n/\lambda)^\lambda)$ is known for an arbitrary number of protocols λ . Even for two protocols, there are shortest paths of length at least $\Omega(\sqrt{n}2^{\Omega(\sqrt{n})})$. The best available upper bound is $2^{O(\lambda^2 n^2)} - 1$ (see [3] for more details). Using a path vector protocol would lead to a superpolynomial message length. We opted for a *stack vector algorithm* (by analogy with distance vector protocols), where the destination, the cost and the protocol stack are stored in the routing table.

B. The algorithm

The initialization algorithm (Algorithm 1) allows each node U to share with all its neighbors $Nei(U)$ the set of protocols $In(U)$ that it can receive. Sending a message (U, H, c) means that the sender can reach destination U at cost c if it receives a packet with protocol stack H . Thus, the initialization phase consists of informing its neighbors that it can receive any packet with protocol $x \in In(U)$ (without nested encapsulations) at cost 0.

Algorithm 1 Initialization algorithm of node U

```

1: for all  $x \in In(U)$  do
2:   for all  $V \in Nei(U)$  do
3:      $H \leftarrow x$ 
4:      $Send(U, H, 0)$  to  $V$ 

```

Each row of the routing table \mathcal{T} of each node is a 5-tuple (D, H, c, V, f) where D is the destination to reach, H is the protocol stack needed to reach D , c is the cost of the remaining path to reach the destination, node V is the next neighbor (next hop) to reach the destination, and f is the adaptation function to perform at the current node. The row is indexed by the pair (D, H) . Thus, $\mathcal{T}(D, H)$ returns the row corresponding to destination D , and stack H . When a node tries to add a new row (D, H, c, V, f) to its routing table \mathcal{T} , it checks if the

tuple (D, H) is already in its routing table. If it is not, then it inserts the new row to the table. Otherwise, it compares the cost of the new route with the old one. The table is updated by replacing the old route by the new one if the new cost is lower. This step is done according to Algorithm 2.

Algorithm 2 Add a row to a routing table \mathcal{T} of node U

```

Input: A row  $(D, H, c, V, f)$ 
1: if  $(D, H) \notin \mathcal{T}$  then
2:   Add  $(D, H, V, c, f)$  to  $\mathcal{T}$ 
3: else if  $\mathcal{T}(D, H).cost > c$  then
4:    $\mathcal{T}(D, H).cost \leftarrow c$ 
5:    $\mathcal{T}(D, H).next\_hop \leftarrow V$ 
6:    $\mathcal{T}(D, H).function \leftarrow f$ 

```

Algorithm 3 is the main algorithm. When receiving a new message (D, H, c) from a neighbor V , node U determines which of its adaptation functions it can apply to stack H (lines 3-5). Then it computes the new cost of the remaining path by adding the cost of (U, f, V) to the old cost (line 6). It tries to add the new row corresponding to the message and the chosen adaptation function to its routing table \mathcal{T} , according to Algorithm 2 (line 7). If the row corresponds to a new route in the table, then it sends the new message (D, H, c) to its neighbors, indicating that it can reach the destination with the new stack H . Observe that if f is a decapsulation ($a \rightarrow ab$), and U can reach D with some protocol stack H such that $Top(H) = a$, then U can handle a packet with stack $H.b$ to reach D . Similarly, if f is a conversion ($a \rightarrow b$), and if U can reach D with some protocol stack H such that $Top(H) = b$ (i.e., $H = H'.b$ for some stack H'), then U can handle a packet with stack $H'a$ to reach D . Thus, if node U can reach a destination D with some stack H , for each $f \in \mathcal{F}(U)$, it should apply the reverse adaptation function \bar{f} (line 4) to H before sending a message to the other nodes. Note that only stacks with heights less than or equal to λn^2 are kept and shared (line 5). In the next section, we prove that this height is sufficient to reach any destination with a shortest feasible path if there exists one.

Algorithm 3 Routing table construction algorithm of node U

```

1: loop
2:   Receive  $(D, H, c)$  from  $V$ 
3:   for all  $f \in \mathcal{F}(U)$  do
4:      $H \leftarrow \bar{f}(H)$ 
5:     if  $H \neq \emptyset$  and  $h(H) \leq \lambda n^2$  then
6:        $c \leftarrow c + w(U, f, V)$ 
7:       Add row  $(D, H, c, V, f)$  to  $\mathcal{T}$ 
8:       # according to Algorithm 2
9:       if  $\mathcal{T}$  has been modified then
10:        for all  $W \in Nei(U)$  do
11:          Send  $(D, H, c)$  to  $W$ 

```

We assume that the packets to route are of the form (destination, protocol stack, payload); for example, $(D, H, data)$. Once all routing tables are computed, if node U receives a packet with protocol stack H and destination D , it first searches for tuple (D, H) in its routing table. If there is no corresponding row, then node U has no route for the destination with the received protocol stack. Otherwise, if the

corresponding row is (D, H, c, V, f) , then it sends the packet $(D, f(H), data)$ to V . The routing procedure is illustrated by Algorithm 4.

Algorithm 4 Routing a packet.

- 1: Receive a packet $(D, H, data)$
 - 2: **if** $(D, H) \notin \mathcal{T}$ **then**
 - 3: No route, discard the packet
 - 4: **else**
 - 5: Let $(D, H, c, V, f) \in \mathcal{T}$
 - 6: Send the packet $(D, f(H), data)$ to V .
-

C. Correctness and complexity

1) *Stack height upper bound and message size:* First, we prove that for any shortest feasible path, the maximum stack height reached along the path is polynomially bounded.

Lemma 1: The shortest feasible path (if any) between two nodes reaches a maximum stack height of at most λn^2 protocols.

Proof: See Appendix A. ■

Corollary 1: The maximum message size is in $O(\lambda \log \lambda n^2)$.

Proof: A message contains node identifiers, that are in $O(\log n)$ if the number of nodes is n . A protocol stack, its height is at most λn^2 according to Lemma 1. The identifier of a protocol is in $O(\log \lambda)$ for a number λ of protocols. We assume that the cost is bounded. ■

2) *Convergence and correctness:* We will prove that in a synchronized setting (for example the LOCAL model in distributed computing [10]) where all nodes receive messages at the same time t , then process them and send them at the same time $t+1$, Algorithm 3 correctly converges in polynomial time⁶ according to the network size and its diameter.

We define the diameter of a network \mathcal{N} , denoted by $\text{diam } \mathcal{N}$, as the length (in number of hops) of the shortest feasible path that maximizes the number of hops:

$$\text{diam } \mathcal{N} \stackrel{\text{def}}{=} \max_{\mathcal{P} \text{ shortest feasible}} |\mathcal{P}|.$$

Proposition 1: Algorithm 3 computes the correct routing tables after $O(\lambda n^2 \text{diam } \mathcal{N})$ rounds.

Proof: See Appendix B. ■

V. A PROTO-PROTOCOL

A. Practical limitations and algorithm adaptation

The main issue for the implementation of the proposed algorithm is that the stack height and the message length bounds are too large, even if they are polynomial. For example, in a network of 100 nodes with only 2 protocols, the stack height may reach 2×10^4 protocols. This is not sustainable for real applications. Moreover, for a stack height h , there are λ^h possible stacks, i.e., in the worst case, there may be $n \left(\frac{1 - \lambda^{\lambda n^2 + 1}}{1 - \lambda} - 1 \right)$ rows in a routing table. This may induce

⁶Time is measured in number of synchronized rounds.

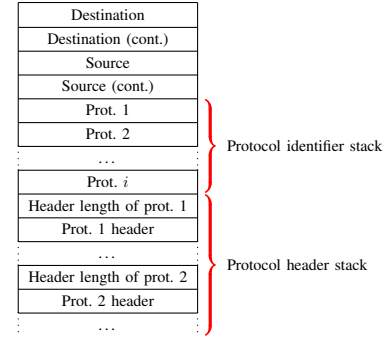


Fig. 2: A header of a packet to route.

an exponential number of exchanged messages. Finally, as demonstrated by the authors of [3], $\text{diam } \mathcal{N}$ can be super-polynomial in the size of the network \mathcal{N} . These bounds are tight, i.e., it is possible to exhibit a network where the shortest feasible path reaches a stack height of λn^2 . Thus, these limitations are not due to our algorithm but are inherent to the problem.

However, the simulations performed in [3] show that such instances are extremely unlikely to appear. Thus, we propose to set the maximum stack height as a parameter of the protocol. A small value of this parameter is enough to compute the shortest path in most cases. We should bound the maximum stack height by a constant h_{\max} for any real implementation. Thus, the condition $h(H) \leq \lambda n^2$ in Algorithm 3 (line 5) should be replaced by $h(H) \leq h_{\max}$. In such a case, the maximum number of different stacks of height h_{\max} would be $\lambda^{h_{\max}}$, and the maximum number of rows in a routing table would be $n \left(\frac{1 - \lambda^{h_{\max} + 1}}{1 - \lambda} - 1 \right)$, which is linear in n and polynomial in λ .

B. Routing proto-protocol specification

The main requirements are that the nodes and the protocols must have a unique identifier. For example, 1 byte for the protocol identifiers and 16 bytes for the node identifiers. The routers involved in the routing protocol should have a specific multicast address. The routing protocol messages that advise a route must contain: the destination of the route, the protocol stack needed for a packet in order to reach the destination, and the cost of the route. In addition, some classical parameters can be exchanged: keep alive duration, emitted message timestamps, etc. The maximum stack height should be set at the start of the process. The routing table must contain entries indexed by the destination and the protocol stack of a received packet. Each entry must contain the cost, the next hop, and the adaptation function to perform before sending the packet to the next hop.

C. Routing a packet

The packet to route contains a stack of headers of different protocols. It can be seen as a meta-header. It must contain:

- The unique identifier of the destination node;
- The source (even if it does not impact the routing process, since it is a per destination/stack routing);

- The stack height: the current stack height of the packet
- The protocol stack: the stack of identifiers of the protocols corresponding to the nested headers. This will speed up the routing process, since it avoids to access to the whole header stack;
- The header stack: each encapsulated header preceded by its length;
- The payload of the inner packet.

Figure 2 illustrates such a meta-header.

VI. SIMULATIONS

In order to evaluate the efficiency of our algorithm, we performed simulations with different parameters.

A. Simulation methodology

All the networks used in the simulations are generated according to the following steps:

- 1) We generate a random undirected graph of a given size according to a preferential attachment mechanism (the Barabási-Albert model [19]), where each added node is attached to 5 existing nodes;
- 2) The graph is then converted into a symmetric directed graph. Each undirected link is converted into two directed links;
- 3) For a given number λ of protocols, there are $3\lambda^2$ possible adaptation functions. Each adaptation function is available on a node with a given probability p .

The algorithm is implemented in Python 3.4.5, using the NetworkX package⁷. The implementation is done in a distributed fashion: each node is simulated by a thread, and a directed link (U, V) is implemented as a queue where U can only write, and V can only read. The simulations were performed on a multi-core server with CPUs 1.59GHz.

The input parameters of the algorithm are: the number of nodes, the probability p of availability of an adaptation function, the number of protocols, and the maximum stack height h_{\max} . The main output results are the convergence time, and the percentage of times where the algorithm finds the shortest path between the network extremities. Note that, if the algorithm does not find the shortest feasible path, it may be because there is no feasible path in the given network. The probability of existence of a feasible path according to different parameters can be found in [3]. All the results are averaged over 1000 runs.

B. Convergence time

Figure 3 shows the convergence time of the algorithm according to different parameters. Figure 3a shows the convergence time according to the network size (number of nodes) in three cases: $p = 0.05$, $p = 0.10$, and $p = 0.20$. The maximum stack height h_{\max} is set to 3, and the number of protocols is set to 2. It appears that the convergence time hugely depends on the parameter p . For $p = 0.05$, the convergence time is 2.74s in a network of 50 nodes; while for $p = 0.20$, the convergence

time is 113s. However, note that if $p = 0.20$, the average number of adaptation functions per node is 2.4, which is unrealistic, since only a few number of nodes should be able to perform conversions/encapsulations. Actually, Figure 3b shows the impact of the probability p (and the average number of adaptation functions per node) on the convergence time (with $h_{\max} = 3$ and 2 protocols). For $p = 0.05$, the convergence time is around 0.7s (resp. around 3s) in a network of 30 (resp. 50) nodes; while for $p = 0.30$, the convergence time is around 50s (resp. around 3min) in a network of 30 (resp. 50) nodes. We can see that the processing time hugely depends upon the number of adaptation functions per node. Figure 3c shows the impact of the parameter h_{\max} on the convergence time. The number of protocols is set to 2 and the network has 50 nodes. Note the logarithmic scale on the y -axis. The maximum stack height hugely impacts the convergence time. With $p = 0.05$ (resp. $p = 0.20$), the convergence time is 0.7s (resp. around 30s) if the maximum stack height is set to 2. However, the convergence time is around 21s (resp. 11min) if the maximum stack height is set to 5.

C. Algorithm efficiency

Figure 4 shows the average number of times where Algorithm 3 found a shortest feasible path between the network extremities, according to different parameters. For each run, we check if a shortest feasible path is computed between the network extremities. Note that, if there is no path, it can be due to two different reasons: i) there is no feasible path between these nodes; ii) the shortest feasible path reaches a stack height larger than h_{\max} .

Figure 4a shows the impact of probability p on the number of found paths. For 30 nodes (resp. 50 nodes) and h_{\max} set to 3, the number of runs where a shortest feasible path was found is 1% (resp. 3%) if $p = 0.05$. While it is 80% (resp. 73%) if $p = 0.30$. Note that the probability of existence of a feasible path is much smaller if $p = 0.05$ than if $p = 0.30$. Figure 4b shows the impact of the maximum stack height h_{\max} on the number of found paths. The probability p is set to 0.20 and there are 2 protocols. Obviously, the algorithm finds more feasible paths if it is allowed to explore paths with higher stacks. However, the difference between stack heights 3 and 5, for example, is negligible. It seems that almost all the shortest feasible paths are of maximum stack height at most 3 (i.e., two nested tunnels).

These simulations shows that the problem is complex, and the convergence time can be prohibitive, especially if the number of adaptation functions per node is large. However, a convergence time of few minutes is sustainable in a network with infrequent topology changes.

VII. CONCLUSION

Nowadays, the Internet encompasses several protocols. The interoperability between these protocols is an important issue, and is ensured thanks to adaptation functions. A path between two nodes may involve different protocols at different stages. The current routing protocols are not able to automatically compute such paths due to the lack of distributed algorithms taking into account the adaptation functions. In this paper, we design the first fully distributed algorithm taking into

⁷<https://networkx.github.io/>

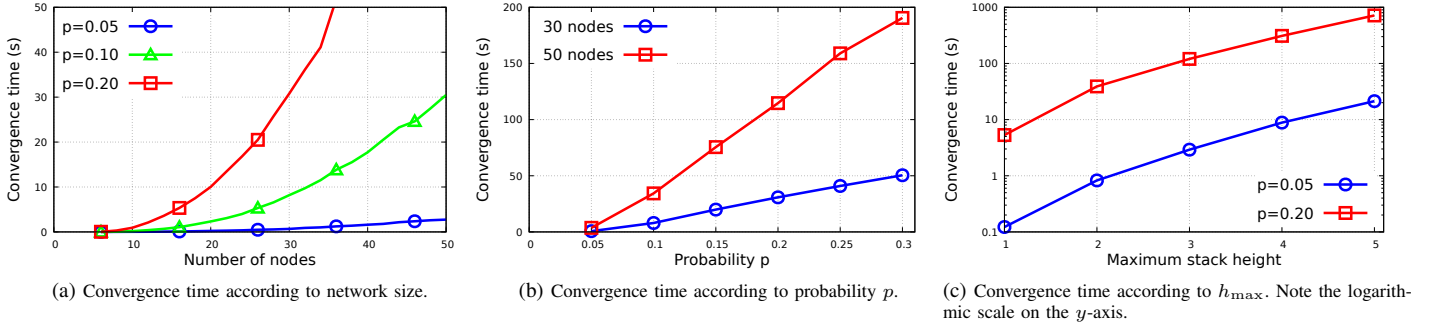


Fig. 3: The convergence time of Algorithm 3 according to different parameters.

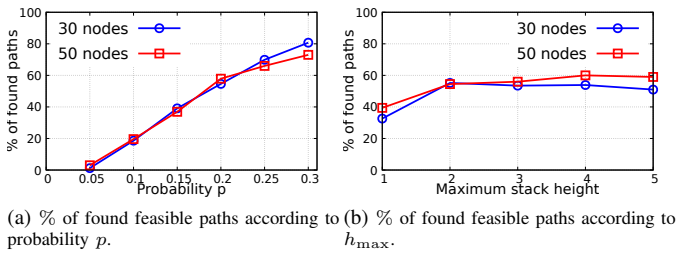


Fig. 4: The % of found feasible paths by Algorithm 3 according to different parameters.

account these functions. Our algorithm builds at each node a local routing table that allows to route a packet following the optimal path. Moreover, we prove that our algorithm converges polynomially in the size of the network and its diameter, and it uses messages of polynomial size, despite the fact that the shortest feasible paths can be of superpolynomial length. We propose a basic implementation of our algorithm as a stack-vector routing protocol, and we evaluate its efficiency through simulations. We believe that this work can have an important impact on protocols such as TSP, on IPv4/IPv6 interoperability, and on automatic tunneling more generally. As a future work, we plan to study the application of our algorithm to secured communication, more precisely to nested encrypted tunnels. In this paper, we assumed that each node had access to the full protocol stack of a packet that it receives. This implies that the data can be encrypted several times but that the headers must not be encrypted. This is not suitable for security reasons. Thus, the main possible improvement to our algorithm would be to adapt it to the case in which all the encapsulated headers are also encrypted. This means that any node has only access to the top (outer) protocol. This would have applications to distributed *onion routing*.

Acknowledgement. The authors would like to thank G raud S nizergues for his invaluable help and explanations about the proof of Lemma 1. The first author was partially supported by the H RA project, funded by The French National Research Agency. Grant no.: ANR-18-CE25-0002.

REFERENCES

- [1] G. Tsirtsis and P. Srisuresh, "RFC 2766 - Network Address Translation - Protocol Translation (NAT-PT)," 2000.
- [2] M. Blanchet and F. Parent, "RFC 5572: IPv6 Tunnel Broker with the Tunnel Setup Protocol (TSP)," 2010.
- [3] M. L. Lamali, N. Fergani, and J. Cohen, "Algorithmic and complexity aspects of path computation in multi-layer networks," *IEEE/ACM Transactions on Networking*, vol. 26, no. 6, pp. 2787–2800, Dec 2018.
- [4] S. Bryant and P. Pate, "RFC 3985: Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture," 2005.
- [5] B. Jiang, P. Nain, D. Towsley, and S. Guha, "On a class of stochastic multilayer networks," in *Sigmetrics*, 2018.
- [6] F. A. Kuipers and F. Dijkstra, "Path selection in multi-layer networks," *Computer Communications*, vol. 32, no. 1, pp. 78–85, 2009.
- [7] M. L. Lamali, H. Pouyllau, and D. Barth, "Path computation in multi-layer multi-domain networks: A language theoretic approach," *Computer Communications*, vol. 36, no. 5, pp. 589–599, 2013.
- [8] F. Iqbal, J. van der Ham, and F. Kuipers, "Technology-aware multi-domain multi-layer routing," *Computer Communications*, vol. 62, pp. 85–96, 2015.
- [9] E. Cshuhaj-Varj , C. Martin-Vide, V. Mitrana, and G. Vaszil, "Parallel communicating pushdown automata systems," *International Journal of Foundations of Computer Science*, vol. 11, no. 04, pp. 631–650, 2000.
- [10] D. Peleg, "Distributed computing: A locality-sensitive approach," *SIAM Monographs on discrete mathematics and applications*, vol. 5, 2000.
- [11] A. Abboud, K. Censor-Hillel, and S. Khoury, "Near-linear lower bounds for distributed distance computations, even in sparse networks," in *International Symposium on Distributed Computing, DISC*. Springer, 2016, pp. 29–42.
- [12] M. Elkin, "Distributed exact shortest paths in sublinear time," in *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC*, 2017, pp. 757–770.
- [13] W. Bang-Li, "A Study on IPsec Tunnel Endpoint Discovery Technology," *Information Security and Communications Privacy*, vol. 4, p. 047, 2010.
- [14] B. Carpenter and C. Jung, "RFC 2529: Transmission of IPv6 over IPv4 Domains without Explicit Tunnels," 1999.
- [15] B. Carpenter and K. Moore, "RFC 3056: Connection of IPv6 Domains via IPv4 Clouds," 2001.
- [16] C. Huitema, "RFC 3068: An Anycast Prefix for 6to4 Relay Routers," 2001.
- [17] F. Templin, "RFC 5579: Transmission of IPv4 Packets over Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) Interfaces," 2010.
- [18] S. Steffann, I. van Beijnum, and R. van Rein, "RFC 7059: A Comparison of IPv6-over-IPv4 Tunnel Mechanisms," 2013.
- [19] A.-L. Barab si and R. Albert, "Emergence of scaling in random networks," *science*, vol. 286, no. 5439, pp. 509–512, 1999.
- [20] G. S nizergues, "A characterisation of deterministic context-free lan-

guages by means of right-congruences.” *Theor. Comput. Sci.*, vol. 70, no. 2, pp. 213–232, 1990.

- [21] A. Amarilli and M. Jeanmougin, “A proof of the pumping lemma for context-free languages through pushdown automata,” *arXiv preprint arXiv:1207.2819*, 2012.

APPENDIX A

Proof of Lemma 1

Proof: We will prove that if there is a feasible path from S to D that reaches a maximum stack height $h_{\max} > \lambda n^2$, then there is a shorter feasible path from S to D that reaches a maximum stack height $h'_{\max} < h_{\max}$. The main ideas of the proof are results from language theory, and are related to the pumping lemma. They can be found in [20] and [21] for example.

Let \mathcal{P} be a feasible path of length ℓ between two nodes. Suppose that \mathcal{P} reaches a maximum stack height $h_{\max} > \lambda n^2$ at some position j (i.e., $h_j = h_{\max}$). For each stack height h such that $h_{\max} - \lambda n^2 \leq h \leq h_{\max}$, let i_h (resp. k_h) be the last position before (resp. the first position after) j reaching the stack height h . More formally:

- $i_h = \max\{i \leq j \mid h_i = h\}$
- $k_h = \min\{k \geq j \mid h_k = h\}$

And let $\langle h \rangle$ be the 3-tuple (U_i, U_k, a_i) where:

- U_i (resp. U_k) is the current node at position i_h (resp. k_h),
- a_i is the current protocol at position i_h (i.e., $a_i = \text{Top}(H_{i_h})$). Note that $\text{Top}(H_{i_h}) = \text{Top}(H_{k_h})$ since f_i is followed by an encapsulation, and f_k follows a decapsulation.

For any h , there are λn^2 possible values of $\langle h \rangle$. However, since $h_{\max} - \lambda n^2 \leq h \leq h_{\max}$, h can take $\lambda n^2 + 1$ different values. Thus, there are two stack heights $h \neq h'$ such that $\langle h \rangle = \langle h' \rangle$. Let $h' < h$ for convenience. By definition of i_h and k_h , the subpath from position i_h to k_h does not involve a protocol stack smaller than h . Thus, it depends only on the part of the stack above h , on the top of the stack, and on the current node at position i_h (recall that i_h is the *last* position before j reaching stack height h , the following positions until k_h involve stacks higher than h). Likewise, the subpath between positions $i_{h'}$ and $k_{h'}$ does not involve a stack smaller than h' . The part of the stack below height h (resp. h') is transparent to the nodes between i_h and k_h (resp. $i_{h'}$ and $k_{h'}$). The same sequence of adaptation functions and nodes leading from i_h to k_h can also lead from $i_{h'}$ to $k_{h'}$. It is then possible to shortcut the part of the path between $i_{h'}$ and i_h together with the part from k_h to $k_{h'}$. Thus, the path:

$$\mathcal{P}' = S f_0 U_1 f_1 \dots U_{i_{h'}-1} f_{i_{h'}-1} U_{i_h} f_{i_h} \dots U_j f_j \dots U_{k_h} f_{k_h} U_{k_{h'}+1} f_{k_{h'}+1} \dots D$$

is feasible, shorter than \mathcal{P} , and reaches a maximum stack height $h_{\max} - (h - h')$. ■

APPENDIX B

Proof of Proposition 1

First we prove the following lemma:

Lemma 2: Let \mathcal{N} be a network. If after a number of rounds t , all the shortest feasible paths reaching a maximum stack height at most h_{\max} are already computed, then all the shortest feasible paths reaching a maximum stack height at most $h_{\max} + 1$ will be computed after $t + O(\text{diam } \mathcal{N})$ rounds.

Proof: Computing a shortest feasible path of maximum stack height $h_{\max} + 1$ from shortest feasible paths of maximum stack height h_{\max} can be done in two steps:

- 1) The new path is a path of maximum stack height h_{\max} following an encapsulation and followed by a decapsulation;
- 2) Concatenation of several new paths of maximum stack height at most $h_{\max} + 1$.

Suppose that after some rounds, all the shortest feasible paths that reach a maximum stack height $h \leq h_{\max}$ are computed. Let $\mathcal{P} = S f_0 U_1 f_1 U_2 f_2 \dots U_\ell f_\ell D$ be a shortest feasible path from node S to D that reaches a maximum stack height h_j at some position j , such that $h_j > h_{\max}$. Let i_h (resp. k_h) be the last position before (resp. the first position after) j that reaches stack height $h = h_j - h_{\max}$, and follows an encapsulation (resp. is followed by a decapsulation). More formally:

- $i_h = \max\{i < j \mid h_i = h_j - h_{\max} \text{ and } h_{i-1} = h_i - 1\}$
- $i_k = \min\{k > j \mid h_k = h_j - h_{\max} \text{ and } h_{k+1} = h_k - 1\}$

Then the path $U_{i_h} f_{i_h} \dots U_j f_j \dots f_{k_h-1} U_{k_h}$ is feasible and is already computed, i.e., node U_{i_h} knows that it can reach U_{k_h} by emitting a packet of some protocol x , and $(U_{k_h}, x, c, U_{i_h+1}, f)$ is in its routing table for some cost c and some protocol x .

On the other hand, node U_{k_h} already knows that it can reach node U_{k_h+1} with some protocol x , since it received the message $(U_{k_h+1}, x, 0)$ from U_{k_h+1} at the first round, thanks to Algorithm 1. Thus, at round t , node U_{i_h+1} knows that it can reach node U_{k_h+1} with some protocol stack. After two rounds, through node U_{i_h} , node U_{i_h-1} will know that it can reach node U_{k_h+1} with some protocol x .

Now it remains to prove that if $\mathcal{P}_1, \dots, \mathcal{P}_d$ are shortest feasible paths already computed, and the last node of \mathcal{P}_i (call it D_i) is the first node of \mathcal{P}_{i+1} , then the shortest feasible path $\mathcal{P}' = \mathcal{P}_1 \dots \mathcal{P}_d$ (if any) is computed after $O(\text{diam } \mathcal{N})$ rounds. It means that the first node of \mathcal{P}_1 knows that it can reach the last node of \mathcal{P}_d by emitting a suitable protocol. By the same argument, the first node of \mathcal{P}_1 will know that it can reach the last node of \mathcal{P}_d after at most $|\mathcal{P}_1, \dots, \mathcal{P}_d|$ rounds. Since, by definition, $|\mathcal{P}_1, \dots, \mathcal{P}_d| \leq \text{diam } \mathcal{N}$, this concludes the proof. ■

Recall now Proposition 1:

Proposition 1: Algorithm 3 computes the correct routing tables after $O(\lambda n^2 \text{ diam } \mathcal{N})$ rounds.

Proof: We prove by induction that after $O(\lambda n^2 \text{ diam } \mathcal{N})$ rounds, all the shortest feasible paths that reach a maximum

stack height $h \leq h_{\max}$ are computed, i.e., the routing tables are able to route packets following these paths.

Basis: Suppose that there is a feasible path $\mathcal{P} = Sf_0U_1f_1U_2f_2\dots U_\ell f_\ell D$ from $S = U_0$ to $D = U_{\ell+1}$ that keeps stack height of 1 (i.e., there is neither encapsulations nor decapsulations along the path). Since, by definition, $|\mathcal{P}| \leq \text{diam } \mathcal{N}$, after at most $\text{diam } \mathcal{N}$ rounds, the routing table of S contains the row $(D, x, c, U_1, w(\mathcal{P}))$ for $0 \leq i \leq \ell$ and some protocol $x \in \text{In}(S)$, since \mathcal{P} is feasible.

Induction: Suppose that after t rounds, all the shortest feasible paths that reach a maximum stack height at most h_{\max} are computed. By Lemma 2, all the shortest feasible paths that reach a maximum stack height of $h_{\max} + 1$ are computed after $t + O(\text{diam } \mathcal{N})$ rounds. By Lemma 1, the maximum stack height of a feasible shortest path (if any) is λn^2 . This concludes the proof. ■