



	<i>centralized/SDN</i>	<i>distributed/traditional</i>	<i>hybrid</i>
	OpenFlow [2], PCE [3], SR [4]	IGP [5, 6], RSVP-TE [1]	<b>Fibbing</b>
<i>forwarding paths:</i>			
- <i>configuration</i>	simple (declarative & global)	complex (indirect & per-device)	<b>simple</b> (declarative & global)
- <i>manageability</i>	high (direct control)	low [7, 8] (need for coordination)	<b>high</b> (direct control)
- <i>path installation</i>	slow (by controller, per-device)	fast (by device, distributed)	<b>fast</b> (by device, distributed)
<i>robustness:</i>			
- <i>network failures</i>	slow (by controller)	fast (local)	<b>fast</b> (local)
- <i>controller failures</i>	hard (ad-hoc synch)	native (distributed)	<b>easy</b> (synch via IGP)
- <i>partitions</i>	hard (uncontrollable devices)	best (distributed)	<b>best</b> (fallback on distributed)
<i>routing policies:</i>	highest (any path)	- low for IGP (shortest paths) - highest for RSVP (any path)	<b>high</b> (any non-loopy paths)

**Table 1: Fibbing combines the advantages of existing control planes, avoiding the main drawbacks.**

deed, IGPs like OSPF and IS-IS are scalable (support networks with hundreds of nodes), robust, and quickly react to failures. Building a SDN controller with comparable scalability and reliability is challenging. It must compute and install forwarding rules for all the switches, and respond quickly to topology changes. Even the simple task of updating the switch rule tables can then become a major bottleneck for a central controller managing hundreds of thousands of rules in hundreds of switches. In contrast, distributed routing protocols naturally parallelize this work. For reliability and scalability, a SDN controller should also be replicated and geographically distributed, leading to additional challenges in managing controller state. Finally, the deployment of SDN as a whole is a major hurdle as many networks have a huge installed base of devices, management tools, and human operators that are not familiar with the technology. As a result, existing SDN deployments are limited in scope, *e.g.*, new deployments of private backbones [8, 9] and software deployments at the network edge [10].

This paper introduces *Fibbing*, a technique that offers *direct control* over the routers’ forwarding information base (FIB) by manipulating the input of a distributed routing protocol. Fibbing relies on traditional link-state protocols such as OSPF [5] and IS-IS [6], where routers compute shortest paths over a synchronized view of the topology. Fibbing controls routers by carefully *lying* to them, removing the need to configure them. It coaxes the routers into computing the target forwarding entries by presenting them with a carefully constructed augmented topology that includes fake nodes (providing fake announcements of destination address blocks) and fake links (with fake weights). In essence, Fibbing inverts the routing function: given the forwarding entries (*i.e.*, the desired output) and the routing protocol (*i.e.*, the function), Fibbing automatically computes the routing messages to send to the routers (*i.e.*, the input).

Fibbing can solve the problem in Fig. 1a adding two fake nodes (Fig. 1b), connected to *A* and *E* with the depicted weights. Both fake nodes advertise that they can reach  $D_1$  directly. Based on the augmented topology, *D* starts to use *A* to reach  $D_1$ , as the new cost (3) is lower

than the original one (6). *A* and *E* also select different paths. Since the fake nodes do not really exist, packets forwarded by *A* or *E* actually flow through *B*. Routers *B* and *C* do not change their forwarding decisions.

Table 1 gives an overview of how Fibbing improves flexibility and manageability by adopting a SDN-like approach while keeping the advantages of distributed protocols (*e.g.*, robustness and fast FIB modifications).

**Fibbing is expressive.** Fibbing can steer flows along *any* set of per-destination loop-free paths. In other words, it can exert full control at a per-destination granularity. For this reason, Fibbing readily supports advanced forwarding applications such as: (a) traffic engineering, (b) load balancing, (c) fast failover, and (d) traffic steering through middleboxes. By relying on destination-based routing protocols, Fibbing does not support finer-grained routing and forwarding policies such as matching on port numbers. Though, those policies can easily be supported via middleboxes.

**Fibbing scales and is robust to failures.** Lying to routers is powerful but challenging. Indeed, Fibbing must be *fast* in computing augmented topologies to avoid loops and blackholes upon network failures. At the same time, Fibbing must compute *small* augmented topologies since routers have limited resources. Finally, Fibbing must be *reliable* and gracefully handle controller failures. We address all three challenges.

**Fibbing differs from previous approaches that rely on routing protocols to program routers.** Prior approaches like the Routing Control Platform [11] rely on BGP as a “poor man’s” SDN protocol to install a forwarding rule for each destination prefix on each router. In contrast, Fibbing leverages the routing-protocol implementation on the routers. Doing so, Fibbing can adapt the forwarding behavior of many routers at once, while allowing them to compute forwarding-table entries and converge on their own. That is, while the controller computes the routing *input* centrally, the routing *output* is still computed in a distributed fashion.

**Fibbing works on existing routers.** We implemented a fully-functional Fibbing prototype and used it to pro-




Figure 2: The four-staged Fibbing workflow.

gram real routers (both Cisco and Juniper). Based on an augmented topology, these routers can install hundreds of thousands of forwarding entries with an average installation time of less than 1ms per entry. This offers much greater scale and faster convergence than is possible with state-of-the-art SDN switches [12, 13], without requiring the deployment of new equipment and per-device actions from the controller. This also means that Fibbing can implement recent SDN proposals, like Google’s B4 [8] and Microsoft’s SWAN [9]—on top of existing networks.

Our earlier work showed that Fibbing can enforce any set of forwarding DAGs [14]. This paper goes further by describing the complete design, implementation, and evaluation of a Fibbing controller managing intra-domain routing. Rather than focusing on specific use cases (like traffic engineering), we describe its support for different higher-level approaches (*e.g.*, [8, 9]). We make the following contributions:

- **Abstraction:** We show how to express and realize high-level forwarding requirements by manipulating a distributed link-state routing protocol (§2).
- **Algorithms:** We propose new, efficient algorithms to compute compact augmented topologies (§3).
- **Implementation:** We describe a complete Fibbing implementation which works with unmodified Cisco and Juniper routers (§4).
- **Evaluation:** We show that our Fibbing controller quickly generates small augmented topologies, inducing minimal load on routers (§5).

## 2. FLEXIBLE FIBBING

Fibbing workflow proceeds in four consecutive stages based on two inputs: the desired forwarding graphs (one Directed Acyclic Graph, or DAG, per destination) and the IGP topology (Fig. 2). The forwarding DAGs can either be provided directly or expressed indirectly, at a high-level, using a simple path-based language. In the latter case, the **Compilation** (§2) stage starts by compiling the requirements into concrete forwarding DAGs. Then, the **Topology Augmentation** (§3) stage computes an augmented topology that achieves these forwarding DAGs. While computing an augmented topology is fast, the resulting topology can be large. As such, the job of the next **Topology Optimization** (§3) stage is to reduce the augmented topology while preserving the forwarding paths. Finally, the **Injection & Mon-**

$pol$	::=	$(s_1; \dots; s_n)$	Fibbing Policy
$s$	::=	$p \mid b$	Requirement
$r$	::=	$p_1 \text{ and } p_2 \mid p_1 \text{ or } p_2 \mid p$	Path Req.
$p$	::=	$\text{Path}(n^+)$	Path Expr.
$n$	::=	$id \mid * \mid n_1 \text{ and } n_2 \mid n_1 \text{ or } n_2$	Node Expr.
$b$	::=	$r \text{ as backupof } ((id_1, id_2)^+)$	Backup Req.

Figure 3: Syntax of Fibbing high-level language.

**itoring** (§4) stage turns fake information into actual “lies” that the controller injects into the network.

In this section, we present the high-level language and compilation process (§2.1), and show that Fibbing can express a wide range of forwarding behaviors (§2.2).

### 2.1 Fibbing high-level language

Fibbing language (Fig. 3) provides operators with a succinct and easy way to specify their forwarding requirements. A Fibbing policy is a collection of *requirements*, naturally expressed as regular expressions on paths. Each requirement is either a *path requirement* which specifies how traffic should flow to a given destination, or a *backup requirement* which specifies how traffic should flow if the IGP topology changes. Each path requirement is recursively defined as a composition of path requirements through logical AND and OR. Operators can express load-balancing requirements using a conjunction of  $n$  requirements. Similarly, they can ensure that traffic to a specific destination will take one of  $n$  paths (*e.g.*, containing a firewall) using disjunction. Path requirements are composed of a sequence of node requirements. A node requirement is either a node (router) identifier or the wildcard  $*$ , representing any sequence of nodes. Like path requirements, node requirements can be combined using logical AND and OR. Whenever no path requirement is stated, the original IGP paths should be used. This way, operators do not have to express all the unmodified paths, only deviations from the IGP shortest paths.

The following example illustrates the main features of the language. It states that traffic between  $E$  and  $D_1$  should be load-balanced on two paths, traffic between  $A$  and  $D_2$  should cross  $B$  or  $C$  and that traffic from  $F$  to  $D_3$  should be rerouted via  $G$  if the link  $(F, H)$  fails.

```
( (E, C, D1) and (E, G, D1);
  ((A, *, B, *, D2) or (A, *, C, *, D2));
  (F, G, *, D3) as backupof ((F, H)); )
```

Fibbing policies are compiled into per-destination forwarding DAGs by finding convenient network paths (if any). Compilation works in two consecutive stages. First, the compiler expands any requirement with wildcards into paths. This step can be computationally expensive as, in general, a network can have a number of paths exponential in the number of nodes. While this is unlikely, especially for networks designed according to best practices, we bounded the number of paths that can be expanded out of a single requirement. We only expand again if no solution is found with the current set of paths. Once all requirements are expanded, the


compiler groups them by destination and computes the Disjunctive Normal Form (DNF) of each requirement. To finally produce a forwarding DAG, the compiler iterates over the disjunction of path requirements and checks whether the resulting graph is loop-free.

## 2.2 Fibbing expressiveness

Beyond steering traffic along a given path (§1), we now show that Fibbing can also: (i) balance load over multiple paths and; (ii) provision backup paths.

### Fibbing can forward any flow on any set of paths.

Fibbing can load-balance traffic over multiple paths to maximize throughput, minimize response time, or increase reliability. For example, consider the network in Fig. 4a where three sources  $S_1$ ,  $S_2$ , and  $S_3$  send traffic to three corresponding destinations. Demands and link capacities are such that link  $(F, G)$  is congested. One way to alleviate congestion is to split traffic destined to  $D_2$  over the top (via  $(B, C)$ ) and bottom (via  $(F, G)$ ) paths. Load-balancing traffic coming from  $E$  on multiple paths is possible under conventional link-state routing (e.g., by re-weighting links  $(F, G)$  to 15). However, this would force the traffic from  $S_2$  and  $S_3$  to spread over both paths, creating congestion. More generally, it is impossible to route the traffic destined to  $D_2$  and  $D_3$  on different links under conventional link-state routing.



**Figure 4: Fibbing supports multi-path forwarding. Here, it avoids the initial congestion (see (a)) by load-balancing traffic for  $D_2$  (see (b)).**

This simple requirement can easily be expressed as:


$$((S_2, E, B, C, H, D_2) \text{ and } (S_2, E, F, G, H, D_2))$$

Fig. 4b shows the augmented topology which realizes this requirement. A fake node announcing  $D_1$  (with a weight of 6) is inserted between  $E$  and  $B$ . After introducing this node,  $E$  has two shortest paths (of cost 7) to reach  $D_2$  and, hence, splits  $D_2$  traffic over  $B$  and  $F$ . In this example, Fibbing enables maximum network efficiency as each link is used to its full capacity.

### Fibbing can provision backup paths for any flow.

Fibbing can provision backup paths to prevent congestion or increased delays after link and node failures. As an illustration, consider the network in Figure 5a. The failure of link  $(E, F)$  leads to congestion since traffic flows for  $D_1$  and  $D_2$  are both rerouted to the same path via link  $(A, B)$ . To prevent congestion, traffic destined to  $D_1$  and  $D_2$  should be split over the two remaining disjoint paths but only upon failure on the path

$(A, D, E, F)$ . This is another example of a requirement that is impossible to achieve with link-state routing, and would require significant control-plane overhead in MPLS. In contrast, it is easily done with Fibbing.



**Figure 5: Fibbing can provision backup paths. Here, possible congestion upon a link failure (see (a)) is avoided by adding a fake node (see (b)).**

Backup paths can be specified in our language as:

$$(A, B, *, D_1) \text{ and } (A, G, *, D_2) \\ \text{as backupof}((A, D) \text{ or } (D, E) \text{ or } (E, F))$$

Fig. 5b shows the corresponding augmented topology, which has a single fake node advertising  $D_2$ . The weights are set to prevent  $A$  from using the fake node to reach  $D_2$  unless a failure occurs along the path  $(A, D, E, F)$ . While successful for this example, Fibbing cannot satisfy all possible requirements for backup paths (§3.3).

## 3. AUGMENTING TOPOLOGY

In this section, we detail the augmentation problem (§3.1), and we show how the Fibbing controller quickly computes small augmented topologies from a set of forwarding DAGs. We rely on a *divide-and-conquer* approach based on three consecutive steps.

**1) Topology initialization (§3.2):** We modify the initial weights in the link-state protocol (if necessary), to guarantee that any set of forwarding DAGs can be enforced by Fibbing. If needed, this operation has to be done only once, when Fibbing is first deployed.

**2) Per-destination augmentation (§3.3):** Starting from an initialized topology, we compute a suitable augmentation, individually for every destination of an input forwarding DAG. We designed two algorithms for this step, achieving different trade-offs between computation time and augmentation size. The fastest one, **Simple**, can compute augmented topologies *within milliseconds*, and works by injecting a dedicated fake node for every router that changes its next-hop. The relatively slower one, **Merger**, reduces the augmentation size by re-using the same fake nodes to program multiple routers. Simple and Merger are suited for different goals. The speed of Simple is useful for quick failure reaction. In contrast, Merger can be run in background to progressively re-optimize the augmented topology. We evaluate the trade-offs achieved by each algorithm in §5.

**3) Optimization across destinations (§3.4):** We merge the augmentations obtained in the per-destination augmentation step to further reduce the number of fake

nodes and edges. Namely, whenever safe, we replace multiple fake nodes announcing different destinations with a single fake node which either announces all the destinations or creates a new path (a shortcut) between routers in the augmented topology.

### 3.1 The Topology Augmentation Problem

We start the description of the topology augmentation algorithms by precisely defining the basic concepts on which they rely and the problem that they solve.

**Fake nodes scoping.** A Fibbing controller can generate both *locally-scoped* lies (targeted to a single router) and *globally-scoped* lies (targeted to all routers). Locally-scoped lies are useful as they enable local actions on one router without creating side effects on other routers. Globally-scoped lies affect the entire network. Hence, if carefully computed, they can reduce the size of the augmented topology. All of our previous examples used globally-scoped lies. We detail how to implement both kinds of lies in the current OSPF in §4.

**Fake edges to forwarding next-hop mapping function.** Fibbing can modify the routing path computed by any IGP router  $r$ . In particular, it can augment the IGP topology so that  $r$ 's shortest path is no longer the one in the original topology but includes some fake sub-paths. Throughout the paper, we assume that a fake edge in the shortest path from any router  $r$  to any destination  $d$  corresponds to the ability to force the next-hop of  $r$  for  $d$  to be any of its neighbors. In the example in Fig. 1, for instance, the fake edge between the  $A$  and its adjacent fake node translates into  $A$  forwarding traffic to  $B$ . We discuss in §4 how to achieve this ability in the current OSPF protocol, as well as in future IGPs.

**Topology augmentation problem.** Since we assume an arbitrary mapping between fake edges and forwarding next-hops, the topology augmentation problem is defined as follows: Given an initial topology  $G$  and a set of forwarding DAGs, compute an augmented topology  $G' \supset G$  such that for each path  $(u, v, \dots, d)$  in the forwarding DAG for  $d$ , the next-hop of  $u$  in one of its shortest paths for  $d$  in  $G'$  is either  $v$  or a fake node.

### 3.2 Topology Initialization

In the topology initialization, we scale the link weights of the original IGP topology  $G$  to guarantee arbitrary per-destination control through Fibbing and help reduce the size of topology augmentations. In particular, we proportionally increase link weights (multiplying them by a constant factor) if they are too low in  $G$ . Moreover, we set very high announcement cost for any destination, at least equal to the length of the longest path in  $G$  times the maximum link weight.

**Topology initialization enables full Fibbing expressivity.** Indeed, it makes the IGP topology *Fibbing compliant*, which provably avoids cases in which a forwarding DAG cannot be implemented by Fibbing

(see [15]). We say that a topology is Fibbing compliant if for every destination  $d$ , the cost of the shortest path from every router (including the ones announcing  $d$ ) to  $d$  exceeds 2. In Fibbing compliant topologies, for any router  $r$  and destination  $d$  the controller can always compute a fake path  $P$  such that (i)  $P$  is shorter than the original shortest path from  $r$  to  $d$ ; and (ii)  $P$  is longer than the original shortest path from any other router  $v \neq r$  to  $d$ . As proved by the following theorem, this implies the ability of Fibbing to forward flows for the same destination on any set of loop-free paths.

**THEOREM 1.** *Any set of per-destination forwarding DAGs can always be enforced by augmenting a Fibbing-compliant topology even only with globally-scoped lies.*

**PROOF.** We prove the statement by showing a simple topology augmentation procedure. Let  $G$  be the initial topology. For every forwarding DAG with destination  $d$ , we add for each node  $r$  in the network a fake node  $f_r$  announcing  $d$ . This generates a new fake path  $(r, f_r, d)$  in the augmented topology. We set the total cost of this newly added fake path to 2. Since  $G$  is Fibbing compliant, then the cost of the shortest path from  $r$  to  $d$  in  $G$  is greater than 2. Hence, the shortest path of every node  $r$  in the augmented topology will be  $(r, f_r, d)$ . The forwarding DAG is then implemented by mapping the fake link on the right physical link.  $\square$

Note that Theorem 1 applies to destinations in the augmented topology. Those destinations do not need to match the destination prefixes announced in the original IGP. Hence, Fibbing allows to control flows for de-aggregation of the original IGP destinations (up to IP address granularity) or even non-overlapping prefixes.


**Topology initialization is non-intrusive.** Since it is based on the adaptation of a few configurable parameters, our initialization procedure can be applied to any link-state routing configuration, and *preserves the original forwarding paths*. It can be carried out in a running network, using known lossless reconfiguration techniques [16]. Moreover, it is strictly needed no more than once in the network lifetime. Indeed, since Fibbing compliance does not depend on the routing requirements or the presence of specific links, any topology remains Fibbing compliant independently of new requirements or the failures of nodes or links. Finally, topologies growing in size can be easily kept Fibbing compliant by ensuring that the new destinations are announced with high costs and the new links have weights consistent with pre-existing ones.

### 3.3 Per-destination augmentation

We now describe Simple and Merger. We use Figure 6 to illustrate the difference of the two algorithms.

#### 3.3.1 Simple

Simple relies solely on locally-scoped lies to avoid having to compute any fake path cost. For every destina-



**Figure 6: Outcome of our per-destination augmentation algorithms. Simple produces larger topologies (with locally-scoped lies) in a very short time, while Merger reduces the size of the augmentation by relying on globally-scoped lies and a longer computation.**

tion  $d$  and corresponding forwarding DAG  $D$ , the algorithm adds fake nodes to each router whose next-hops in the original topology differ from those in  $D$ . Precisely, for every router  $r$  that changes its next-hop for  $d$ , Simple adds a fake node  $f_{r,d}$  and a fake link  $(r, f_{r,d})$ . Node  $f_{r,d}$  announces  $d$  to  $r$  with a locally-scoped lie. We set the total cost of path  $(r, f_{r,d}, d)$  to 2. Since the topology is Fibbing compliant,  $r$  is ensured to change its shortest path. Also, since the lie is locally-scoped, other routers are not affected by it.

Figure 6b shows the output of Simple for the example of Figure 6a. Nodes  $A, B, C$  and  $X$  are required to change their respective next-hops. Moreover,  $A$  needs to load-balance on  $B$  and  $C$ . Thus, Simple creates five locally-scoped nodes (two connected to  $A$ ), all providing fake paths to the destination with a cost of 2.

### 3.3.2 Merger

To reduce the number of fake nodes, Merger relies on globally-scoped lies that can change the forwarding behavior of multiple routers at once. When applied to Figure 6a, Merger creates only two fake nodes (see Fig. 6c) to change the next-hops of  $A, B, C$  and  $X$ , instead of the five used by Simple. The added fake nodes create load-balancing on  $A$  (cost 136) via  $B$  and  $C$ , as required.

Merger performs the topology augmentation for any destination  $d$  in two phases. First, it adds an excessive number of fake nodes, and computes the lower and upper bounds for their respective cost. Second, it merges fake nodes whenever possible, based on the value of the computed bounds. We now provide an intuitive description of those two phases. Additional details about them and Merger correctness proofs are reported in [15].

**Step 1. Fake bounds computation.** Merger starts by adding fake nodes to every router  $r$  that is required to change the next-hop (according to the input forwarding DAG) for  $d$ . In this case, one new fake node  $f_{r,d}$  is connected to  $r$  for every new  $r$ 's next-hop in the input forwarding DAGs. However, to enable merging of globally-visible fake nodes, Merger calculates lower and upper bounds for every newly-added fake path  $(r, f_{r,d}, d)$ . Since the initial positioning of fake nodes is as in Simple, we illustrate bound computations referring to Figure 6b.

The *upper bound*  $ub(f_{r,d})$  represents the maximum value of the fake path cost that changes  $r$ 's shortest path to  $(r, f_{r,d}, d)$ . It is easy to compute by statically considering the non-augmented topology  $G$ . Indeed, it is equal to  $dist(r, d, G) - 1$ , where  $dist(u, v, G)$  is the cost of the shortest path from  $u$  to  $v$  in  $G$ .

The *lower bound*  $lb(f_{r,d})$  represents the minimal value of the fake path cost that does not change the shortest path of any real node different from  $r$ . To compute it, we divide nodes in two sets, depending on whether the input forwarding DAG prescribes to change their respective next-hops or not.

For the next-hop preserving nodes whose shortest path does not traverse  $r$  in the input topology, we impose that their original shortest path is not modified by  $f_{r,d}$ . For example, when computing  $lb(f_0)$  in Figure 6b, we ensure that  $f_0$  does not change the shortest path of  $F$ , by constraining  $lb(f_0) + 25 > dist(F, d, G) = 110$ , that is,  $lb(f_0) = 86$ . More generally, for every fake node  $f_{r,d}$  and every next-hop preserving neighbor  $n$  of  $r$ , we impose that  $lb(f_{r,d}) > dist(n, d, G) - dist(n, r, G)$ .

For next-hop changing nodes (connected to other fake nodes), the final value of their shortest paths is not known in advance, but is determined by the augmentation itself. That is, the lower bound of any fake node generally depends on the lower bound of other fake nodes. For example, in Figure 6b,  $f_4$  changes the shortest path of  $X$  only if  $lb(f_4) < dist(X, C, G) + lb(f_3)$ . To avoid that real nodes pass through fake nodes not directly connected to them, Merger runs a *lower bound propagation procedure*. This procedure takes as input the lower bounds initialized with values from next-hop preserving nodes. It then fixes one lower bound at the time, following a specific order and adjusting the others to be consistent with the fixed one. This order guarantees that each lower bound must be considered only once. Sometimes, lower bounds cannot be made consistent. Indeed, Theorem 1 does not provide guarantees if fake nodes are connected only to next-hop changing nodes. We solve these cases by using locally-visible lie.

**Step 2. Fake nodes merging.** In this step, Merger tries to merge fake nodes together. More precisely, it iterates over every simple path from a source to  $d$  in

the input forwarding DAG. For each of those paths, it merges pairs of fake nodes whenever safe.

To assess when it is safe to merge a fake nodes  $f'$  into  $f''$ , Merger sequentially performs three checks. We illustrate these checks by considering the required path  $(A, B, X, D)$  and the merge of  $f_1$  into  $f_2$  in Figure 6b.

First, Merger assesses whether the IGP shortest paths are compliant with the considered source-sink path in the DAG. In our example, it verifies that the shortest path from  $A$  and  $B$  (*i.e.*,  $(A, B)$ ) is a sub-path of the required  $(A, B, X, D)$ . If  $A$  had predecessors in the required path not connected to a fake node, this check would have been repeated for those predecessors as well.


Second, Merger checks the possibility to use  $f''$  as part of the shortest path of the real node connected to  $f'$  without changing the current next-hops of any node. To this end, the algorithm assess the existence of feasible post-merging bounds for  $f''$ . More precisely, it re-computes the modified lower bound of  $f''$  as the minimum value (if any) that forces the new shortest path of the real node connected to  $f'$  and its next-hop preserving predecessors in the required path through  $f''$ , without affecting nodes previously not crossing  $f''$ . In our example, the lower bound of  $f_2$  is modified to exclude the constraint of not changing  $A$ 's next-hop, hence it is decreased to  $lb(f_2) = 81$  (it was greater before, for  $A$ 's next-hop to be  $f_1$ ). The upper bound of  $f''$  is also modified to ensure that the real node connected to  $f'$  and its next-hop preserving predecessors use the fake path via  $f''$ . In our example,  $ub(f_2)$  is modified to 129, as the cost of the original shortest path from  $A$  to the destination is 135 and those between  $A$  and  $B$  is 5.

Third, Merger simulates the merge to assess whether all bounds can be consistently adjusted network-wide, given that the merging  $f'$  into  $f''$  would change the bounds of  $f''$  and remove one fake node. To this end, we re-run the lower bound propagation procedure, devoting special attention to fake nodes used for load-balancing. In our example, for instance, we constrain the lower bound of  $f_0$  to be equal to the cost of path  $(A, B, f_2, d)$ , meant to be used by  $A$  after the merging.

If all the three checks pass, then Merger actually perform the merge, by removing  $f'$  and updating the bounds of all other fake nodes (including  $f''$ ) according to the values computed during the last check.

### 3.3.3 Dealing with backup requirements

Backup requirements can be specified by providing additional sets of (tagged) forwarding DAGs to our algorithms. Let  $G'$  be an augmented topology computed to accommodate primary requirements. To deal with backup ones, slight variants of Simple and Merger can be run after the computation of  $G'$ . The main modification of Simple consists in setting the cost of the fake paths added for backup requirements to 3 instead of 2. In contrast, backup requirements are supported in Merger by imposing that lower bounds are always greater than the cost of the shortest path in  $G'$ .



(a) Fake destination merging. (b) Fake shortcuts creation.

**Figure 7: Cross-destination optimizations.**

Contrary to primary requirements (see Theorem 1), Fibbing may not enforce backup requirements, even in a Fibbing-compliant topology (see [15]). Indeed, if the cost of the original shortest path of a node  $r$  is equal to a fake one (used for a primary requirement) in  $G'$ , then backup paths different from the original shortest path cannot be implemented on  $r$ . In those cases, we notify the operator on the impossibility to implement the given backup paths.

## 3.4 Cross-Destination Optimization

Fake nodes computed on a per-destination basis may be redundant. We reduce such redundancy in two ways, namely, by (i) merging fake nodes connected to the same real node, and (ii) replacing fake destination announcements with fake paths connecting real nodes.

**Cross-destination merging.** After per-destination augmentations, two fake nodes  $f_1$  and  $f_2$  announcing different destinations  $d_1$  and  $d_2$  can be connected to the same node  $r$  and used to force traffic to the same real link  $(r, n)$ . Those fake nodes can always be merged. Indeed, we can replace  $f_1$  and  $f_2$  with a new fake node  $f'$  such that (i)  $cost(r, f') = \min\{cost(r, f_1), cost(r, f_2)\}$ , and (ii)  $f'$  announces both  $d_1$  and  $d_2$ , with  $cost(f', d_i) = cost(r, f_i, d_i) - cost(x, f')$  for  $i = 1, 2$ . For example, assume that in the network in Figure 6 additional destinations are attached to  $C$  and  $F$  as in Figure 7a. The result of the cross-destination merging is shown in Figure 7a, where both  $A$  and  $X$  have a single fake neighbor announcing multiple destinations (rather than multiple fake neighbors each announcing a single destination). This reduces the number of fake nodes from 4 to 2.

**Creating shortcuts.** One of the most appealing features (unmatched by competitor solutions) of Fibbing is that a single lie can change the paths for multiple destinations. To this end, we need however to replace fake destination announcements with fake paths connecting real nodes together, *i.e.*, *fake shortcuts*.

Currently, we use fake shortcuts only if a real link is never traversed in different directions. Consider, for example, the link between  $X$  and  $D$  in Figure 7a. It is traversed from  $X$  to  $D$  for the destinations attached to  $D$  and  $F$ , and never from  $D$  to  $X$ . In those cases, we try to transform  $X$ 's fake neighbor into a fake shortcut, as in Figure 7b. Let  $u$  and  $v$  be the two real nodes at the endpoints of the shortcut. First, we check if a

shortcut cost  $c$  exists such that all the shortest paths<sup>1</sup> are kept the same with and without the shortcut. If this condition is met, given a fake node  $f$  used to enforce the subpath  $(u, \dots, v)$  in the input forwarding DAGs, all fake destinations announced by  $f$  are replaced by a fake shortcut  $(u, f, v)$  and the cost of  $(u, f, v)$  is set to  $c$ . Figure 7b illustrates that we found such a value  $c = 20$  for  $X$  in our example. Also, we use asymmetric weights to prevent the fake shortcut from being traversed in the opposite direction. To this end, the cost of path  $(v, f, u)$  is set to a very high value, *e.g.*, by setting a high weight of the directed link  $(v, f)$ . In Figure 7b, we indeed set the weight of the link from  $D$  to the fake node in the shortcut to 500.

## 4. IMPLEMENTATION

We built a complete prototype of Fibbing in Python (algorithmic part) and C (interaction with OSPF) by extending Quagga [17]. Fibbing code base spans over 2300 (resp. 400) lines of Python (resp. C) code. It is available at <http://www.fibbing.net>. In this section, we present our prototype (§4.1), describe how Fibbing works with current OSPF routers (§4.2), and propose two small modifications to link-state protocols that would make Fibbing even more efficient (§4.3). Finally, we describe how to ensure controller reliability (§4.4).

### 4.1 Fibbing Controller

Our prototype consists of three main components:

**Fake topology generator** applies (i) the compilation algorithms (§2) to turn forwarding requirements into forwarding DAGs and (ii) the augmentation algorithms (§3) to convert the forwarding DAGs into fake nodes and links. The topology generator uses a JSON interface to register for update events produced by the event manager. Upon network updates, the generator automatically recomputes the augmented topology either using Simple to ensure fast convergence, or Merger to reduce the size of the topology. To ensure fast convergence *and* a small augmented topology, the generator also pre-computes augmentations with Merger, *e.g.*, those needed for any single link failure, and stores them in a *deltas database*.

**Link-state translator** interacts with the routers by establishing routing adjacencies to inject lies and track topology changes. Thanks to the flooding mechanism used by link-state protocols, a single adjacency is sufficient to send and receive all routing messages to/from all routers. Though, maintaining several adjacencies is useful for reliability. In this case, the translator simply injects the lies via all adjacencies. While this slightly increase the flooding load, doing so does *not* impact the

routers memory as each message has a unique identifier and routers only maintain one copy per ID in memory.

**Event manager** maintains an update-to-date view of the network topology by (i) parsing the routing messages collected by the translator and (ii) constructing a network graph for the topology generator. The event manager checks whether each new event (*e.g.*, a node failure or addition) affects any of the forwarding requirements. If so, it first checks the deltas database for a pre-computed lie, and otherwise notifies the topology generator to request a new augmented topology.

### 4.2 Fibbing with Unmodified OSPF

Our Fibbing prototype works with unmodified OSPF-speaking routers (tested on Cisco and Juniper). To create lies, our prototype leverages the *Forwarding Address* (FA) [18] field of OSPF messages. Suppose the controller wants routers to think that destination  $d$  is directly attached to the router with IP address  $y$ . Then, the controller injects the route for  $d$  with a forwarding address of  $y$  and the desired cost for the fake edge from  $y$  to  $d$ . Router  $y$  ignores the message, and all other routers compute the cost of the route as the sum of their cost to  $y$  plus the cost in the message.

**Locally-scoped lies in OSPF.** To support locally-scoped lies, we reserve a set of IP addresses to be used as FAs. All those addresses are propagated network-wide in OSPF. However, every router is configured not to install routes to all those addresses. Consequently, only the directly-connected router can reach any of those addresses, and accept routes specifying that address as FA. Both allocation and configuration of FA-associated IP addresses can be done just once in the network lifetime.

**Globally-scoped lies in OSPF, and limitations.** OSPF readily supports globally-scoped lies by simply propagating OSPF messages with the FA set to an IP address announced in OSPF. However, some subtle constraints hold in an unmodified OSPF network due to how FAs are resolved on the router. Prominently, (i) OSPF fake nodes are actually fake routes to the specified FA, hence both their positioning and the path to be used for reaching them are constrained; and (ii) OSPF routers discard any OSPF message whose FA is one of its own IP address and computes its shortest path according to the topology without the fake node. The combination of these two constraints limit the power of globally-scoped lies in OSPF, making them insufficient to implement all possible forwarding DAG. An example of those cases is reported in [15].

**Overcoming OSPF limitations.** To support the full expressiveness of Fibbing, our prototype controller uses an OSPF-compliant implementation of Merger with a combination of locally and globally-scoped lies. This implementation uses globally-scoped lies whenever possible, and falls back to locally-scoped lies for any requirements that cannot be met that way.

<sup>1</sup>By definition of the shortest path, we only need to check the paths from  $v$  and from  $u$  neighbors to each destination.



### 4.3 Proposed Protocol Enhancements

Other link-state protocols, like IS-IS [6], do not support forwarding addresses, and even the OSPF implementation of lies has limitations. However, minor protocol extensions can enable more flexible Fibbing in future routers. Fully-fledged Fibbing needs for the routing protocol to support two functions: (i) the creation of adjacencies (with fake nodes) on the basis of a received message; and (ii) a third-party next-hop mechanism which allows to specify in a route the forwarding next-hop to be used if that route is selected. Support for these functions can be added to protocol specifications (without impacting current functionalities), and can be deployed through router software updates.

Preliminary discussions with router vendors confirm that these changes are reasonable and could be integrated into current protocol implementations. Moreover, backwards compatibility is easily achieved as legacy routers would simply ignore any Fibbing-specific protocol features. Our algorithms can be modified to account for the fact that the shortest path of a legacy router is never changed by any fake node.

### 4.4 Controller Replication

As any network component, a Fibbing controller can fail at any time. Reliability can be ensured by running multiple copies the Fibbing controller in parallel and connecting them to different places in the network.

No state needs to be synchronized between the replicas besides the input forwarding requirements (mostly static anyway). Indeed, Fibbing algorithms are deterministic, hence all replicas will always compute exactly the same augmented topology. The only dynamic state maintained by a Fibbing replica is the network graph. This state, however, is implicitly synchronized through the shared topology offered by the underlying IGP: the link-state flooding mechanism keeps the network graph up to date and eventually consistent across all replicas.

The determinism of our algorithms enables all replicas to inject (the same) lies at the same time. However, this would increase the amount of flooded information.

To limit control-plane overhead, our implementation relies on a primary-backup architecture with a single active replica and an inexpensive election process. A pre-defined range of router IDs is reserved for controller replicas. The replica with the lowest router ID across all running ones is the active replica. It is the only one injecting lies, while the others only compute the topology augmentation. When the controller is booted or when an active replica fails, running replicas receive IGP messages on the current network topology. Based on those messages, every replica independently infers the possible presence of other replicas; it also checks whether it is the new active replica by comparing its router ID with the one of the other running replicas.

## 5. EVALUATION

We now evaluate Fibbing along three axis. First, we show that existing routers are perfectly capable of handling the extra load induced by Fibbing (§5.1). We then demonstrate the efficiency of Fibbing’s augmentation algorithms in terms of speed and size of the topology (§5.2). Observe that Fibbing behaves as a plain IGP at the network level. Hence, given its negligible impact on single routers and the efficiency of our controller, current ISP networks can be seen as the best large-scale evaluation for Fibbing. We therefore complete the evaluation by illustrating how Fibbing can be used in a realistic case (§5.3).

### 5.1 Router measurements

By increasing the size of the link-state routing topology, Fibbing could increase the CPU and memory overhead on the routers, or slow down protocol convergence. Our experiments demonstrate that the impact on load and convergence time is negligible. All our measurements were performed using OSPF on a recent Cisco ASR9K running IOS XR v5.2.2 equipped with 12GB of DRAM assigned to the routing engine, as well as on a (7-year-old) Juniper M120 running JunOS v9.2, equipped with 2GB of DRAM. Both routers are representative of typical edge devices (*i.e.*, aggregation routers) found in commercial networks. We draw the same conclusions on both router platforms, and focus on measurements collected on the Cisco device in the following.

**Fibbing induces very little CPU and memory overhead on routers.** We first measured the memory increase caused by a growing number of fake nodes (Table 2). Two processes are impacted by the presence of fake nodes: (i) the RIB process, which maintains information about all the routes known to each destination, and (ii) the OSPF process which maintains the entire OSPF topology. Even with a huge number of fake nodes (100,000), the total overhead on both processes was only 154MB—a small fraction of the total memory available. We collected the CPU utilization on the router every minute immediately after we started injecting fake nodes. The utilization was systematically low, at most 4%. This is easily explained as Fibbing relies on OSPF Type-5 LSAs which do not cause the routers to recompute their shortest paths to each other.

# fake nodes	RIB memory (MB)	OSPF memory (MB)
1,000	0.09	0.56
5,000	1.58	5.19
10,000	3.56	10.96
50,000	19.67	56.37
100,000	39.78	113.17

**Table 2: Routers easily sustain Fibbing-induced load, even with huge augmented topologies.**




Figure 8: Evaluation of our augmentation algorithms.

### Fibbing can quickly program forwarding entries.

In a second experiment, we measured how long it took for a router to install a growing number of forwarding entries (Table 3). We injected a growing number of fake nodes, one per destination, and measured the total installation time, by tracking the time at which the router updated the last entry in its FIB. The time to process and install one entry was constant (around  $900\mu\text{s}$ ), independent of the number of entries. This result is several orders of magnitude better than any OpenFlow switches currently on the market [12, 13]. Since installation of forwarding entries is distributed, routers can install their entries in parallel, meaning *Fibbing can program thousands of network-wide entries within 1 second.*

# fake nodes	installation time (s)	avg time/entry ( $\mu\text{s}$ )
1,000	0.89	886.00
5,000	4.46	891.40
10,000	8.96	894.50
50,000	44.74	894.78
100,000	89.50	894.98

Table 3: Programming a forwarding entry in a router with Fibbing is fast, sub 1ms.

**Fibbing does not have any impact on routing-protocol convergence time.** Finally, we compared the total time for routers to converge with and without fake nodes. We failed a link and measured the time for the last FIB entry to be updated considering two cases: (i) no lie was injected and (ii) one lie per destination was injected. Similar to the previous experiments, we repeated the measurements for a growing number of destinations and lies, between 100 and 100,000. In all our experiments, the presence of lies did not have any visible impact. The total convergence times with or without lies were systematically within 4ms, with the router being even faster to converge in the presence of lies in some cases.

## 5.2 Topology Augmentation Evaluation

We now evaluate Simple and Merger (§ 3) according to: (i) the time they take to compute an augmented topology for a given requirement, and (ii) the size of the resulting augmented topology. Results are depicted in Fig. 8. Our evaluation is based on simulation performed on realistic ISP topologies [19], whose sizes range from 80 nodes to over 300. On these topologies, we generated forwarding requirements by randomly changing the next-hop of randomly selected nodes. Destinations of requirement DAGs were also randomly generated.

### Fibbing augments network topologies within ms.

Fig. 8a shows the time (on the y-axis) taken by Simple and Merger for an increasing number of nodes that must change their next-hop (on the x-axis). The plot refers to simulations we ran on the biggest Rocketfuel topology (AS1239). The time taken by Simple to compute the per-destination augmentation varies in the order of milliseconds, ranging from 0.5 ms to 8 ms. While Merger took more time (as expected), its performance is still one order of magnitude lower than the second. For both algorithms, the computation time does not vary much with the number of nodes changing their next-hops.

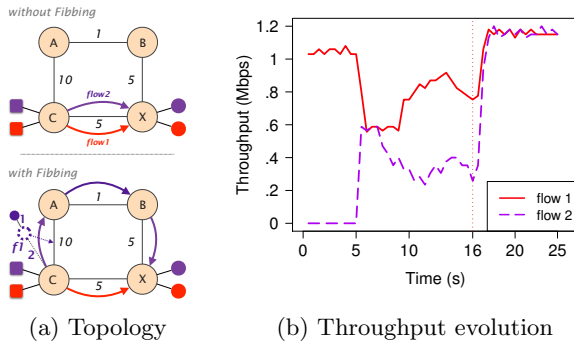
### Merger and cross-optimization effectively reduce the size of the augmented topology.

Fig. 8b plots the fake topology size (on the y-axis) when the number of nodes that have to change their next-hop increases (on the x-axis) for a single destination on all topologies. The plot shows that Merger reduces the number of fake nodes by about 25% in the average case and almost 50% in the best case. Fake topology reduction is further corroborated by our cross-destination optimization procedures (see §3.4). Fig. 8c shows a cumulative distribution function (CDF) of the topology augmentation size computed by Simple, Merger, and Merger with cross-destination optimization. The figure refers to simulations with a number of destinations varying between 1 and 100 with 26% of the nodes changing their next-hop. In more than 90% of our simulations, cross-destination optimization achieves a reduction of the augmented topology. Depending on the experi-

ment, such a reduction is up to about 10% with respect to Merger without cross-destination optimization, and 20% with respect to Simple.

### 5.3 Case Study

We now show the practicality of Fibbing by improving the performance of a real network consisting of four routers (Cisco 3700 running IOS v12.4(3)) connected in a square with link of 1 Mbps capacity (see Fig. 9a). In this network, we introduce two sources (bottom left) that send traffic to two destinations (bottom right) using `iperf`. The first source is introduced at time  $t = 0$ , the second one at time  $t = 5$ . OSPF weights are configured such that all traffic flows along link  $(C, X)$ .



**Figure 9: Case study on how Fibbing can alleviate congestion. At  $t = 16$ , a fake node is added to shift one flow to the upper path in the network, increasing the total available bandwidth.**

Such a network suffers from two inherent inefficiencies: (i) the upper path is never used and (ii) the two flows systematically traverse the same path, competing for bandwidth, no matter what the link weights are. Fig. 9b plots the throughput of each flow. Immediately after the introduction of the second flow at  $t = 5s$ , the two flows start competing for the available bandwidth. To improve network efficiency, the Fibbing controller injects a fake node  $f_1$  connected to  $C$  and announces one destination at time  $t = 16$ . A few ms after the injection, we see that the throughput of both flows double as each of them now traverses a different path.

## 6. REACTION TO FAILURES


We now analyze Fibbing’s reaction to different kinds of failures. We distinguish between network (affecting real router or router-to-router links) and controller (shutting down replicas or replica-to-router links) failures. Also, we separately deal with failures inducing network partitions and non-partitioning ones.

**Fibbing quickly reacts to non-partitioning failures.** Upon network failures, forwarded flows fall in one of the following three cases. First, some flows are not impacted by the failure as their pre-failure forwarding path is not disrupted. Second, flows for which no

input requirements have been specified require only the IGP to establish a new path, but no action from the Fibbing controller. Reaction to failures is extremely fast in this case (sub-second even in large networks), thanks to fast convergence [20] and local fast re-route [21] features commonly supported by current IGP implementations. Third, the remaining flows are forwarded on paths modified by the Fibbing controller. They require reaction from the controller, both to remove possible blackholes or loops due to previously injected lies [14], and to avoid requirement violations due to new IGP paths. Theoretically, the total failure reaction time is equal to the sum of the notification time (for the controller to be notified of the failure), the processing time (for the controller to compute the new topology augmentation) and the IGP convergence time (for all routers to install the new lies). Our evaluation (§5) shows that the processing time is negligible, especially for the Simple algorithm. Moreover, the notification time is bounded by the IGP convergence time, as flooding is faster than re-convergence. Thus, in the worst case, the total reaction time is twice the IGP convergence time, that is, still below 2 seconds [20]. Also, in the average case, the notification time is smaller than IGP convergence, because the controller is notified about the failure before all other routers complete convergence; hence, the controller injects new lies *during* the IGP convergence, and the total failure reaction time is slightly higher than IGP convergence without Fibbing.

In addition, if one or more controller replicas fail but others are running, we have no impact on forwarded flows, unless the failed replica is the active one and some of its injected lies expire before the new active replica is elected. Even in the latter case, the new active replica is quickly elected, in a time which is approximately equal to the detection and flooding of the failure event by the IGP. The short election time makes it unlikely that lies expire before the new active replica is elected, and limits the period with possible disruptions.

**Fibbing can implement both fail-open and fail-close semantics to deal with partitions.** Even if unlikely, catastrophic events like a simultaneous failure of all the controller replicas or network partitions may happen. As for any centralized solution, a major risk in those cases is to leave the network uncontrolled. This happens, for example, if some routers are not reachable by a controller replica after a network partition. With respect to pure SDN solutions, Fibbing has the additional possibility to delegate control to the underlying IGP. This way, Fibbing can implement both the fail-open or fail-close semantics, on a per-destination basis. For non-critical (optimization) requirements like traffic engineering ones, the corresponding destinations can be injected in the IGP, so that connectivity can be preserved as long as the partition leaves at least one source-destination path. For stringent requirements like security ones (*e.g.*, firewall traversal), Fibbing can im-



**Figure 10: Case study on how Fibbing reacts to failures, and can successfully implement fail-close (flow 1) and fail-open (flow 2) semantics.**

plement fail-close semantics by not announcing the corresponding destinations in the IGP. As such, the corresponding flows stop to be forwarded in the absence of the controller. To quickly reach this configuration, we can set a low validity time of the injected lies, making them rapidly expire if not refreshed. This then comes at the cost of additional control-plane overhead.

**We confirmed Fibbing resilience.** We consider again the topology in Figure 9a, and we connect two controller replicas respectively to routers  $A$  and  $B$ . The active replica is initially the one connected to  $A$ . We assume a strict policy on the red flow forcing it to cross the link  $(C, X)$ . We then configure a fail-close semantics to it, and a fail-open to the other flow. Starting from a state in which both replicas and all links are up, we successively fail (i) the active replica at time  $t = 5$ ; (ii) link  $(A, B)$  at  $t = 12$ ; and (iii) link  $(B, X)$  at  $t = 20$ . Finally, we re-establish both failed links, one at the time (at  $t = 36$  and  $t = 48$ ).

The results of this experiment, collected via `iperf`, are reported in Figure 10. Concretely, the *failure of the active replica has no impact* on the forwarded flows. Indeed, the initially passive replica (connected to  $B$ ) quickly detects the failure of the other replica, and start refreshing the injected lies by the failed controller. When  $(A, B)$  fails, the active replica needs to remove the fake node  $f1$ : Since the physical path  $(C, A, B, X)$  is not available anymore, this fake node is creating a loop between  $C$  and  $A$  for the violet flow. Upon failure detection, the controller then sends the LSA to remove  $f1$ , re-establishing the connectivity for the disrupted flow in approximately  $1s$ . Note that this time can be lowered by relying on fast failure detection mechanisms (like BFD). When  $(B, X)$  also fails, we create a partition that makes it impossible for the running replica to interact with routers  $A, C$ , and  $X$ . After about  $1s$ , the injected lies disappear, because they are not refreshed anymore by any controller. Consistent with the configured failure semantics, the red flow is blackholed (to avoid the IGP routing it over policy-violating paths)

while the violet flow keeps using the IGP shortest path. Finally, re-adding the failed links allows the running replica to re-take control of the network: it re-builds a (safe) path for the red flow upon  $(B, X)$  restoration, and re-optimizes the distribution of both flow over the available paths when  $(A, B)$  is restored.

## 7. FREQUENTLY ASKED QUESTIONS

We now provide answers to high-level concerns often raised against Fibbing. Since empirical analyses are hardly applicable to those concerns (*e.g.*, debuggability), we describe qualitative considerations.

**Is Fibbing a long-term solution?** *Yes.* We believe Fibbing is here to stay. In the short run, Fibbing offers programmability and is easy to deploy, at very little cost. A network that ultimately needs even greater flexibility could deploy finer-grained SDN functionality at the edge, and solutions like Fibbing in the core, as advocated by major industry [22] and academic actors [10, 23]. By combining the best of centralized and distributed routing, Fibbing fits the needs of the network core (flexibility, robustness, low overhead) better than current forwarding paradigms.

**Does Fibbing make networks harder to debug?** *No.* Fibbing relies on “tried and true” protocols. This has several implications. First, Fibbing routing matches the current mental model of operators, a major advantage with respect to other SDN proposals. Moreover, Fibbing is compatible with any existing management, monitoring, and debugging tools. Finally, the Fibbing controller can expose a higher-level interface for debugging, including a mapping between the injected lies and their usage (matched requirements and how).

**Does Fibbing sum the complexities of centralized and distributed approaches?** *No.* Fibbing uses the underlying IGP in a very simple way. The IGP output is easy to predict and provides the controller with a powerful API to program routers. As a result, the design of the Fibbing controller is significantly simpler than for existing SDN controllers (*e.g.*, [24, 25, 26, 27]) since heavy tasks such as path computation and topology maintenance are offloaded to the routers. Even basic primitives for controller replication and replica consistency are mainly delegated to current distributed routing protocols (see §6).

**Does Fibbing impact security?** *No.* The lies introduced by the Fibbing controller can easily be authenticated, *e.g.*, using MD5-based authentication [28, 29].

**Since Fibbing can only program loop-free paths, can it support middleboxes chaining?** *Partially.* Forwarding loops can be encountered when steering traffic through a chain of middleboxes (*e.g.*, [30] and [31]). These requirements *can* be satisfied in Fibbing with local support from routers to break the loops. For instance, a router could match on the input interface in

addition to the destination IP address using policy-based routing, a feature widely available on existing routers [32, 33] and provisioned centrally using BGP flowspec [34, 35]. Alternatively, middlebox traffic steering could be implemented through SDN functionality at the network edge, while still using Fibbing in the core.

## 8. RELATED WORK

Fibbing contributes to the larger debate about centralized and distributed control over routing by identifying a new point in the design space.

**Centralized configuration of distributed routing protocols:** A centralized management system can perform traffic engineering by optimizing the link weights in link-state routing protocols [36, 37]. Fibbing is more general, since it can implement *any* forwarding paths by injecting fake nodes and links into the link-state routing topology. The extra flexibility enables even better load balancing, as well as a wider range of functionality.

**Centralized control using existing routing protocols as a control channel:** RCP [11] is a logically-centralized platform that uses BGP to install forwarding entries into routers. RCP must install forwarding entries one-by-one, on each device. In contrast, Fibbing can adapt the forwarding behavior of many routers at once, with little input (*e.g.*, one fake node), and let them compute their own forwarding entries.

**Centralized control over the routing/forwarding tables:** In SDN, a central controller installs packet-processing rules directly in the switches, possibly reacting to the reception of specific packets. While more flexible (*e.g.*, enabling stateful control logic) than Fibbing, SDN requires updating the switch-level rules one-by-one, and forgoes the scalability and reliability benefits of distributed routing. Recently, the IETF developed I2RS [38] which offers a new management interface for centralized updates the routing information bases (RIBs) in the routers. Still, I2RS must push RIB entries individually to each router.

The Fibbing language for expressing requirements is similar in spirit to Merlin [39], but the mechanism for satisfying the requirements (*i.e.*, fake nodes/links) is entirely different. Our main contributions are the Fibbing techniques and algorithms, not the language.

For the networks which require the extra flexibility provided by OpenFlow, Fibbing helps during the transition by providing access to the FIBs of legacy routers to any SDN controller [40]. This contrasts to techniques like Panopticon [41], where programmability is only available in the SDN-enabled parts of the network.

## 9. CONCLUSIONS

The advent of SDN makes it clear that network operators want their networks to be more programmable and easier to manage centrally. In this paper, we show how Fibbing can achieve those objectives, by centrally and

automatically controlling forwarding without forgoing the benefits of distributed routing protocols. Fibbing is expressive, scalable, and works with existing routers. In future work, we plan to look at extensions of IGP protocols (*e.g.*, for source-destination routing [42] or network service header awareness) to enable finer-grained control via Fibbing. Abstractly, Fibbing shows how centralized and distributed approaches can be profitably combined. We believe that new research can further explore this direction, for example, investigating an alternative division of tasks between centralized and distributed network components.

## Acknowledgements

We are grateful to SIGCOMM anonymous reviewers and our shepherd, Teemu Kotonen, for insightful comments. We thank Jo Segart from BELNET and Dave Ward, Clarence Filsfils and Kris Michielsen from Cisco Systems for their support in testing Fibbing on real routers. This work has been partially supported by the EC Seventh Framework Programme (FP7/2007-2013) grant no. 317647 (Leone) and by the ARC grant 13/18-054 from Communauté française de Belgique.

## 10. REFERENCES

- [1] D. Awduche *et al.*, “RSVP-TE: Extensions to RSVP for LSP Tunnels,” RFC 3209, 2001.
- [2] N. McKeown *et al.*, “OpenFlow: enabling innovation in campus networks,” *ACM SIGCOMM CCR*, vol. 38, no. 2, pp. 69–74, 2008.
- [3] A. Farrel, J.-P. Vasseur, and J. Ash, “A Path Computation Element (PCE)-Based Architecture,” RFC 4655, 2006.
- [4] C. Filsfils *et al.*, “Segment Routing Architecture,” Internet Draft, 2014.
- [5] B. Clouston and B. Moore, “Definitions of Managed Objects for HPR using SMIv2,” RFC 2238, 1997.
- [6] D. Oran, “OSI IS-IS Intra-domain Routing Protocol,” RFC 1142, 1990.
- [7] A. Pathak, M. Zhang, Y. C. Hu, R. Mahajan, and D. A. Maltz, “Latency inflation with MPLS-based traffic engineering,” in *IMC*, 2011.
- [8] S. Jain *et al.*, “B4: Experience with a Globally-Deployed Software Defined WAN,” in *SIGCOMM*, 2013.
- [9] C.-Y. Hong *et al.*, “Achieving High Utilization with Software-Driven WAN,” in *SIGCOMM*, 2013.
- [10] M. Casado *et al.*, “Fabric: A retrospective on evolving SDN,” in *HotSDN*, 2012.
- [11] M. Caesar *et al.*, “Design and implementation of a routing control platform,” in *NSDI*, 2005.
- [12] X. Jin *et al.*, “Dynamic scheduling of network updates,” in *SIGCOMM*, 2014.

- [13] C. Rotsos, N. Sarrar, S. Uhlig, R. Sherwood, and A. W. Moore, "OFLOPS: An Open Framework for Openflow Switch Evaluation," in *PAM*, 2012.
- [14] S. Vissicchio, L. Vanbever, and J. Rexford, "Sweet little lies: Fake topologies for flexible routing," in *Hotnets*, 2014.
- [15] S. Vissicchio, O. Tilmans, L. Vanbever, and J. Rexford, "Central Control over Distributed Routing (Extended Version)," Technical Report, 2015.
- [16] L. Vanbever, S. Vissicchio, C. Pelsser, P. Francois, and O. Bonaventure, "Seamless Network-Wide IGP Migrations," in *SIGCOMM*, 2011.
- [17] "Quagga routing suite," [www.nongnu.org/quagga](http://www.nongnu.org/quagga).
- [18] J. Moy, "OSPF Version 2," RFC 2328, Apr. 1998.
- [19] N. Spring, R. Mahajan, and D. Wetherall, "Measuring ISP topologies with Rocketfuel," in *SIGCOMM*, 2002.
- [20] P. Francois, C. Filsfil, J. Evans, and O. Bonaventure, "Achieving Sub-second IGP Convergence in Large IP Networks," *ACM SIGCOMM CCR*, vol. 35, no. 3, 2005.
- [21] C. Filsfil, P. Francois, M. Shand, B. Decraene, J. Uttaro, N. Leymann, and M. Horneffer, "Loop-Free Alternate (LFA) Applicability in Service Provider (SP) Networks," RFC 6571, 2012.
- [22] T. Koponen *et al.*, "Network Virtualization in Multi-tenant Datacenters," in *NSDI*, 2014.
- [23] "Time for an SDN Sequel? Scott Shenker Preaches SDN Version 2," [www.sdxcentral.com/articles/news/scott-shenker-preaches-revised-sdn-sdnv2/2014/10/](http://www.sdxcentral.com/articles/news/scott-shenker-preaches-revised-sdn-sdnv2/2014/10/).
- [24] "ONOS: Open Network Operating System," <http://onosproject.org/>.
- [25] T. Koponen *et al.*, "Onix: A distributed control platform for large-scale production networks," in *OSDI*, 2010.
- [26] "Project Floodlight," <http://www.projectfloodlight.org/floodlight/>.
- [27] N. Foster *et al.*, "Languages for software-defined networks," *IEEE Comm. Mag.*, 2013.
- [28] "Cisco OSPF MD5 Authentication," <http://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/13697-25.html>.
- [29] "Juniper OSPF MD5 Authentication," [http://www.juniper.net/documentation/en\\_US/junos14.2/topics/topic-map/ospf-authentication.html](http://www.juniper.net/documentation/en_US/junos14.2/topics/topic-map/ospf-authentication.html).
- [30] Z. A. Qazi *et al.*, "Simple-fying middlebox policy enforcement using sdn," in *SIGCOMM*, 2013.
- [31] S. K. Fayazbakhsh *et al.*, "Enforcing network-wide policies in the presence of dynamic middlebox actions using flowtags," in *NSDI*, 2014.
- [32] "Cisco. Configuring Policy-Based Routing," [http://www.cisco.com/c/en/us/td/docs/ios/12\\_2/qos/configuration/guide/fqos\\_c/qcfpbr.html](http://www.cisco.com/c/en/us/td/docs/ios/12_2/qos/configuration/guide/fqos_c/qcfpbr.html).
- [33] "Juniper. Configuring Filter-Based Forwarding to a Specific Outgoing Interface or Destination IP Address," [http://www.juniper.net/techpubs/en\\_US/junos12.2/topics/topic-map/filter-based-forwarding-policy-based-routing.html](http://www.juniper.net/techpubs/en_US/junos12.2/topics/topic-map/filter-based-forwarding-policy-based-routing.html).
- [34] "Cisco. Implementing BGP Flowspec," [http://www.cisco.com/c/en/us/td/docs/routers/asr9000/software/asr9k\\_r5-2/routing/configuration/guide/b\\_routing\\_cg52xasr9k/b\\_routing\\_cg52xasr9k\\_chapter\\_011.html](http://www.cisco.com/c/en/us/td/docs/routers/asr9000/software/asr9k_r5-2/routing/configuration/guide/b_routing_cg52xasr9k/b_routing_cg52xasr9k_chapter_011.html).
- [35] "Juniper. Enabling BGP to Carry Flow-Specification Routes," [https://www.juniper.net/documentation/en\\_US/junos12.3/topics/example/routing-bgp-flow-specification-routes.html](https://www.juniper.net/documentation/en_US/junos12.3/topics/example/routing-bgp-flow-specification-routes.html).
- [36] B. Fortz and M. Thorup, "Internet traffic engineering by optimizing OSPF weights," in *INFOCOM*, 2000.
- [37] B. Fortz, J. Rexford, and M. Thorup, "Traffic engineering with traditional IP routing protocols," *IEEE Comm. Mag.*, vol. 40, no. 10, pp. 118–124, 2002.
- [38] A. Atlas, J. Halpern, S. Hares, and D. Ward, "An Architecture for the Interface to the Routing System," Internet Draft, 2013.
- [39] R. Soulé *et al.*, "Merlin: A language for provisioning network resources," in *CoNEXT*, 2014.
- [40] S. Vissicchio, L. Vanbever, and O. Bonaventure, "Opportunities and research challenges of hybrid software defined networks," *ACM SIGCOMM CCR*, vol. 44, no. 2, pp. 70–75, 2014.
- [41] D. Levin, M. Canini, S. Schmid, F. Schaffert, and A. Feldmann, "Panopticon: Reaping the Benefits of Incremental SDN Deployment in Enterprise Networks," in *USENIX ATC*, 2014.
- [42] F. Baker, "IPv6 Source/Destination Routing using OSPFv3," Internet Draft, 2013.