Path Switching in Content Centric and Named Data Networks

Ilya Moiseenko
Cisco Systems
ilmoisee@cisco.com

Dave Oran
Network Systems Research & Design
daveoran@orandom.net

ABSTRACT

ICN communication is inherently multipath and potentially multi-destination. Content Centric and Named Data Networks at present do not offer a mechanism to direct traffic onto a specific path in multipath or a specific destination in a multi-destination environment, because the forwarding plane multiplexes packets across nexthops dynamically. This makes it challenging to provide practical multipath traceroute and ping applications, or implement multipath-aware congestion control, traffic engineering or SDN solutions. The symmetry of forward and reverse paths in Content Centric and Named Data Networks allows one to compute an end-to-end path label in a Data message on the reverse path and subsequently use this label to forward an Interest message through a specific nexthop. ICN Path Switching is a method of high-speed Interest forwarding in Content Centric and Named Data networks based on exact matching of a nexthop label retrieved from the Interest's path label against a nexthop ID in the ICN Forwarder's Adjacency database. ICN Path Switching maintains all major characteristics of CCN / NDN architectures, such as multicasting, caching, flow balance, etc. Simulations demonstrate that path labels are consistent with ICN control plane routing state in the presence of route updates. Analysis of ICN Path Switching with regards to Multiprotocol Label Switching (MPLS) and Segment Routing architectures suggests that it offers similar advantages at lower complexity with the potential to simplify network operations.

CCS CONCEPTS

• **Networks** → **Network protocol design**;

KEYWORDS

ICN; NDN; CCN; high-speed forwarding; traffic engineering; transition; co-existence; MPLS; source routing;

ACM Reference Format:
Ilya Moiseenko and Dave Oran. 2017. Path Switching in Content Centric and Named Data Networks. In *Proceedings of ICN '17*. ACM, New York, NY, USA, 11 pages.
https://doi.org/10.1145/3125719.3125721

ICN '17, September 26–28, 2017, Berlin, Germany

ACM ISBN 978-1-4503-5122-5/17/09…$15.00
https://doi.org/10.1145/3125719.3125721

1 INTRODUCTION

Content Centric Networking (CCN) and Named Data Networking (NDN) are two similar general-purpose, information-centric network (ICN) architectures [9, 14]. ICN's Interest/Data exchange is receiver-driven, with a consumer (receiver) endpoint requesting Data that is provided by a producer endpoint. As a consequence of ICN's stateful forwarding, the Interest's forward (upstream) path and the corresponding Data message's reverse (downstream) path are symmetric. ICN communication is inherently multipath and potentially multi-destination, because an ICN Interest message can be forwarded towards the producer application or cache. In regards to packet forwarding, FIB lookups of the same Interest name in the ICN forwarder may result in multiple Interest message being forwarded over different nexthops. This contrasts with the forwarding behavior of IP networks where all packets of the same flow are forwarded to the same interface[1] in the absence of changes in the control (routing) plane. This dynamic behavior makes it difficult to direct traffic over a specific network path in ICN.

One important open question is how to exploit ICN's multipath capabilities, including how to couple multipath routing/forwarding with associated network objectives including:

- Ability to discover, monitor and troubleshoot multipath network connectivity based on names and name prefixes (somewhat analogous to IP traceroute and ping, or MPLS tree trace [13]).
- Ability to accurately measure network performance, which generally requires multiple packets follow the same path under control of an application.
- Deliver congestion control that provides fairness and maximized network utilization. For instance, allocating resources for consumers / flows when each consumer / flow may have multiple paths to the requested content source (of which there may be multiple instances), each consumer/flow is assured to have a fair share of the network resources, and network utility is maximized.

One approach to this ICN challenge is to try to make multi-path entirely the responsibility of ICN forwarders, with no role for the endpoint consumers. Looking at the network objectives, however, this is difficult.

By not exposing path information to the consumers (both reporting path information to consumers and allowing consumers to use path information to explicitly "steer" packets), diagnostic capabilities similar to traceroute and ping become very difficult to perform.

Network-only approaches either using virtual-circuit style forwarding, or keeping local state about prior packet transmissions in order to properly allocate new packets to paths are known to

[1] or alternative equivalent interfaces in the case of Equal-cost Multipath splitting

be complicated, not scale well, and do not give consumers enough information to accurately control their rate of transmission.

This paper presents techniques of Path Discovery and Path Steering in ICN networks that are enabled by the symmetry of forward and reverse paths in CCN / NDN networks. Path discovery is achieved by a consumer endpoint transmitting an ordinary Interest message and receiving a Content (Data) message containing an end-to-end path label constructed on the reverse path by the forwarding plane. Path Steering is achieved by a consumer endpoint including a path label in the Interest message, which is forwarded to a nexthop through the corresponding egress interfaces in conjunction with longest name prefix match (LNPM) FIB lookup. In section 2.4 we describe what ICN applications, techniques and use cases become possible or might benefit from ICN Path Discovery and Steering.

Another very important issue in ICN is how to forward packets faster. CCN and NDN rely on expensive LNPM FIB lookup to forward Interest messages. Despite many efforts [2, 22–24, 26, 27] to scale ICN FIB lookup to the performance of forwarding in state-of-the-art IP and Multiprotocol Label Switching (MPLS) networks or bypass FIB lookup [6–8, 12, 17], name-based forwarding remains a major problem. It is also premature to conclude that conventional ICN FIB lookup can be accelerated with specialized hardware (e.g. ASIC, FPGA). LNPM FIB lookup against hierarchical tokenized names has daunting technical challenges, due to long variable size message names and TCAM memory size limitations in modern commercial-grade network processing chips.

ICN Path Switching is a new method of high-speed Interest forwarding in Content Centric and Named Data networks. It employs exact matching of a *nexthop label* retrieved from a path label present in the Interest message against a *nexthop ID* in the ICN Forwarder's Adjacency database. In addition to the use of Path Discovery & Steering, ICN Path Switching includes a novel technique that bypasses the LNPM FIB lookup while keeping the path(s) consistent with the actual routing plane state.

The remainder of the paper is structured as follows. Section 2 describes how the concepts of ICN Path Discovery / Steering can be embedded in standard CCN / NDN forwarding plane. Section 3 explains ICN Path Switching. Section 4 contains side-by-side simulations of ICN Path Switching and regular CCN / NDN forwarding plane in dynamic network environments. Section 5 discusses new security challenges introduced by ICN Path Switching. Section 6 compares ICN Path Switching to MPLS, Segment Routing and Dynamic Source Routing. Section 7 concludes.

2 ICN PATH DISCOVERY & STEERING

This section explains the concepts of a path label, path discovery and path steering in Content Centric and Named Data Networks as an extension of regular LNPM FIB-based forwarding. The next section describes how path steering works without LNPM FIB-based forwarding.

End-to-end Path Discovery in an ICN network is achieved by constructing a *path label* as part of an ICN Content (Data) message as the message traverses the reverse path of transit ICN forwarders (Figure 1). The path label is updated by adding the nexthop label of the interface at which the Content(Data) message has arrived to the existing path label. Eventually, when the Content(Data) message arrives at the consumer, the path label identifies the exact path the Content(Data) message took to reach the consumer.

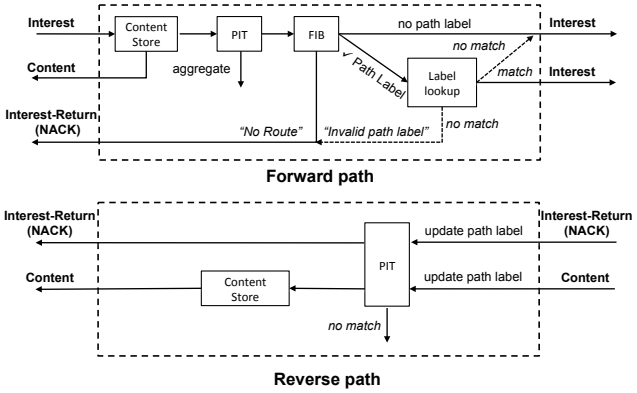

Figure 1: Path Discovery and Steering in regular CCN / NDN forwarding plane.

Due to the symmetry of forward and reverse paths in ICN networks, a consumer application can reuse this newly discovered path label to fetch the same or similar (e.g. next chunk or next Application Data Unit) Content (Data) message over the discovered network path. *Path Steering* in an ICN network is achieved by processing the Interest message's path label at each transit ICN forwarder and forwarding the Interest through the specified nexthop among those identified as feasible by LNPM FIB lookup (Figure 1).

2.1 Adjacency database

The Adjacency database is a required data structure in a router. It is used to store and access nexthop information in the router's forwarding plane. In general, adjacency database entries are pointed to by corresponding FIB entries to express the information associated with a nexthop. In the case of layer 2 protocols like Ethernet, they also store the mapping of a nexthop IP address or some other nexthop identifier to a destination MAC address. The CCN / NDN FIB maps the prefix to the nexthop ID as shown in Figure 2. The Adjacency database is updated in case of L1/L2 state changes or through configuration changes.

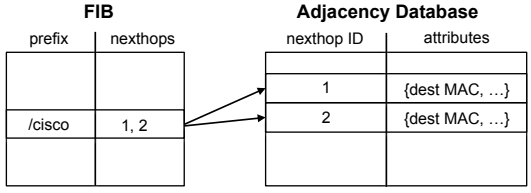

Figure 2: CCN / NDN FIB and Adjacency database.

2.2 Example of path discovery & steering

Figure 3 illustrates the process of updating a path label, path discovery and path steering in a simple network topology with a single forwarding split point.

Three ICN messages are exchanged in this example.

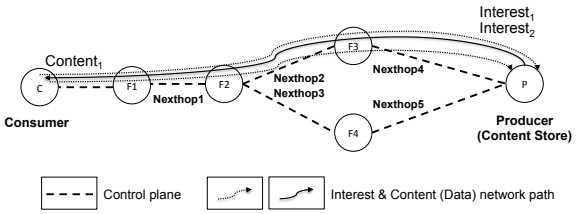

Figure 3: Basic example of path discovery and steering.

(1) $Interest_1$ addresses the content served by a producer application. $Interest_1$ is forwarded by $Forwarder_1$ via $Nexthop_1$, $Forwarder_2$ via $Nexthop_2$, $Forwarder_3$ via $Nexthop_4$ with regular LNPM FIB lookups and eventually reaches the producer application serving the matching $Content_1$.

(2) $Content_1$ is forwarded downstream on the reverse path. $Forwarder_3$ receives $Content_1$ from $Nexthop_4$, encodes $Nexthop_4$ in the empty path label of $Content_1$, and forwards $Content_1$ downstream via regular PIT matching.

$Forwarder_2$ encodes $Next-hop_2$ in the existing path label of $Content_1$ and forwards $Content_1$ downstream.

$Forwarder_1$ encodes $Next-hop_1$ in the existing path label of $Content_1$ and forwards $Content_1$ downstream.

(3) Consumer has received $Content_1$ and discovered a new path label encoding $Nexthop_4$, $Nexthop_2$, and $Nexthop_1$. Now consumer transmits $Interest_2$ with the same path label in order to fetch the next chunk of the content with the intention of taking the same path to the producer application (or ICN Forwarder cache).

LNPM FIB lookup on $Forwarder_1$ returns one nexthop: $Nexthop_1$. $Forwarder_1$ decodes $Nexthop_1$ from the path label of $Interest_2$, compares it to the available nexthops, finds a match and forwards $Interest_2$ through $Nexthop_1$.

LNPM FIB lookup on $Forwarder_2$ returns two nexthops: $Nexthop_2$ and $Nexthop_3$. $Forwarder_2$ decodes $Nexthop_2$ from the path label of $Interest_2$, compares it to the available nexthops, finds a match and forwards $Interest_2$ through $Nexthop_2$.

LNPM FIB lookup on $Forwarder_3$ returns one nexthop: $Nexthop_4$. $Forwarder_3$ decodes $Nexthop_4$ from the path label of $Interest_2$, compares it to the available nexthops, finds a match and forwards $Interest_2$ through $Nexthop_4$.

2.3 Path label

The challenge of encoding a path label can be formulated more formally as a problem of encoding N numbers (i.e. nexthop IDs as nexthop labels) $\{x_1...x_n\}$ in a space-efficient manner such that that numbers $\{x_1...x_n\}$ can be recovered afterwards.

2.3.1 Bloom filter. One possible mechanism is to use a Bloom filter path description to store the hops in the path [12]. A Content (Data) message carries a Bloom filter header field constructed from the node+link identifiers chosen at each split point. The consumer includes this Bloom filter in future Interests for the corresponding path. At split points, the forwarder looks up, in the received Bloom Filter, the node+link for each nexthop in the FIB entry. A hit indicates which nexthop should be used. The first disadvantage of this approach is that Bloom Filters are probabilistically incorrect. The typical implementation produces false positives: a node+link is reported as being a member of the Bloom Filter set though it is not. A second disadvantage is that the identifiers in the Bloom Filter must be unique network-wide since the lookup at any hop can find any member of the Bloom Filter set. Network-wide unique identifiers increase the cost of the scheme, so, unless the network-wide unique identifier provides a benefit, the local identifiers would be a superior approach.

2.3.2 Pairing function encoding. The Cantor pairing function [5] allows one to uniquely encode two natural numbers into a single natural number $\langle x, y \rangle$. An inverse Cantor pairing function also exists that allows one to recover both values from a single natural number $\langle x, y \rangle$.

A pairing function can be used recursively to encode an arbitrary number of natural numbers (i.e. pairings of pairings). This makes it possible to encode the whole network path in a single natural number. Unlike a Bloom Filter, an inverse pairing function is completely deterministic.

Unfortunately, our experiments with hop-by-hop pairing demonstrated that the path label value grows quickly and exceeds 64 bits due to the quadratic nature of the Cantor pairing function. Since 512 or 1024 bits (and a *bignum* square root calculation) are likely to be necessary to support larger network sizes it is doubtful the forwarding plane could perform encoding and decoding operations fast enough. Other pairing and encoding functions [19][25] might have more compact distribution and, therefore, be more suitable for the purpose of recursive path label construction.

2.3.3 Polynomial encoding. A faster way of manipulating path labels whose size exceeds 64 bits is to allocate to each transit ICN forwarder a fixed size portion of the bit array. Allocating 12 bits (i.e. 4095 as a 'generator polynomial') to each intermediate ICN forwarder seems to match the scalability of today's commercial routers that support up to 4096 physical and logical interfaces [1] and usually do not have more than a few hundred active ones. In this approach an ICN forwarder that receives a Content (Data) message encodes the nexthop label in the next available slot and increments label index (Figure 4). Conversely, an ICN forwarder that receives an Interest message reads the current nexthop label and decrements label index. The advantages of this approach are 1) fixed size of the path label, and 2) ease of encoding / decoding. The disadvantage is the hard limit on the maximum number of network hops.[2]

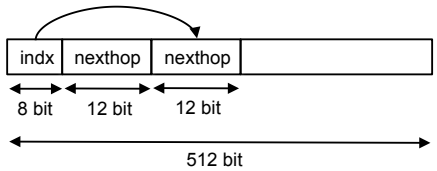

Figure 4: 512 bit path label can encode a 42-hop long path.

[2] A fixed size path label might not be a practical problem in CCN, since the size can be pre-set during Interest message encoding based on the consumer specified value of the hop-limit.

2.3.4 Stack encoding. The path label can also be encoded similarly to an MPLS label stack. In this approach, an ICN forwarder that receives a Content (Data) message encodes the nexthop label and pushes it onto the path label stack. Conversely, an ICN forwarder that receives an Interest message pops the nexthop label from the path label stack (Figure 5). One advantage of this approach lies in its potential hardware compatibility with MPLS technology. The obvious disadvantage is the path-length dependent size of the label stack.

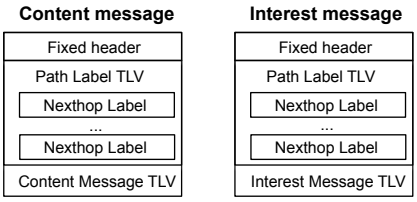

Figure 5: Path label encoded as a stack of nexthop labels.

2.4 Path Discovery & Steering use cases

The example above illustrates how the mechanism of path discovery and steering works in conjunction with LNPM FIB lookup. While adding some overhead to the forwarding plane[3], path steering mechanism enables many important applications and techniques that are not possible with the regular CCN / NDN forwarding method.

- An ICN Ping application can reliably measure path RTT by ensuring that subsequent requests are sent over the specified network path instead of traversing alternating network paths [20].
- An ICN Traceroute application can iteratively discover multiple network paths by directing the requests onto the partially explored paths [21].
- Consumer endpoint multipath-aware congestion control can discover and load balance across multiple available network paths [15].
- A consumer endpoint can mitigate content poisoning attacks by directing the Interests onto the network paths away from the poisoned caches.
- The path steering mechanism can serve as a foundation for overlay networks over CCN / NDN. For example, an overlay broadcast network might be used to scale multiparty synchronization protocols, such as ChronoSync [28].
- Traffic engineering (TE) and SDN solutions can be built to distribute precomputed path labels to the ingress routers to manage network traffic.

2.5 Effect of route updates

Over time, the state of interfaces and FIB on forwarders may change such that, at any particular ICN forwarder, a given nexthop is no longer valid for a given prefix. In this case, the path label will point to a now-invalid nexthop. This will be detected by failure to find a match between the decoded nexthop ID and the nexthops of the FIB entry after LNPM FIB lookup (Figure 1).

One useful action at this point is to respond to the Interest with an Interest-Return (NACK) carrying a new "Invalid path label" response code and include the current path label in it. Each transit ICN forwarder processing the Interest-Return (NACK) messages updates the path label in the same manner as Content (Data) messages (Figure 1), so that the consumer receiving the Interest-Return (NACK) can easily identify which path label is no longer valid.

An alternative behavior for the router detecting the inconsistency is to forward the Interest by means of LNPM FIB lookup. The consumer endpoint, if it cares, can keep enough information about outstanding Interests to determine if the path label sent with the Interest fails to match the path label in the returned Content (Data), and replace stale path labels.

To summarize, the matching of the decoded nexthop ID to the FIB entry nexthops guarantees that FIB changes invalidating a path label are detected in all cases, and that the Interest always follows a valid path to the producer.

3 ICN PATH SWITCHING

Because of the limited network MTU size, content objects meaningful to applications are in general larger than a single MTU. This requires multiple Interest/Data exchanges by the consumer in order to obtain useful results. One example is a video streaming application fetching multiple Content (Data) messages (e.g. chunks) of a video stream. Path label(s) become available to the consumer application after a single RTT when the first Content (Data) messages start arriving. When employed by a congestion control scheme like [15], any subsequent Interest sent by the consumer can include (one of) the discovered path label(s). Depending on the traffic pattern, Interests carrying path labels could comprise more than 90% of Interest messages in the network. Traffic engineering (TE) and SDN solutions could raise this ratio to 100%.

ICN Path Switching further leverages the Path Label idea to bypass the LNPM FIB lookup, which is the most expensive operation in conventional ICN forwarding. The key constraint is to ensure that when skipping the regular FIB lookup, the paths remain consistent with the corresponding FIB and routing plane routes.

A 10,000 ft view on ICN Path Switching is as follows:

- Content (Data) messages are forwarded based on PIT lookup as in regular CCN / NDN. This preserves ICN flow balance and multicast properties.[4]
- Interest messages undergo the same Content Store and PIT lookup as in regular CCN / NDN. This preserves ICN caching and Interest aggregation properties.
- ICN Path Switching enables the same new use cases, techniques and applications as ICN Path Discovery & Steering (Section 2.4) with an additional benefit of "fast path" Interest forwarding.

3.1 Forwarding plane

Figure 6 illustrates the changes in CCN/NDN forwarding plane introduced with Path Switching technology. With these changes

[3] one could argue that this is in fact not additional overhead, at least compared to many other techniques for choosing among matching nexthops of a FIB entry

[4] In contrast to [8], we do not consider a PIT-less design.

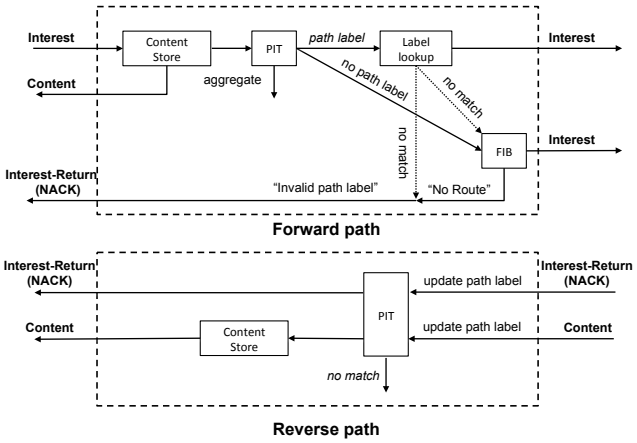

Figure 6: Path Switching CCN / NDN forwarding plane.

the ICN Path Switching forwarding plane performs the following actions:

(1) When an Interest arrives at an ICN forwarder, it is matched to Content Store and PIT to find a cached Data message or aggregate an Interest. If there is no match in CS or PIT, the ICN forwarder attempts to forward the Interest.
(2) If the Interest does not carry a path label, the ICN forwarder performs regular LNPM FIB lookup and determines the nexthop. The corresponding nexthop label is added to the new PIT entry (Figure 7) and the Interest is forwarded.
(3) If the Interest carries a path label, the ICN forwarder performs a lookup of the nexthop label in the forwarder's adjacency database.
(4) If a match is found, 1) the nexthop label is added to the new PIT entry (Figure 7), 2) the Interest's path label is updated (e.g. topmost nexthop label is removed), and 3) the Interest is forwarded through the nexthop with matching nexthop label.
(5) If nexthop label lookup fails, the Interest cannot be path switched. In this case, the ICN forwarder either performs LNPM FIB lookup, discards the path label and forwards the Interest upstream, or replies with "Invalid path label" Interest-Return (NACK) message (Section 2.5).
(6) If LNPM FIB lookup fails, a "No-Route" Interest-Return (NACK) message is sent with the Interest's path label.
(7) When a Content (Data) or Interest-Return (NACK) message arrives at an ICN forwarder, a nexthop label of the ingress interface is added to the path label. The Content (Data) or Interest-Return (NACK) message is matched to PIT and Content Store (if applicable).
(8) In addition to the regular PIT matching procedure, the nexthop label of an interface at which the Content (Data) or Interest-Return (NACK) message has arrived must match the nexthop label stored in the PIT entry (Figure 7).

Interest	Ingress Interfaces	Egress Interfaces	Nexthop Labels

Figure 7: PIT entry is extended to store nexthop label(s).

3.2 Effect of route updates

ICN Path Switching does not depend on any specialized control plane protocol, because the path label is constructed entirely by the data plane. However, without care in the path matching procedure, Interests might be forwarded according to the path label even if the nexthop ID has been removed from the corresponding FIB entry or the whole FIB entry has been removed from the FIB, or the nexthop associated with a different FIB entry no longer valid for the corresponding name.

Our design decouples stable nexthop IDs as used in the management and control plane from nexthop labels used in ICN Path Switching. We add an additional "nexthop label" in the adjacency database entry (Figure 8). Any time a nexthop is removed from a FIB entry it receives a new random nexthop label. If the whole FIB entry is removed from the FIB, all nexthops referenced by the FIB entry receive new random nexthop labels. The adjacency database scales with the number of adjacencies the router has and does not depend on the size of FIB; nor does it increase in size on label regeneration events as there is only one label per adjacency entry.

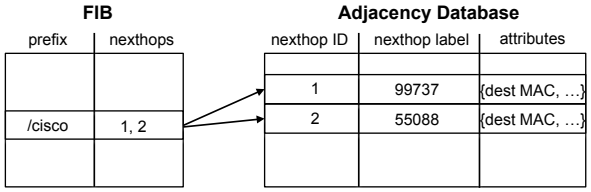

Figure 8: Adjacency entry stores temporary label.

This technique instantly and reliably breaks the path every time a route update changing the FIB occurs. On the forward path, a nexthop label decoded from an Interest path label will not match any of the entries in the adjacency database. On the reverse path, a nexthop label of an interface at which the Content (Data) or Interest-Return (NACK) message has arrived upon will not match the nexthop label in the PIT entry. The purpose of nexthop label matching with the PIT entry is to prevent the consumer endpoint from discovering a potentially invalid path label — a path that does not exist in the control plane. Alternative approaches to dropping a Content (Data) or Interest-Return (NACK) message on PIT entry nexthop label mismatch are: (1) to remove the path label and forward the message downstream, (2) encode the nexthop label taken from the PIT entry and forward the message downstream, or (3) perform LNPM FIB lookup to verify that the interface at which the message has arrived upon is still a valid nexthop, encode the nexthop label and forward the message downstream.

Regeneration of a nexthop label creates a side effect if 1) the same nexthop is referenced from multiple FIB entries with no common name prefix, 2) at least one of these FIB entries is removed or updated, and 3) there is ongoing ICN communication over the other name prefix. In this case, impacted nexthops get new random nexthop labels, which causes "Invalid path label" Interest-Return (NACK) messages for Interest messages carrying the other name prefix. A consumer might have to wait a few RTTs before the prior path can be rediscovered.

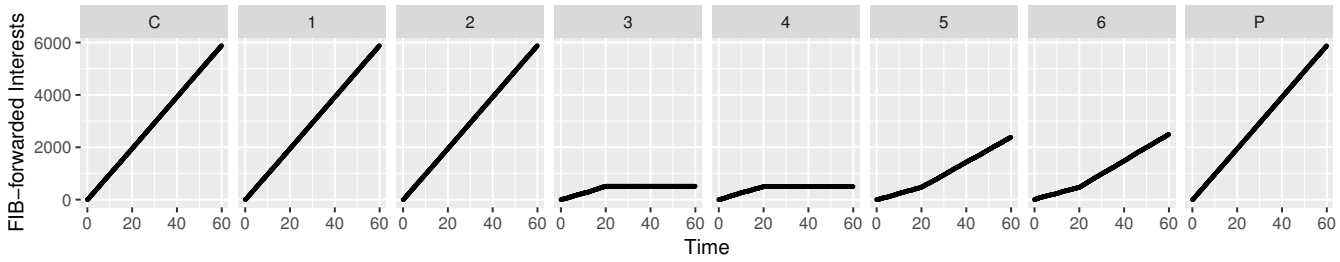

(a) Per-node statistics of regular CCN / NDN forwarding plane.

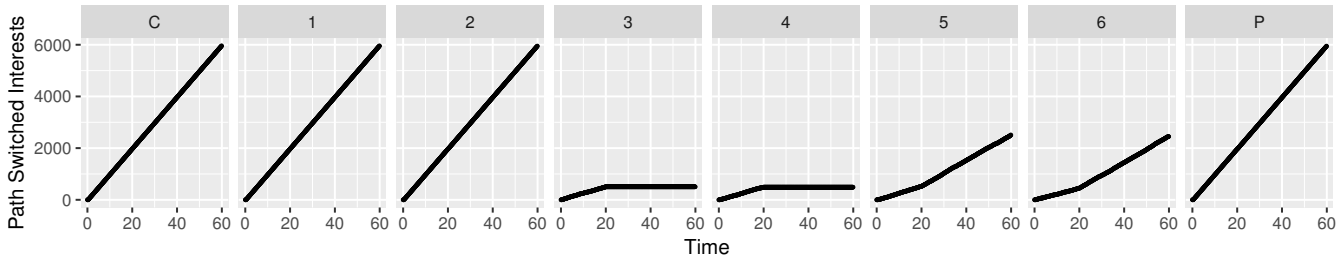

(b) Per-node statistics of path switching CCN / NDN forwarding plane.

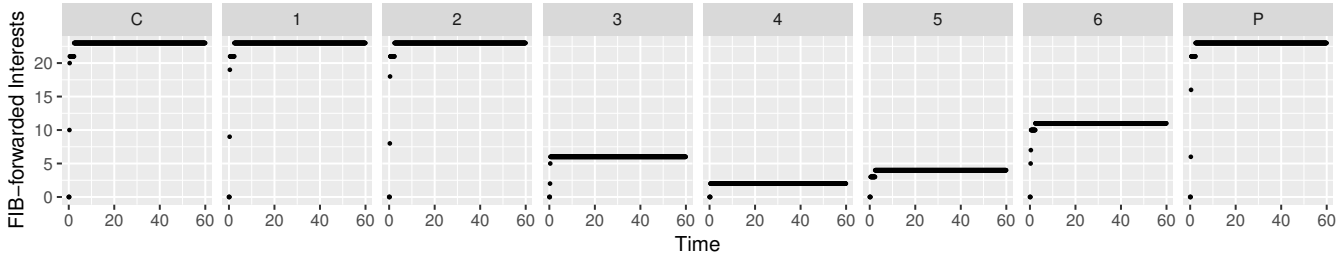

(c) Per-node statistics of path switching CCN / NDN forwarding plane.

Figure 9: Regular vs. path switching CCN / NDN forwarding plane before and after single name prefix route update.

4 EVALUATION

The main risk of using ICN Path Switching in the fast path of the ICN forwarding plane is the potential inconsistency between the contents of the FIB as contructed by the control plane at each ICN forwarder and path labels used by the ICN applications and forwarders. In this section, we conduct two simulations in ndnSIM 2.2[16] to study the behavior of ICN Path Switching in the presence of route updates. We also provide performance measurements comparing the regular NDN forwarding daemon (NFD) to a path switching NFD. These measurements, while promising, are not comprehensive and not meant to be representative of real-life performance. The present work is focused on validating the most critical aspects of the protocol.

4.1 Single name prefix route update

Figure 10 illustrates a simple multipath network topology where a producer node is reachable through four equal cost multiple paths (ECMP). As a baseline, we use a regular NDN forwarding plane with a forwarding strategy selecting a random ECMP nexthop. A Consumer endpoint transmits Interests at a constant rate of 100 messages per second for 60 seconds, which roughly corresponds to fetching 6 MB of data. There is no congestion and no packet loss in this simulation. Two route updates on Node 2 remove the adjacency with Node 3 and Node 4 after 20 seconds of simulation.

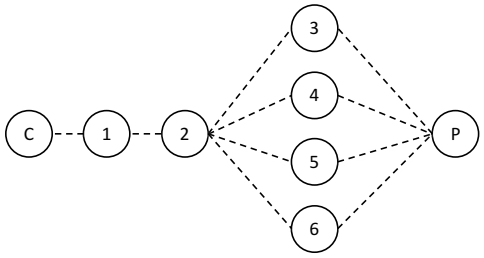

Figure 10: Producer is reachable through 4 equal cost paths.

Figure 9a shows how a regular NDN / CCN forwarding plane with a forwarding strategy selecting a random ECMP nexthop balances the traffic across the available egress interfaces. During the initial 20 seconds Node 2 forwards an equal number of Interests to each of the four ECMP nexthops. After the route update, Node 2 forwards an equal number of Interest to the two remaining ECMP nexthops. Producer receives all Interests.

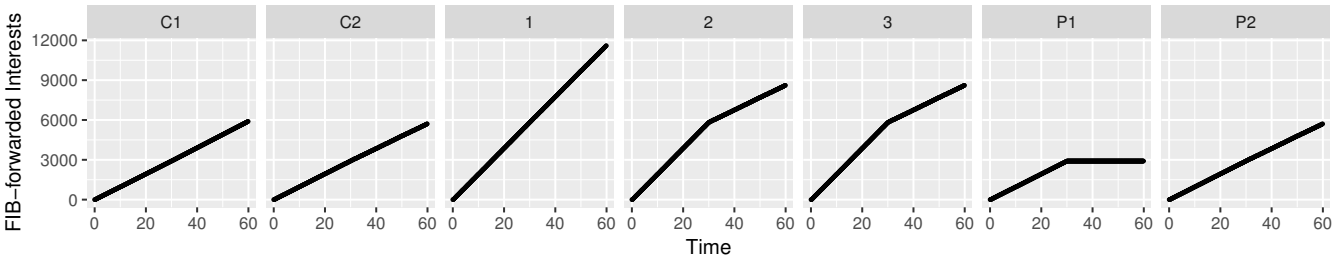

(a) Per-node statistics of regular CCN / NDN forwarding plane.

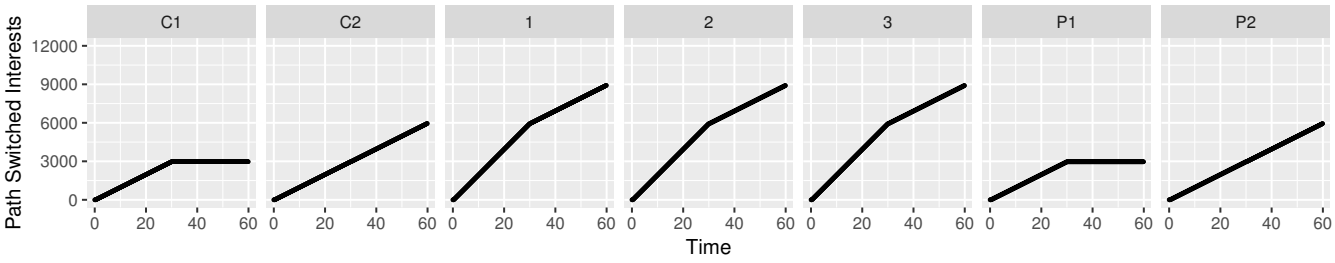

(b) Per-node statistics of path switching CCN / NDN forwarding plane.

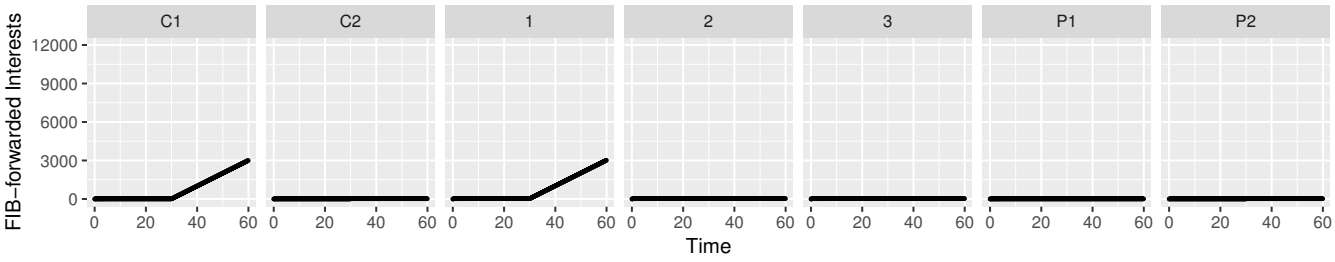

(c) Per-node statistics of path switching CCN / NDN forwarding plane.

Figure 11: Regular vs. path switching CCN / NDN forwarding plane before and after multiple name prefix route update.

The same experiment (e.g. topology, routing plane, forwarding strategy, etc.) is repeated for ICN Path Switching.

Forwarder transmits an "Invalid path label" Interest-Return (NACK) message if label lookup fails on the forward path, and drops Data (Content) and Interest-Return (NACK) messages if label lookup fails on the reverse path. Stack encoding is used to encode a path label.

Consumer supports basic Path Discovery & Steering. First 20 Interests are transmitted without any path label and all path labels retrieved from the Content (Data) messages are stored in a *path table* at the consumer. All the remaining Interests are transmitted with a randomly selected path label from the path table. Upon the receipt of an "Invalid path label" NACK message, the consumer removes the path label from the path table.

A more sophisticated design of a "Path Switching-aware" consumer application would most likely: a) transmit a small percentage of Interests in the path steering phase without any path label in order to periodically probe for new paths, b) maintain per-path statistics such as RTT, packet loss, jitter, etc. to rank and distribute load across available paths, and c) perform multipath aware congestion control, such as [15].

Figure 9 shows two key metrics of ICN Path Switching side by side. The number of Interests forwarded through LNPM FIB lookup (Figure 9c) is low and corresponds to the number of Interests transmitted by the consumer endpoint in the path discovery phase. The number of Interests forwarded in the result of the exact match of nexthop labels (Figure 9b "Path switched Interests") closely tracks the number of Interests forwarded in the regular NDN / CCN forwarding plane (Figure 9a).

This experiment demonstrates that ICN Path Switching properly reacts to the changes in ICN control plane state. Whether and to what extent multiple paths are discovered and used for path steering or switching depends on the dynamic behavior of ICN forwarders. A forwarding strategy using a random ECMP is just one of a number of dynamic forwarding methods that produce multiple paths. Conversely, if all ICN forwarders use a forwarding strategy providing only one "best route" egress nexthop, a consumer endpoint can discover only one end-to-end path.

4.2 Multiple name prefix route update

Often FIB entries with no common name prefix have a common nexthop. Regeneration of a nexthop label of a common nexthop might impact other ongoing Interest/Data transmissions (Section 3.2).

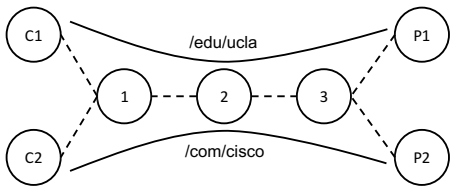

Figure 12: Separate name prefixes use a common nexthop.

Figure 12 illustrates two applications communicating over the shared nexthops. In a first application, Consumer 1 and Producer 1 exchange Interest/Data messages in the /edu/ucla namespace. In a second application, Consumer 2 and Producer 2 exchange Interest/Data messages in the /com/cisco namespace. Node 1 and Node 2 have two FIB entries: /edu/ucla and /com/cisco, both referencing the same nexthop.

Both consumers transmit Interests at a constant rate of 100 messages per second for the duration of 60 seconds. There is no congestion and no packet loss in the network. A route update on Node 2 removes /edu/ucla FIB entry after 30 seconds of simulation.

Figure 11a shows how a conventional CCN / NDN forwarding plane reacts to the route update. Producer 1 does not receive any new Interests after 30 seconds of simulation. Node 1 still receives and forwards all the Interests from both consumers and Node 2 drops the half of the Interests it receives after 30 seconds of simulation.

Figures 11b and 11c show how an ICN Path Switching forwarding plane reacts to the route update. Producer 1 does not receive any new Interests after 30 seconds of runtime. Consumer 1 does not transmit Interests with path labels after 30 seconds of runtime, because it has received an "Invalid path label" NACK message from Node 2. After 30 seconds Consumer 1 only transmits Interests without any path label. These Interests are forwarded by Node 1 in the result of LNPM FIB lookup, and dropped at Node 2. Although Consumer 2 receives an "Invalid path label" at the same time as Consumer 1, Interests from Consumer 2 are forwarded to the Producer 2 without any significant disruption, because Consumer 2 "re-discovers" the path. After 30 seconds of simulation, Nodes 1, 2, and 3 do not receive as many path switched Interests as in the beginning, because Consumer 1 does not have any valid path label to use.

This experiment demonstrates that route updates leading to the regeneration of a nexthop label do not introduce a significant negative impact on other ongoing Interest/Data transmissions in the unrelated namespaces.

4.3 Forwarding scalability

Our initial evaluation plan was to apply the same code changes made in NDNsim to NDN Forwarding Daemon (NFD) and measure its performance in both modes. However, the performance tests did not demonstrate any radical improvement due to the limitations of NFD architecture, such as dynamic memory allocations, non-zero copy packet I/O, shared underlying data structure for both PIT and FIB, etc. These constraints dominate the results of performance tests and render all forwarding path improvements minor.

NFD	Interests	Total Interest processing delay minus total PIT insertion delay
Regular	100,000	377120 microseconds
Path Switching	100,000	1057 microseconds

Table 1: PIT-FIB benchmark of path switching NFD.

The results in Table 1 were obtained with a modified PIT-FIB benchmark, which comes with the NFD software. This benchmark models ICN traffic only with regards to PIT and FIB operations (e.g. lookups and insertions). It does not model Content Store operations or contain any other logic of NFD, so it essentially produces overly optimistic estimation of NFD forwarder performance. This is not critical for our comparative study.

We model an ICN traffic of 100,000 unique Interests with varying name length in the range from one to eight name components. There is no mixed use of path labels in the same run — either all Interests carry a path label or none. FIB was populated with 100,000 name prefixes with length not exceeding eight name components. As in full NFD, dynamic memory allocation during the PIT insertion contributed the major part of Interest processing delay, so we measured it separately to clean the results in Table 1.

This experiment demonstrates that ICN Path Switching has much better scaling properties than regular NDN forwarding plane with LNPM FIB lookup, because the speed of nexthop label lookup does not depend on the size and the contents of FIB.

5 SECURITY CONSIDERATIONS

A path in ICN Path Switching is invalidated by renumbering nexthop label(s). Can a malicious consumer mount an attack by transmitting Interests with path labels which differ only in a single now-invalid nexthop label in order to "brute force" a valid nexthop label? If such an attack succeeds, a malicious consumer would be capable of steering Interests over a network path that potentially does not match the paths computed by the routing plane.

When a label lookup fails, an "Invalid path label" Interest-Return (NACK) message is delivered to the consumer. This contains a path label identical to the one included in the corresponding Interest message. A malicious consumer can analyze message's Hop Count field to infer which specific nexthop label had failed[5] and launch a brute force attack to discover a valid nexthop label on the router that changed its nexthop label(s).

This threat can be mitigated by the following countermeasures:

- A nexthop label of larger size is harder to crack. If nexthop labels are not allocated in a predictable fashion by the routers, brute forcing a 32-bit nexthop label requires on average 2^{31} Interests.
- An ICN forwarder can periodically update nexthop labels to limit the maximum lifetime of paths.[6]

[5] A malicious consumer might not be able to reconstruct a path label encoded with a pairing function.

[6] Route flapping at a path-switching enabled forwarder cannot degrade forwarding performance to a level worse than in a conventional NDN / CCN forwarder.

- A void Hop Count field in an "Invalid path label" Interest-Return (NACK) message does not give out the information on which specific nexthop label had failed. An attacker might need to brute force all nexthop labels in all combinations.

An implementation of ICN Path Switching where a label lookup failure in the forward path is handled by LNPM FIB lookup (Section 2.5) is potentially less secure, because in some cases a malicious consumer will be able to infer a failed nexthop label by comparing the path label sent with the Interest to the path label received in the Content (Data) message.

ICN protocols can be susceptible to a variety of cache poisoning attacks, where a colluding consumer and producer arrange for bogus content (with either invalid or inappropriate signatures) to populate router caches. These are generally confined to on-path attacks, but ICN path switching introduces a new attack vector whereby a path label is used inappropriately by a consumer to cause bad content of its choosing to poison caches.

Given the route exists between a consumer (*C_eve*) and a producer (*P_eve*), *C_eve* can transmit an Interest for the name /eve/foo and learn the path label to *P_eve*. *C_eve* can then send an Interest for /com/nytimes/index with the newly discovered path label to *P_eve*. Since LPNM FIB lookups are bypassed when a valid path label is present in an Interest, *P_eve* can inject an off-path and potentially bogus /com/nytimes/index object in on-path caches In order to foil this attack, objects returned with a path label must have their CS entries annotated with the corresponding path label and only used to satisfy Interests with a matching path label. To also ameliorate cache pollution, such CS entries should not evict entries for the same object with no path label, or a different path label.

It is also theoretically possible to launch a similar attack without a cooperating producer such that the caches of on-path routers become poisoned with the content from off-path routers (i.e. physical connectivity, but no route in a FIB for a given prefix). We estimate that without any prior knowledge of the network topology, the complexity of this type of attack is in the ballpark of Breadth-First-Search and Depth-First-Search algorithms with the additional burden of transmitting 2^{31} Interests in order to crack a nexthop label on each hop. Relatively short periodic update of nexthop labels and anti-"label scan" heuristics implemented in the ICN forwarder may successfully mitigate this type of attacks.

5.1 Cryptographic protection of a path label

If the countermeasures listed above do not provide sufficient protection against malicious mis-steering of Interests, the path label can be made opaque to the consumer endpoint via hop-by-hop symmetric cryptography applied to the path labels (Figure 13). This method is viable due to the symmetry of forward and reverse paths in CCN / NDN networks and ICN path switching requiring only reads/writes of the topmost nexthop label (i.e. active nexthop label) in the path label. This way a path switching ICN forwarder receiving a Data (Content) message encrypts the current path label with its own non-shared symmetric key prior to adding a new nexthop label to the path label. The Data (Content) message is forwarded downstream with unencrypted topmost (i.e active) nexthop label and encrypted remaining content of the path label. As a result, a consumer endpoint receives a Data (Content) message with a unique path label 'exposing' only the topmost nexthop label[7]. A path switching ICN forwarder receiving an Interest message performs label lookup using the topmost nexthop label, and decrypts the path label with its own non-shared symmetric key and forwards the message upstream.

Cryptographic protection of a path label does not require any key negotiation among ICN forwarders, and is no more expensive than MACsec or IPsec. It is also quite possible that strict hop-by-hop path label encryption is not necessary and can be replaced with path label encryption only on the border routers of the trusted administrative or routing domains.

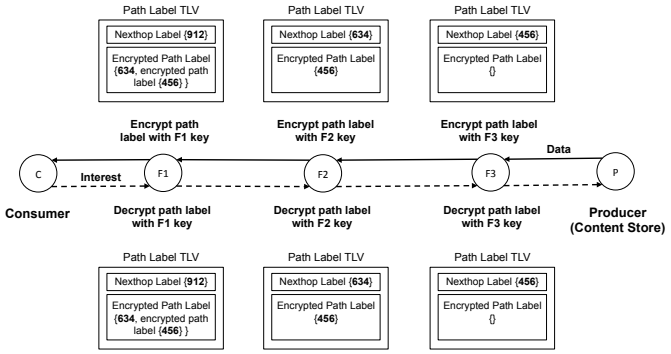

Figure 13: Path label protection with hop-by-hop symmetric cryptography.

6 RELATED WORK

Many organizations critically depend on network performance and availability. For Internet service providers, cloud and content providers, service level agreements (SLA) in terms of packet loss, delay, jitter, and available bandwidth represent a key business differentiator. The goal of meeting SLA requirements drives network evolution towards architectures supporting traffic steering with more flexibility than in traditional IP-networks.

6.1 MPLS

In traditional IP routing, packets undergo FIB lookup at each hop. In a Multiprotocol Label Switching (MPLS) network, packets carrying data (i.e. IP packets) are assigned to a label-switched path (LSP) on entry to a given routing domain and subsequently forwarded via the LSP using labels at each node. For interior nodes of the network, the forwarding decision is based only on these labels [18].

ICN Path Switching, unlike MPLS, does not depend on any additional control plane protocol because the path label is constructed in the data plane. In contrast, an MPLS data plane strongly depends on timely distribution of label bindings, which can be a major issue from the point of view of network operations.

The MPLS control plane maps an end-to-end Label Switched Path (LSP) to link-local labels. Some of the most common MPLS label distribution protocols are Label Distribution Protocol (LDP) and Resource Reservation Protocol with Traffic Engineering (RSVP-TE).

[7] most likely the label added by the localhost ICN forwarder

As a tunneling technique, MPLS requires specialized ingress and egress LERs in order to organize a path in MPLS core. ICN Path Switching could be used as a tunneling mechanism, but in general it does not require specific entry and exit routers because of its built-in path discovery mechanism.

In the absence of a dedicated control plane, ICN Path Switching does not have an analogue of MPLS Fast Reroute mechanism, which relies on pre-existing knowledge of valid alternative labels in case of a forwarding failure. Fast Reroute mechanism can be implemented in ICN by extending any ICN IGP protocol to support path labels.

6.2 Segment Routing

Segment Routing (SR) is a source routing based tunneling technique that allows a host or an edge router to steer a packet through the network by using a list of segments [4]. A segment is an identifier for a topological instruction (steering the packet over a given path) or a service instruction (delivering the packet to a service).

SR can be directly applied to the MPLS architecture with no change to the forwarding plane [3]. SR can be applied to the IPv6 architecture, with a new type of routing header.

Because the information encoding the path that the packet has to traverse is included in the packet, intermediate routers do not have to maintain state for all steered paths that the network offers. Additionally, as a shortest-path segment includes all the ECMP paths to the related node, SR supports the ECMP nature of IP by design. These two features can provide substantial gains in network scalability [4].

The control plane of SR defines how the segment ID (SID) information is communicated among devices in the network. In a SR network, Node and Adjacency SIDs are advertised via the link state IGP protocol, such as ISIS and OSPF. By leveraging the rapid convergence properties of these IGPs, the segment database on each router can be quickly updated after any topology change. SR extensions of ISIS and OSPF distribute the labels and hence permit end-to-end encapsulation without requiring a specialized label distribution protocol.

Both ICN Path Steering and Segment Routing have the concept of a source route and no requirement for a specialized label distribution protocol in their designs. This arguably leading to much simpler operations and higher scalability as intermediate routers do not maintain state for all steered paths that the network offers.

6.3 Dynamic Source Routing

Dynamic Source Routing protocol (DSR) is a routing protocol designed specifically for use in multi-hop wireless ad hoc networks [10, 11]. The DSR protocol is composed of two mechanisms: Route Discovery and Route Maintenance.

Route Discovery is the mechanism by which a source node wishing to send a packet to a destination node obtains a source route to the destination. Each Route Request contains a record listing the address of each intermediate node through which this particular copy of the Route Request has been forwarded. When the initiator receives a Route Reply, it caches this route in its Route Cache for use in sending subsequent packets to this destination.

Route Maintenance is the mechanism by which a source node is able to detect, while using a source route to the destination, if the network topology has changed such that it can no longer use a particular route because a link along the route no longer works. When forwarding a packet using a DSR source route, each node forwarding the packet is responsible for confirming that data can flow over the link from that node to the nexthop.

Our scheme shares number of similarities to DSR's method of using an index into a fixed array of egress interfaces. The main difference between DSR and ICN Path Switching is that DSR is a full-fledged routing protocol that modifies the internal state of participating network nodes (i.e. Route Cache), whereas forwarding failures in ICN Path Switching do not trigger any changes in ICN FIB[8].

The roles of path discovery in ICN Path Switching and route discovery in DSR are roughly equivalent. In DSR, however, each node, upon receiving a Route Request message, rebroadcasts the packet to its neighbors, so a Route Request message is flooded throughout the network. Flooding does not happen in ICN Path Switching, because ICN path discovery only uses preexisting routes.

7 CONCLUSION

In this paper we have presented a flexible scheme to accomplish packet steering for the Content Centric or Named Data Networking ICN architectures. In order to optimally exploit the valuable multi-path and multi-destination capabilities of ICN networks, we argue that relying entirely on dynamic forwarding operating solely in the routers has important limitations. Instead we propose that explicit packet steering by consumers provides better network diagnostics and measurement (through path-steered ping and traceroute), better congestion control (through multi-path congestion control algorithms like [15]), and provides an exposed forwarding structure (the path label) that can be used for traffic engineering or SDN-like fine-grained control. Having designed the basic path label machinery, we show that it can further be utilized to accomplish path switching, which can dramatically reduce the per-packet forwarding overhead of ICN routers. We show that it achieves these gains without sacrificing performance or correctness under routing churn, and that any added vulnerabilities to attack against malicious path steering can be foiled with simple protections. The additional mechanisms needed to add Interest steering and path switching to NDN or CCN provide an attractive tradeoff between complexity and improvements in both the functional capability of the architecture and forwarding performance.

ACKNOWLEDGMENTS

The authors thank the ACM ICN reviewers and our shepherd Marc Mosko for the comments that helped to improve the paper.

REFERENCES

[1] 2017. Maximum Number of Interfaces and Subinterfaces for Cisco IOS Routers: IDB Limits. (2017).
[2] Alexander Afanasyev, Cheng Yi, Lan Wang, Beichuan Zhang, and Lixia Zhang. 2015. SNAMP: Secure namespace mapping to scale NDN forwarding. In *Computer Communications Workshops (INFOCOM WKSHPS), 2015 IEEE Conference on*. IEEE, 281–286.
[3] C. Filfils et. al. 2017. Segment Routing Architecture. (2017). https://tools.ietf.org/html/draft-ietf-spring-segment-routing-11

[8] although the fallback ICN dynamic forwarding algorithms often do have mechanisms that modify FIB entries

[4] Clarence Filsfils, Nagendra Kumar Nainar, Carlos Pignataro, Juan Camilo Cardona, and Pierre Francois. 2015. The Segment Routing Architecture. In *2015 IEEE Global Communications Conference (GLOBECOM)*. IEEE, 1–6.

[5] G. Cantor. 1878. Ein Beitrag zur Begrundung der transfiniter Mengenlehre. In *J. Reine Angew. Math.*, Vol. 84. 242–258.

[6] JJ Garcia-Luna-Aceves and Maziar Mirzazad Barijough. 2016. Content-centric networking using anonymous datagrams. In *IFIP Networking Conference (IFIP Networking) and Workshops, 2016*. IEEE, 171–179.

[7] JJ Garcia-Luna-Aceves and Maziar Mirzazad-Barijough. 2016. A Light-Weight Forwarding Plane for Content-Centric Networks. *IEEE ICNC 2016* (2016).

[8] J.J. Garcia-Luna-Aceves, Maziar Mirzazad-Barijough, and Ehsan Hemmati. 2016. Content-Centric Networking at Internet Scale Through The Integration of Name Resolution and Routing. In *Proceedings of the 3rd ACM Conference on Information-Centric Networking (ACM-ICN '16)*. ACM, New York, NY, USA, 83–92. DOI: https://doi.org/10.1145/2984356.2984359

[9] Van Jacobson, Diana K. Smetters, James D. Thornton, Michael F. Plass, Nicholas H. Briggs, and Rebecca L. Braynard. 2009. Networking Named Content. In *Proc. of CoNEXT*.

[10] David Johnson, Y Hu, and D Maltz. 2007. RFC 4728: The dynamic source routing protocol (DSR) for mobile ad hoc networks for IPv4. (2007).

[11] David B Johnson. 1994. Routing in ad hoc networks of mobile hosts. In *Mobile Computing Systems and Applications, 1994. WMCSA 1994. First Workshop on*. IEEE, 158–163.

[12] Petri Jokela, András Zahemszky, Christian Esteve Rothenberg, Somaya Arianfar, and Pekka Nikander. 2009. LIPSIN: line speed publish/subscribe inter-networking. *ACM SIGCOMM Computer Communication Review* 39, 4 (2009), 195–206.

[13] K. Kompella, G. Swallow. 2006. RFC 4379: Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures. (2006). https://tools.ietf.org/html/rfc4379

[14] L. Zhang et al. 2010. *Named Data Networking (NDN) Project.* Technical Report NDN-0001.

[15] Milad Mahdian, Somaya Arianfar, Jim Gibson, and Dave Oran. 2016. MIRCC: Multipath-aware ICN Rate-based Congestion Control. In *Proceedings of the 2016 conference on 3rd ACM Conference on Information-Centric Networking*. ACM, 1–10.

[16] Spyridon Mastorakis, Alexander Afanasyev, Ilya Moiseenko, and Lixia Zhang. 2016. *ndnSIM 2: An updated NDN simulator for NS-3.* Technical Report. Technical Report NDN-0028, NDN.

[17] Michele Papalini, Antonio Carzaniga, Koorosh Khazaei, and Alexander L Wolf. 2014. Scalable routing for tag-based information-centric networking. In *Proceedings of the 1st international conference on Information-centric networking*. ACM, 17–26.

[18] E Rosen, Arun Viswanathan, Ross Callon, and others. 2001. RFC 3031 Multiprotocol label switching architecture. (2001).

[19] Arnold L. Rosenberg. 2002. Efficient Pairing Functions - And Why You Should Care. In *Proceedings of the 16th International Parallel and Distributed Processing Symposium (IPDPS '02)*. 134–.

[20] S. Mastorakis, J. Gibson, I.Moiseenko, R. Droms, D. Oran. 2016. ICN Ping Protocol. (2016). https://tools.ietf.org/html/draft-mastorakis-icnrg-icnping-00

[21] S. Mastorakis, J. Gibson, I.Moiseenko, R. Droms, D. Oran. 2016. ICN Traceroute Protocol Specification. (2016). https://tools.ietf.org/html/draft-mastorakis-icnrg-icntraceroute-00

[22] Thomas C Schmidt, Sebastian Wolke, Nora Berg, and Matthias Wahlisch. 2016. Let's collect names: How PANINI limits FIB tables in name based routing. In *IFIP Networking Conference (IFIP Networking) and Workshops, 2016*. IEEE, 458–466.

[23] Won So, Ashok Narayanan, and David Oran. 2013. Named Data Networking on a router: fast and dos-resistant forwarding with hash tables. In *Proceedings of the ninth ACM/IEEE symposium on Architectures for networking and communications systems*. IEEE Press, 215–226.

[24] Tian Song, Haowei Yuan, Patrick Crowley, and Beichuan Zhang. 2015. Scalable name-based packet forwarding: From millions to billions. In *Proceedings of the 2nd International Conference on Information-Centric Networking*. ACM, 19–28.

[25] M. Szudzik. 2006. An elegant pairing function. (2006). http://szudzik.com/ElegantPairing.pdf

[26] Haowei Yuan and Patrick Crowley. 2015. Reliably scalable name prefix lookup. In *Architectures for Networking and Communications Systems (ANCS), 2015 ACM/IEEE Symposium on*. IEEE, 111–121.

[27] Haowei Yuan, Tian Song, and Patrick Crowley. 2012. Scalable NDN forwarding: Concepts, issues and principles. In *Computer Communications and Networks (ICCCN), 2012 21st International Conference on*. IEEE, 1–9.

[28] Zhenkai Zhu and Alexander Afanasyev. 2013. Let's chronosync: Decentralized dataset state synchronization in named data networking. In *Network Protocols (ICNP), 2013 21st IEEE International Conference on*. IEEE, 1–10.