

On the Role of Routing in Named Data Networking*

Cheng Yi[†]
University of Arizona
yic@cs.arizona.edu

Jerald Abraham
University of Arizona
jeraldabraham@cs.arizona.edu

Alexander Afanasyev
UCLA
afanasev@cs.ucla.edu

Lan Wang
University of Memphis
lanwang@memphis.edu

Beichuan Zhang
University of Arizona
bzhang@arizona.edu

Lixia Zhang
UCLA
lixia@cs.ucla.edu

ABSTRACT

A unique feature of Named Data Networking (NDN) is that its forwarding plane can detect and recover from network faults on its own, enabling each NDN router to handle network failures locally without relying on global routing convergence. This new feature prompts us to re-examine the role of routing in an NDN network: does it still need a routing protocol? If so, what impact may an intelligent forwarding plane have on the design and operation of NDN routing protocols? Through analysis and extensive simulations, we show that routing protocols remain highly beneficial in an NDN network. Routing disseminates initial topology and policy information as well as long-term changes in them, and computes the routing table to guide the forwarding process. However, because the forwarding plane is capable of detecting and recovering from failures quickly, routing no longer needs to handle short-term churns in the network. Freeing routing protocols from short-term churns can greatly improve their scalability and stability, enabling NDN to use routing protocols that were previously viewed as unsuitable for real networks.

Categories and Subject Descriptors

C.2.2 [Network Protocols]: Routing protocols

General Terms

Analysis, Performance

Keywords

NDN; routing; routing scalability; adaptive forwarding

*This work was partially supported by the National Science Foundation (No. 1039615, 1040868, 1040036) and Cisco. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the sponsors.

[†]Dr. Cheng Yi is currently with Google Inc.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ICN'14, September 24–26, 2014, Paris, France.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-3206-4/14/09 ...\$15.00.

<http://dx.doi.org/10.1145/2660129.2660140>.

1. INTRODUCTION

Named Data Networking (NDN) [13, 34] is a new network architecture that changes the basic network service semantics from “delivering packet to a given destination” to “retrieving data with a given name.” NDN communication is receiver-driven: a data consumer sends *Interest* packets carrying the names of desired data; any node in the network can return *Data* packets that have matching names to satisfy the Interests. This two-way Interest-Data packet exchange takes the same network path but in opposite directions.

Symmetric Interest-Data exchange and in-network forwarding state enable a unique feature of NDN – *adaptive forwarding* ([32, 31]). More specifically, a node expects a Data packet to come back from the same interface where it forwarded the Interest within a reasonable time period (e.g., round-trip time), otherwise it should get a timeout or receive a NACK packet [31], which signals a failure of this attempt. Upon detection of a failure, the node can then send the Interest to other interfaces to explore alternate paths. This built-in failure detection and recovery capability works on the forwarding plane, with no intervention from the control plane. Our earlier work [31] shows that NDN’s adaptive forwarding can handle link failures, prefix hijacking, and congestion control more effectively than IP networks.

Having an intelligent and adaptive forwarding plane raises new research questions. Today’s IP networks put all intelligence into routing, which disseminates topology and policy information, computes routes, detects and recovers from failures while the data plane merely forwards packets according to the FIB. When the data plane has its own adaptability, are routing protocols still needed? If so, for what purpose and to what extent? If some of routing’s tasks can be offloaded to forwarding, would that bring positive impact on routing protocols’ design and operation, e.g., making routing more scalable and stable?

In this paper we investigate the role of routing in NDN networks. Through analysis, design, and extensive simulation, we find that a routing protocol is highly beneficial in bootstrapping the forwarding plane for effective data retrieval, and in efficient probing of new links or recovered links. However, *NDN routing does not need to converge fast following network changes*, which can be handled by adaptive forwarding more promptly. This enables one to significantly improve the scalability and stability of the routing system using larger keep-alive timer values that ignore short-term failures. Furthermore, routing algorithms that would not work well in today’s IP networks may work fine in an

NDN network due to routing’s reduced role in bootstrapping adaptive forwarding.

The rest of this paper is organized as follows. Section 2 reviews NDN with a focus on the adaptive forwarding plane. Section 3 discusses the role of routing in both IP and NDN. The coordination of NDN routing and forwarding is explained in Section 4. Section 5 evaluates the performance of the coordination. Section 6 discusses other possible routing schemes for NDN. Section 7 presents related work and Section 8 concludes the paper.

2. NDN FORWARDING OVERVIEW

Each NDN packet carries a *name* field that uniquely identifies a piece of data, e.g., /ndn/papers/routing.pdf/seg1. NDN routers forward Interests based on the names, and keep forwarding state for each pending Interest. When Data packets arrive, routers use names to match them to corresponding pending Interests and forward them accordingly. Each Interest also carries a *nonce* field that can be used to detect forwarding loops. In this section we briefly review NDN’s forwarding process and how it handles link failures.

2.1 Forwarding Process

There are three key data structures in NDN’s node model, i.e., *Forwarding Information Base* (FIB), *Pending Interest Table* (PIT) and *Content Store* (CS). FIB serves as the forwarding table. It is different from the FIB in IP routers in that it is indexed by name prefixes instead of IP prefixes, and each FIB entry may provide multiple interfaces instead of a single best interface for each name prefix. Unlike FIB, PIT and CS are unique to NDN. Both PIT and CS are indexed by names. A PIT entry records incoming and outgoing interface(s) of an Interest, and is used to guide Data forwarding. CS is a temporary cache of Data packets that can speed up the satisfaction of Interests.

In NDN, the forwarding process works as follows. When a router receives an Interest, it first checks the Interest name against the CS and returns the Data if there is a match. Otherwise, the router checks the Interest name against the PIT. If a PIT entry already exists, i.e., the Interest has already been forwarded but no Data has been returned yet, the router simply adds the incoming interface of the Interest to the PIT entry. If no PIT entry exists, the router adds a new PIT entry and further looks up the Interest name in the FIB using longest prefix match. If a matching FIB entry is found, the Interest is forwarded by a *forwarding strategy* module [13]. Otherwise, the router cannot satisfy the Interest and may send a NACK back to the incoming interface of the Interest [31]. When a router receives a Data packet, it checks the Data name against the PIT. If a PIT entry is found, the Data is stored in the CS and further forwarded to the incoming interfaces of the corresponding Interests, which have been recorded in the PIT. Otherwise, the Data is dropped since it is either unrequested or no longer wanted.

The forwarding strategy associated with the name space of an Interest determines whether and how to forward the Interest. It may take such information as ranking from the routing protocol, interface status, round-trip time (RTT) and congestion level into consideration. In this paper we adopt the forwarding strategy proposed in [31]. For each name prefix, each interface is assigned a color code depending on its current working status. It is *Green* for a working interface, *Red* if the interface is not working, and *Yellow*




Figure 1: A simple network example.

if the status is uncertain. The forwarding strategy always prefers Green interfaces over Yellow ones, and never uses Red interfaces to forward Interests.

2.2 Failure Recovery

NDN’s two-way symmetric traffic flow enables fast fault detection. Routers can calculate RTT for each Interest-Data exchange, which can be used as a prediction for future Interests. After forwarding an Interest, a router starts a timer based on the average of previous RTTs; potential network problems can be detected if no Data is received before the timer expires. With Interest NACKs [31], fault detection and notification is even faster. When network problems are detected, routers can explore alternative paths freely without worrying about loops, since loops can be detected by checking the nonce field carried in Interests. Fast fault detection and loop-free forwarding are the two unique features that make NDN’s forwarding plane smart and adaptive – routers are able to handle network faults such as prefix hijacking, failures and congestion locally at the forwarding plane [31].

We use the simple example in Figure 1 to illustrate how NDN routers handle link failures. The cost of the links are marked in the figure; routers rank the interfaces using the cost of their best paths towards the destination. When there is no failure in the network, *A* uses *B* as its primary next hop for content provided by *D*. Interface *A-B* will be marked Green as long as Data continues to flow from *B* to *A*. When link *B-D* fails, *A* will keep sending Interests to *B* at first. However, *B* cannot satisfy the Interests due to the failure, so it will send NACKs back to *A*. Upon receiving a NACK, *A* will mark *A-B* Yellow and retry the next best interface, in this case *A-C*. Since there is no failure on this path, Data will flow back through path *D-C-A*. *A* will then mark interface *A-C* Green and start using *C* as the primary next hop.

3. ROLE OF ROUTING

Since NDN’s forwarding model is a strict superset of the IP model, any routing scheme that works well for IP should also work well for NDN [13]. However, today’s IP routing protocols suffer from issues such as slow convergence or poor scalability. On the other hand, NDN has a smart and powerful forwarding plane, which is able to take over part of routing’s responsibility in IP. In this section, we first review IP routing, and then rethink the role of routing in NDN.

3.1 Routing in IP

IP’s routing plane is intelligent and adaptive, but its forwarding plane is stateless and strictly follows routing. Therefore the routing plane is also regarded as the control plane. Routing is responsible for building the routing table and

maintaining it in face of network changes, including both long-term topology and policy changes as well as short-term churns. When there is a change in the network, routers need to exchange routing updates with each other in order to reach new global consistency. The time period after a change happens and before all routers agree on the new routing state is called the *routing convergence period*. IP routing protocols need to converge fast in order to reduce packet loss and resume packet delivery after network changes.

However, fast routing convergence is challenging in large operational networks. The fundamental reason is that it conflicts with other design goals for routing protocols, i.e., routing stability and scalability. Routing stability ensures stable routing paths within the network. It is important for applications that suffer from RTT fluctuation; it also helps routers achieve traffic engineering goals. Routing scalability is essential for supporting a large number of nodes, links and prefixes¹ in the network. For link-state routing, each router knows the entire topology. These protocols can converge fast, but at the cost of poor stability and limited scalability. For distance/path-vector routing, routers do not have a full knowledge of the topology. They are able to achieve better scalability, but the convergence time may be as long as tens of minutes. Below we use link-state routing as an example to explain the issues with today’s IP routing protocols.

The routing convergence period can be divided into four phases, i.e., *failure detection*, *update propagation*, *route computation* and *FIB update*. In link-state routing, routers periodically exchange HELLO messages to maintain connection: if no HELLO message is received within the DEAD interval, the link is considered down. Previous research ([7, 11]) recommended setting the HELLO interval to be on the order of milliseconds in order to detect failures quickly. However, this not only increases overhead but also affects routing stability, since a temporarily congested link may be mistakenly considered fluctuating down and up. After a link failure is detected, attached routers need to generate routing updates and propagate them to the rest of the network; when a router receives a routing update, it needs to recompute the routing table. To achieve fast routing convergence, all these steps should be done as quickly as possible. However, if the network is unstable (e.g., there is a flapping link), generating routing updates and recomputing routing table frequently will increase bandwidth and computation overhead as well as harm routing stability. At the same time, shortest path first (SPF) computation time increases with the size of the network; FIB update time depends on the number of prefixes. To achieve fast convergence, both the network size and the number of prefixes need to be limited, leading to poor scalability.

There are mechanisms to improve link-state routing stability and scalability. Dynamic timers improve routing stability by limiting the rate of update generation and SPF computation. However, these timers are increased exponentially each time, potentially increasing convergence time significantly when the network is unstable. Therefore, short initial timers have been suggested [11]. *Area* was introduced to improve routing scalability [22]. However, it leads to sub-optimal paths between areas and increases the complexity of configuration. Although inter-area routing can utilize

¹Supporting large number of prefixes is particularly important in NDN since the number of name prefixes will be orders of magnitude larger than the number of IP prefixes in today’s Internet.

distance-vector or path-vector routing algorithms that may scale better, they converge much slower.

In summary, it is hard to achieve fast convergence, stability and scalability simultaneously in a routing protocol. If failures can be handled without global routing convergence, the requirement on fast convergence can be relaxed, making it possible to improve routing stability and scalability.

3.2 Routing in NDN

In NDN, the forwarding plane is the actual control plane since the forwarding strategy module makes forwarding decisions on its own. This fundamental change prompts us to rethink the role of routing in NDN. The first question is whether NDN still needs routing protocols. Conventionally, routing protocols are responsible for disseminating topology and policy information, computing routes and handling short-term network changes. For NDN to work without routing, routers need to be able to do the following things efficiently: 1) retrieve Data when the network is stable; 2) handle link failures; and 3) handle link recovery. Can NDN achieve these solely with the forwarding plane?

Another question that arises is if NDN does need routing protocols, how will they be different from today’s existing routing protocols? With the intelligent and adaptive forwarding plane, can some of the routing plane’s functionality be offloaded to the forwarding plane, and which? In addition, how will the design and operation of routing protocols benefit from this shift of functionality? In the next section we try to give answers to these questions.

4. ROUTING AND FORWARDING COORDINATION

In this section, we seek answers to the questions raised in 3.2. Previous research [31] shows that NDN routers are able to handle link failures effectively without routing. In this section we focus on whether NDN routers can retrieve Data and react to link recovery *efficiently* without routing. We show that NDN does need routing protocols to help bootstrap the forwarding process and handle link recovery. In addition, we specify how the routing plane coordinates with the forwarding plane, and present a simple method to improve routing stability and scalability in NDN.

4.1 Interface Ranking

The forwarding plane design presented in [31] assumes interfaces are ranked by routing preference. Can NDN routers retrieve Data efficiently without routing to rank the list of available interfaces? The answer is negative. In the extreme case, we can implement a forwarding strategy that floods every Interest to all available interfaces. This way we can always retrieve Data quickly through the best paths. However, it will also incur significantly high overhead. We can also implement forwarding strategies that randomly explore the interfaces or try them one-by-one in a round-robin fashion. Given enough time, routers should be able to find working paths since all possible paths will be explored. One big issue with this method is that path exploration may take extremely long time as shown in Section 5.3.

Consequently, NDN routers need good interface ranking to help bootstrap the forwarding process. The responsibility of providing interface ranking lies in the routing protocols. Existing routing algorithms such as link-state or

Pseudocode 1 ProbingDue Algorithm

```
1: function PROBINGDUE(FibEntry, Intf)
2:   if Intf  $\neq$  FibEntry.RoutingPreferredIntf then
3:     if FibEntry.LastProbingTime + M  $\leq$  Now() or
4:       FibEntry.PacketsSinceLastProbing  $\geq$  N then
5:       Return True
6:     end if
7:   end if
8:   Return False
9: end function
```

distance/path-vector routing can be used to rank the interfaces². The details are explained as follows.

4.1.1 Link-State Routing

Link-state routing protocols store the entire network topology in the link-state database (LSDB), making it possible to compute optimal interface ranking. Suppose a node N has n interfaces $I_1 \dots I_n$. For Data provided by node M , we rank these interfaces using $C_{N,k}^M$, which is the cost of the best path from N to M through interface I_k . One simple method to compute $C_{N,k}^M$ for all destinations through I_k is to remove all interfaces except I_k from N 's LSDB, and run *Dijkstra's algorithm* to compute the shortest paths. This may not be the best method since it will end up calling Dijkstra's algorithm once for every interface. It is just used to illustrate how interface ranking can be done in link-state routing. Optimization of the algorithm is possible but out of the scope of this paper.

4.1.2 Distance/Path-Vector Routing

In distance-vector or path-vector routing, routers announce cost of the complete routing path towards each destination to their neighbors. When router N receives a routing announcement for Data provided by M from interface I_k , it simply adds the link cost of I_k to the received path cost to obtain its path cost $C_{N,k}^M$. The interfaces are then ranked by the path costs to M through them.

Note that a router may not receive routing announcement from all interfaces, since these routing protocols often incorporate split-horizon route announcement to prevent routing loops. If router N learns a route towards M through interface I_k , it will not advertise its route to M over I_k . Interfaces that do not receive routing announcement are assigned infinite cost to ensure they stay at the end of the ranked interface list. They will only be used as the last resort if all higher-ranked interfaces fail to retrieve Data.

These interfaces are useful in many situations. For example, in BGP if a provider P uses a customer C as the next hop, it will not make routing announcement to C . If C 's best path fails, it will not have an alternative path until routing converges, in which case P will announce its alternative path to C . RBGP [14] is proposed to address this issue by allowing P to announce its alternative path to C even without failures. NDN, on the other hand, is able to achieve the same effect without changing the routing protocol.

4.2 Probing

²The case of BGP is more complex because it also takes routing policy into consideration. How to accommodate routing policy in interface ranking is part of our future work.

Pseudocode 2 Probing Algorithm

```
1: function PROBE(Interest, FibEntry, PitEntry)
2:   interface  $\leftarrow$  FibEntry.RoutingPreferredIntf
3:   if interface  $\notin$  PitEntry.Outgoing and
4:     interface  $\notin$  PitEntry.Incoming then
5:     if interface.Available then
6:       Interest.Nonce  $\leftarrow$  GenerateNonce()
7:       Transmit(interface, Interest)
8:       Add interface to PitEntry.Outgoing
9:       FibEntry.LastProbingTime  $\leftarrow$  Now()
10:      FibEntry.PacketsSinceLastProbing  $\leftarrow$  0
11:    end if
12:  end if
13: end function
```

It has been shown that NDN routers can handle link failures locally at the forwarding plane [31]. In this subsection we answer the question of whether the same applies to link recovery. Routers can detect link failures quickly by observing Interest-Data exchanges or Interest NACK. However, there is no explicit signal for link recovery from the forwarding plane. Again let's take Figure 1 as an example. After interface $B-D$ recovers from a failure, interface $A-B$ becomes the best interface for A to retrieve data from D . However, A will continue using interface $A-C$ because the forwarding strategy prefers Green interfaces over Yellow ones. In this case, A needs to probe interface $A-B$ by sending a copy of an Interest to it. If the probing Interest successfully brings Data back, interface $A-B$ will be marked Green and be used to forward subsequent Interests to D .

There is a research question of when to perform probing. An Interest copy is used for probing so that regular Data retrieval will not be affected if probing is unsuccessful. However, this causes extra Interest and Data in the network. There is a trade-off between how fast a link recovery is detected and the amount of overhead caused by probing. In CCNx [2], a prototype implementation of NDN, routers probe alternative interfaces periodically in order to detect better paths. This enables routers to detect link recovery at the forwarding plane. Fast recovery detection is achievable through aggressive probing. However, it will incur significant overhead.

In fact, routing is able to help with the dilemma. If there is a routing protocol, it will be able to detect link recovery and converge to it by ranking the new best interface (Yellow) higher than the currently used one (Green). Thus we can take advantage of routing by only probing a Yellow interface if its ranking is higher than the Green interface(s). This way we can keep the probing overhead low, and switch back to the optimal paths as soon as routing converges. Routing convergence time is not a concern because the alternative paths found by the forwarding plane are of good quality [31]. Note that probing is also useful in failure handling if the alternative paths found by the forwarding strategy are not the optimal ones.

We propose a probing algorithm as presented in Pseudocode 1 and 2. After forwarding each Interest, the strategy module calls **ProbingDue** to check whether probing is needed. Two thresholds are introduced to further limit the probing overhead. For each FIB entry, M is the minimum time interval, and N is the minimum number of packets forwarded between two consecutive probings. The algorithm

Table 1: Topologies Used in the Simulations.

Topology	Before Processing		After Processing	
	Node #	Link #	Node #	Link #
Abilene	12	30	11	28
AS1239-PoP	52	168	32	128
AS701-PoP	83	438	47	366
AS1239-Router	284	1882	N/A	N/A

returns true only if at least M time has elapsed or at least N packets have been forwarded since the last probing. The setting of M and N depends on the traffic load and the probing overhead network operators are willing to tolerate. The probing algorithm (Pseudocode 2) sends a copy of the Interest to the routing preferred interface using a different nonce. The nonce is changed so that routers will not drop the probing Interest after seeing the original Interest.

4.3 Routing Stability and Scalability

Link-state routing protocols exhibit poor stability and scalability in IP due to the fast routing convergence requirement. However, there is a simple method to address these issues in NDN. Since NDN routers can handle network failures at the forwarding plane, short-lived failures can be masked from the routing protocols. Research shows that the duration of network failures follows a long-tailed distribution, and over 50% of failures last less than one minute ([20, 28]). Therefore, the number of routing events can be significantly reduced if routing protocols do not need to react to the short-lived failures. As a result, the bandwidth and CPU cycles consumed by routing updates can be reduced, and there will be less routing fluctuation. In addition, since there is no fast routing convergence requirement, larger networks and more name prefixes become affordable. In summary, both routing stability and scalability can be significantly improved.

For link-state routing, we can implement the idea by increasing the HELLO and DEAD interval. For example, if we set the DEAD interval to be one minute, over 50% of the link failures will be ignored by the routing protocol. Alternatively, we can increase the suppression timer for routing update generation and SPF computation to achieve the same effect. Although this idea looks simple, it can be applied to any existing IP routing protocol to improve its stability and scalability. We will evaluate the effectiveness of this method in the next section.

5. EVALUATION

In this section we use extensive simulations to show that NDN’s packet forwarding performance under network failures is hardly affected by routing convergence time; by masking short-lived failures from routing protocols, one can effectively reduce routing overhead while maintaining high packet delivery performance in NDN networks.

5.1 Simulation Setup

Unless otherwise specified, we run experiments in the QualNet simulator [4] which provides complete implementations of OSPF and RIP routing protocols. We implement basic NDN operations and the forwarding strategy presented in [31] in the simulator. We also make necessary changes to the routing protocols as described in Section 4.1.

We use the Abilene topology [1] and selected Rocketfuel topologies [26] in the experiments. A summary of the topologies is presented in Table 1. We process the first three topologies to remove all single-homed nodes, because if the link of a single-homed node fails, the node will be disconnected from the network and thus cannot provide any insightful result. For OSPF, we use propagation delay as the cost of the links. Unless otherwise specified, we report results from the AS1239-PoP topology. Results for other topologies are similar and lead to the same conclusions. The AS1239-Router topology is only used to show the improvement of routing scalability³.

We inject random link failures into the topologies. A shifted Pareto distribution is used to generate time-to-failure and time-to-recover values for each link independently. We use 120 seconds as the mean-time-to-recover, and 1000 seconds as the mean-time-to-fail; the scale parameter of the Pareto distribution is set to be 208 so that 50% of the failures last less than one minute [18, 20]. When a link fails, both directions of the link stop working. With this model, multiple network events (failures and recovery) can happen concurrently.

5.2 Failure Handling

In this set of experiments, we compare the packet delivery performance of NDN and IP in failure scenarios under different settings. We also evaluate the forwarding overhead of NDN when a prefix becomes unreachable due to failures.

5.2.1 Impact of Routing Convergence Time

In this experiment, we run OSPF as the routing protocol and study the impact of HELLO interval on packet delivery performance. We inject random link failures into the network as described in Section 5.1. In order to measure packet delivery performance in NDN and IP, we run simple applications among all pairs of nodes in the network. For NDN, each node announces a distinct name prefix and provides content under this prefix. Each node also acts as a consumer requesting data from all other nodes. A consumer sends one Interest towards each name prefix every second. If Data is not received, a consumer will retransmit the Interest every second up to twice. Different consumers request different pieces of Data from the same name prefix so that they do not affect each other. Caching is also disabled so that we can focus on routing and forwarding behaviors⁴. For IP, each node acts as both client and server. Each client sends one UDP request to each server every second⁵. The server responds with UDP packet carrying the content. Similar to NDN consumers, these clients also retransmit requests if replies are not received. The sizes of the UDP packets are the same as those in NDN.

Figure 2 and 3 present the packet delivery rate for each node pair in IP and NDN under different HELLO interval settings. Figure 2 shows that HELLO interval has a huge impact on the packet delivery performance in IP. The shorter HELLO interval, the faster packet delivery can be resumed.

³We do not run packet-level simulations on the topology due to performance limitations of the simulator.

⁴If consumers request the same content and caching is enabled, NDN would perform even better.

⁵The packet rate is much lower than real Internet traffic due to performance limitation of the simulator. In fact, the IP packet delivery performance will be worse if the packet rate is higher.




Figure 2: Packet delivery performance in IP under different HELLO interval.




Figure 3: Packet delivery performance in NDN under different HELLO interval.




Figure 4: CDF of packet loss rate under different routing protocols.

The median packet delivery rate of IP is 99%, 91% and 72% when the HELLO interval is 1S, 10S and 60S respectively. Figure 2 also shows that NDN with 60S HELLO interval even works slightly better than IP with 1S HELLO interval.

Figure 3 shows the impact of HELLO interval on the packet delivery rate of NDN. When the HELLO interval increases from 1S to 10S, the performance degradation is negligible. When the HELLO interval increases from 10S to 60S, the packet delivery rate decreases slightly. This is because only two consumer retransmissions are allowed. The packet delivery performance can be further improved by allowing more consumer retransmissions. Overall, the HELLO interval has little impact on the packet delivery performance in NDN.

We also evaluate the packet delivery performance under different routing protocols. Figure 4 shows the CDF of packet loss rate of NDN and IP when OSPF and RIP are used. Although RIP is generally considered to have poor routing convergence properties, it performs quite well with NDN. NDN with RIP performs much better than IP with OSPF or RIP. The performance difference between OSPF and RIP in NDN is due to the difference in interface ranking. Recall that RIP may not provide cost for all interfaces, thus OSPF is able to provide better interface ranking.

5.2.2 Comparison with IPFRR

In the previous section we evaluate the packet delivery performance of plain IP, which relies on routing to handle network failures. However, IP networks may adopt solutions that handle network failures without routing convergence, e.g., IPFRR. In this experiment, we compare NDN against Loop-Free Alternate (LFA) [8], the only commercially available IPFRR solution. We implement LFA in a custom simulator, and repeat the link failure experiment in [31] without routing convergence. Only two consumer retransmissions are allowed for NDN. In each run of the experiment, we associate each link with a probability of failure, and randomly generate link failures. We run the experiments 1000 times for each link failure probability.

Figure 5(a) shows the average reachability of NDN and LFA with 95% confidence interval under different failure probability. We only consider the situations where the source and destination are not physically disconnected by the failures. The figure shows that NDN is always able to recover

from much more failure scenarios than LFA. Figure 5(b) shows the CDF of stretch of alternative paths found by NDN and LFA. The 98-percentile of path stretch for NDN and LFA is 1.06 and 1.13 respectively. In conclusion, NDN is able to cover more failure scenarios and find better alternative paths than LFA.


5.2.3 Prefix Unreachable

Previous experiments show that NDN performs well in handling link failures. When a node fails, however, the name prefix served by the node may become unreachable. In such cases, path exploration may lead to extra Interests all over the network. In this experiment we evaluate NDN’s exploration overhead when a name prefix becomes unreachable. In each run of the experiment we fail one node and let all other nodes request content from this node before routing convergence⁶. Both NDN and IP applications will retransmit the same request twice. For each flow, we count the number of hops that each packet traverses in both NDN and IP, and compute the hop count ratio of NDN over IP. We run the experiment for every node failure scenario and present the CDF of the ratio in Figure 6.


In IP, retransmitted requests will be sent to the same paths, whereas in NDN, retransmitted Interests may trigger path exploration, leading to large overhead. Surprisingly, NDN incurs less overhead than IP in 26% of the cases. This is because retransmitted Interests do not always trigger path exploration in NDN. If a node has already explored all its interfaces, a further retransmission will only get a NACK back to the application without being further forwarded. In contrast, IP routers will always forward the packets before routing convergence. The ratio is smaller than 5 in 93% of the cases. Only in some rare cases does NDN cause excessively high exploration overhead.

The exploration overhead becomes significant when popular content becomes unreachable, as many consumers will be requesting the content and their Interests will trigger many attempts by routers to find working paths. But on the other hand, popular content is usually hosted and served by multiple servers placed at different locations. In addition, popular content is more likely to be cached by routers. Thus its

⁶After routing converges, routers will learn about the failure and stop forwarding the requests.



(a) Reachability.



(b) CDF of path stretch.

Figure 5: Comparison between NDN and IPFR.




Figure 6: CDF of hop count ratio of NDN over IP when prefix is unreachable.




Figure 7: CDF of time to find working paths with and without routing.




Figure 8: CDF of number of hops that probing Interests and Data traverse.

chance of becoming unreachable is slim. The overall impact in large scale networks needs further investigation.

5.3 Forwarding without Routing

In this experiment we show how NDN forwarding performs without routing. Since routers have no idea how to rank the interfaces without input from routing, we implement a forwarding strategy that prefers Green interfaces over Yellow ones, and randomly picks a Yellow interface if no Green interface exists. All interfaces are initialized to be Yellow. If Data is brought back from an interface, the interface will be marked Green and used to forward subsequent Interests.

In each experiment run, we pick one node as the consumer and another as the content provider. Assuming the consumer keeps retransmitting Interests until Data is received, we measure how long it takes to receive the data. We enumerate all combinations of consumers and providers and draw the CDF in Figure 7. In 89% of the cases, the consumer retrieves Data within one second. However, it can take up to 40 seconds to find a working path in some rare cases. The situation can get worse as the network becomes larger. In contrast, Data retrieval always follows the best paths when routing protocol can provide interface ranking. Therefore, although NDN has a powerful forwarding plane that is able to handle link failures on its own with only local information, the interface ranking provided by a routing protocol can make the local search more effective.

5.4 Routing and Forwarding Coordination

In this set of experiments we evaluate how NDN's routing and forwarding plane benefit from each other.

5.4.1 Probing Overhead

With the help of routing protocols, routers only need to perform probing when a better link is presented by routing. We evaluate probing overhead in this experiment. In each run of the experiment, we fail one link and run applications to let routers find working paths. Then we bring the link back up again, and run applications after routing convergence to measure the number of hops that probing Interests and Data traverse. Interest NACKs are counted as probing Interests. Applications are only run between node pairs whose traffic is affected by the failure. We run the experiment on all link failure scenarios and report the CDF in Figure 8. In 36% of the cases, probing Interests and Data only traverse 2 hops; they traverse no more than 6 hops in 94% of the cases. Probing Interests traverse more hops than Data in some rare cases, because a probing Interest does not necessarily bring Data back, and some of them may loop back to previously visited nodes and trigger NACKs. This experiment shows that by taking advantage of routing, probing only incurs very small overhead.

5.4.2 Routing Overhead

In this experiment, we evaluate the routing overhead of OSPF under different HELLO and DEAD interval settings. Specifically, we measure the number of HELLO messages, link-state (LS) updates and SPF computations for each node. HELLO and LS update messages constitute the majority of routing messages triggered by failures and recovery. We set the HELLO interval to be 1S, 10S and 60S; the DEAD interval is always four times the HELLO interval. Random




Figure 9: Routing overhead under different HELLO intervals in AS1239 PoP-level topology.




Figure 10: Routing overhead under different HELLO intervals in AS1239 router-level topology.

link failures are injected into the network as described in Section 5.1, and each experiment is run for 3000 seconds. Only LS updates and SPF computations triggered by failures and recovery are counted⁷. The numbers obtained in this experiment are the same for both NDN and IP.

Figure 9(a) shows the number of HELLO messages sent by each node under different HELLO interval in AS1239-PoP topology. As the HELLO interval increases from 1 second to 60 seconds, the number of HELLO messages sent by each node is decreased by 98% as one would expect. Figure 9(b) and 9(c) present the number of triggered LS updates and SPF computations for each node. As the HELLO intervals increase, less failure events will be detected by OSPF. No routing update will be generated and propagated for the undetected failures, and thus no SPF computation will be performed. If we increase the HELLO interval from 1 second to 60 seconds, the number of LS updates is decreased by 52% to 80%, and the number of SPF computations is decreased by 77% to 82%. Therefore, we can effectively reduce the overhead caused by HELLO messages, LS updates and SPF computation by increasing the HELLO interval.

We run the same experiments in AS1239 router-level topology to illustrate how the method works in large ISP networks. The CDF of number of triggered LS updates and SPF computations are presented in Figure 10. The median numbers of LS updates and SPF computations are decreased by 87% and 90% when HELLO interval increases from 1 sec-

ond to 60 seconds. In conclusion, routing overhead can be significantly reduced by masking short-lived failures from the routing protocol. Since less LS updates are generated and propagated and less SPF computations are performed, routing becomes more stable and scalable.

6. DISCUSSION

Routing is a necessary subsystem for any large scale network. Like IP, NDN itself does not dictate what kinds of routing algorithms or protocols to use. However one can take advantage of NDN’s adaptive forwarding plane to improve the stability and scalability of existing routing protocols, as well as enable routing protocols that are deemed difficult to adopt in IP networks.

Traditional Routing Protocols: With adaptive forwarding, routing in NDN only assumes a supporting role. It provides a reasonable starting point for forwarding which can then effectively explore different choices. The job of routing becomes more of disseminating topology and policy information than distributed computation of best paths. This new division of labor between routing and forwarding makes routing protocols simpler and more scalable. Traditional routing protocols such as OSPF, RIP, and BGP can benefit greatly from NDN’s adaptive forwarding plane. They can be tuned for synchronizing among routers long-term topology and policy information without handling short-term churns.

Centralized Routing: Routing protocols have been designed to operate in a distributed manner to avoid single point of failure. With the increasing complexity in network management, however, Software-Defined Networking (SDN) has emerged to enable centralized management and control of

⁷Notice that OSPF also floods refresh link-state announcements periodically even in the absence of network event. These refresh updates are not counted since they are not affected by routing convergence behaviors.

networks, including logically centralized routing scheme. It is much easier to change the routing configurations on a central controller than on all participating routers, and to implement sophisticated traffic engineering schemes at the controller than on individual routers. Routing overhead can also be greatly reduced, since routing updates only need to be sent to the controller instead of being flooded to the entire network, and only the controller needs to perform SPF computations. However, one of the biggest concerns about centralized routing is the potentially prolonged convergence delay, which includes failure detection at local router, report to the controller, route recompilation at the controller, and dissemination of new routes to individual routers. NDN’s adaptive forwarding removes the demands on convergence delay, making centralized routing feasible.

Coordinate-based Routing: In coordinate-based routing, instead of disseminate the network topology to routers, the coordinates of nodes are disseminated. The main characteristics of the network topology are embedded in the coordinates. Routers do greedy routing based on coordinates, i.e., forward packets to the neighbor whose distance (computed using coordinates) to the destination is the shortest among all neighbors. One example of such routing scheme is hyperbolic routing [23]. The advantages of this routing scheme include smaller routing tables (i.e., only need to know the destination’s coordinates and neighbor routers’ coordinates) and minimal routing updates (i.e., link failures and recovery do not affect a node’s coordinates). However, in IP networks, this routing scheme is not guaranteed to be able to deliver packets. It is possible that the forwarding process runs into a local minimal, where all neighbors are farther to the destination than the current router. Path stretch may also get large. NDN’s adaptive forwarding can fix these problems and make this routing scheme a possibility.

7. RELATED WORK

A massive amount of research has been conducted on how to gracefully accommodate routing changes with minimum impact on packet delivery in IP networks. One category of solutions rely on routing protocols to adapt to the changes. Francois et al. show that sub-second link-state routing convergence in large intra-domain networks is achievable by tuning various timers [11]. But this method incurs extra routing overhead and may also cause routing instability.

Fast reroute (FRR) mechanisms handle link failures by pre-computing alternative paths. MPLS FRR mechanisms provide backup paths in MPLS-enabled networks to protect specific links from failures [3]. Similarly, IPFRR mechanisms (e.g., [8]) provide temporary alternative paths before routing convergence in pure IP networks. However, it is hard for the FRR mechanisms to cover all possible failure scenarios; nor can they handle multiple link failures well.

Another category of solutions handle network failures via multipath forwarding. Path splicing [21] is an end-to-end multipath solution that provides link recovery controlled by end hosts. Each router provides multiple routing tables and let end hosts specify which one to use at each router. Path splicing may take a long time to find alternative paths, and sometimes may not be able to find them even if they exist [31]. MRC [16] provides multiple routing configurations to handle network failures. Different from path splicing, MRC lets routers switch configurations when failures are

detected. However, it may not handle multiple concurrent failures well due to the limited path choices.

There are also solutions that carry routing or forwarding information in the packet headers. Failure carrying packets (FCP) [17] puts failure information into the packet headers, and let routers recompute the routing tables on-the-fly upon receipt of FCP. However, the method increases computation overhead, and the sizes of FCP headers may become arbitrarily large. Liu et al. propose Data-Driven Connectivity (DDC) [19] to ensure forwarding connectivity at the data plane. DDC organizes the network as a destination-oriented directed acyclic graph (DAG) to avoid loops, and uses two bits in the packet header to notify link reversal. DDC has its own control plane algorithm, therefore cannot make use of existing routing protocols.

NDN keeps more states and does more processing at the forwarding plane than IP. However, these forwarding states also bring many benefits, such as native support of synchronous and asynchronous multicast, loop-free multipath data retrieval, efficient recovery from packet loss, flow balance and congestion control, which makes the forwarding plane more robust and efficient. The purpose of this paper is to assess how routing protocols can benefit from such a forwarding plane assuming it’s already in place. There are a number of other work with promising results on how to build such an NDN forwarding plane that can operate at very fast speed [33, 25, 30]. On the other hand, the management and stability of the forwarding state on an Internet scale still need further improvement as argued in [29].

A considerable amount of research has been conducted on routing and forwarding in the context of NDN and ICN in general. Hoque et al. proposed NLSR [12], a link-state NDN routing protocol that runs on top of NDN. It is the first distributed routing protocol for NDN. INFORM [10] is a dynamic Interest forwarding mechanism based on Q-routing. It is able to discover cached Data copies in the network that are not announced through routing protocols. Tortelli et al. proposed COBRA [27], a bloom-filter based intra-domain routing algorithm for NDN. It is simple and efficient as no routing message is required between NDN nodes. Saino et al. applied cache-aware hash routing techniques to ICN and showed that inter-domain traffic can be reduced significantly with hash routing [24]. Carzaniga et al. investigated multi-tree routing in ICN [9]. Their proposed routing scheme supports both content delivery and event notification.

Routing scalability is also a critical issue for NDN to operate at Internet scale. Kutscher et al. discussed the routing scalability issue for Information Centric Networking (ICN) in general [15]. Afanasyev et al. [5] investigated the routing scalability issue specifically for NDN and proposed a solution based on map-n-encap. α Route [6] is a novel name-based routing scheme for ICN. It utilizes distributed hash table to achieve scalable routing table size.

8. CONCLUSION

In this paper we study the role of routing in NDN. NDN’s adaptive forwarding plane leads to a new division of labor between routing and forwarding planes. While the latter can detect and recover from link failures quickly independent from the former, the former helps bootstrap adaptive forwarding and handle link recovery. We specify how NDN routing coordinates with forwarding through interface rank-

ing and probing mechanisms. Our analysis and simulations show that NDN routing protocols can benefit from the forwarding plane due to the relaxed requirement on timely detection of failures and convergence delay. Consequently NDN routing stability and scalability can be greatly improved. Moreover, the adaptive forwarding plane also enables new routing schemes that may not work well in IP to be used in an NDN network.

9. REFERENCES

- [1] Abilene TM. <http://www.cs.utexas.edu/~yzhang/research/AbileneTM/>.
- [2] CCNx. <http://www.ccnx.org/>.
- [3] MPLS Traffic Engineering Fast Reroute – Link Protection. http://www.cisco.com/en/US/docs/ios/12_0st/12_0st10/feature/guide/fastrout.html.
- [4] QualNet. <http://web.scalable-networks.com/content/qualnet/>.
- [5] A. Afanasyev, C. Yi, L. Wang, B. Zhang, and L. Zhang. Scaling ndn routing: Old tale, new design. Technical Report NDN-0004, NDN, July 2013.
- [6] R. Ahmed, M. Bari, S. Chowdhury, M. Rabbani, R. Boutaba, and B. Mathieu. α Route: A name based routing scheme for Information Centric Networks. In *Proceedings of IEEE INFOCOM*, 2013.
- [7] C. Alaettinoglu, V. Jacobson, and H. Yu. Towards Milli-Second IGP Convergence. Internet Draft draft-alaettinoglu-isis-convergence-00.txt, Nov. 2000.
- [8] A. Atlas and A. Zinin. RFC 5286: Basic Specification for IP Fast Reroute: Loop-Free Alternates, 2008.
- [9] A. Carzaniga, K. Khazaei, M. Papalini, and A. L. Wolf. Is information-centric multi-tree routing feasible? In *Proceedings of ACM SIGCOMM ICN Workshop*, 2013.
- [10] R. Chiochetti, D. Perino, G. Carofiglio, D. Rossi, and G. Rossini. Inform: A dynamic interest forwarding mechanism for information centric networking. In *Proceedings of ACM SIGCOMM ICN Workshop*, 2013.
- [11] P. Francois, C. Filsfils, J. Evans, and O. Bonaventure. Achieving Sub-Second IGP Convergence in Large IP Networks. *ACM SIGCOMM CCR*, 35(3), July 2005.
- [12] A. K. M. M. Hoque, S. O. Amin, A. Alyyan, B. Zhang, L. Zhang, and L. Wang. Nlsr: Named-data link state routing protocol. In *Proceedings of ACM SIGCOMM ICN Workshop*, 2013.
- [13] V. Jacobson, D. K. Smetters, J. D. Thornton, M. F. Plass, N. H. Briggs, and R. L. Braynard. Networking Named Content. In *Proceedings of ACM CoNEXT*, 2009.
- [14] N. Kushman, S. Kandula, D. Katabi, and B. Maggs. R-BGP: Staying Connected in a Connected World. In *Proceedings of USENIX NSDI*, 2007.
- [15] D. Kutscher, S. Eum, K. Pentikousis, I. Psaras, D. Corujo, D. Saucez, T. C. Schmidt, and M. Wählisch. ICN Research Challenges. Internet draft, 2014.
- [16] A. Kvalbein, A. Hansen, T. Cicic, S. Gjessing, and O. Lysne. Fast IP Network Recovery Using Multiple Routing Configurations. In *Proceedings of INFOCOM*, 2006.
- [17] K. Lakshminarayanan, M. Caesar, M. Rangan, T. Anderson, S. Shenker, and I. Stoica. Achieving Convergence-Free Routing using Failure-Carrying Packets. In *Proceedings of ACM SIGCOMM*, 2007.
- [18] S. Lee, Y. Yu, S. Nelakuditi, Z. li Zhang, and C. nee Chuah. Proactive vs Reactive Approaches to Failure Resilient Routing. In *Proceedings of IEEE INFOCOM*, 2004.
- [19] J. Liu, A. Panda, A. Singla, B. Godfrey, M. Schapira, and S. Shenker. Ensuring Connectivity via Data Plane Mechanisms. In *Proceedings of USENIX NSDI*, 2013.
- [20] A. Markopoulou, G. Iannaccone, S. Bhattacharyya, C.-N. Chuah, Y. Ganjali, and C. Diot. Characterization of Failures in an Operational IP Backbone Network. *IEEE/ACM Transactions on Networking*, 16(4):749–762, August 2008.
- [21] M. Motiwala, M. Elmore, N. Feamster, and S. Vempala. Path Splicing. In *Proceedings of ACM SIGCOMM*, 2008.
- [22] J. Moy. RFC 2328: OSPF Version 2, 1998. <http://www.ietf.org/rfc/rfc2328.txt>.
- [23] F. Papadopoulos, D. Krioukov, M. Bogua, and A. Vahdat. Greedy Forwarding in Dynamic Scale-Free Networks Embedded in Hyperbolic Metric Spaces. In *Proceedings of IEEE INFOCOM*, 2010.
- [24] L. Saino, I. Psaras, and G. Pavlou. Hash-routing schemes for information centric networking. In *Proceedings of ACM SIGCOMM ICN Workshop*, 2013.
- [25] W. So, A. Narayanan, and D. Oran. Named Data Networking on a Router: Fast and Dos-resistant Forwarding with Hash Tables. In *Proceedings of ANCS*, 2013.
- [26] N. Spring, R. Mahajan, D. Wetherall, and T. Anderson. Measuring ISP topologies with Rocketfuel. *IEEE/ACM Transactions on Networking*, 12(1):2–16, 2004.
- [27] M. Tortelli, L. A. Grieco, G. Boggia, and K. Pentikousis. Cobra: Lean intra-domain routing in ndn. In *Proceedings of IEEE CCNC*, 2014.
- [28] D. Turner, K. Levchenko, S. Savage, and A. C. Snoeren. A Comparison of Syslog and IS-IS for Network Failure Analysis. In *Proceedings of ACM IMC*, 2013.
- [29] M. Wahlisch, T. Schmidt, and M. Vahlenkamp. Lessons from the past: Why data-driven states harm future information-centric networking. In *Proceedings of IFIP Networking*, 2013.
- [30] Y. Wang, Y. Zu, T. Zhang, K. Peng, Q. Dong, B. Liu, W. Meng, H. Dai, X. Tian, Z. Xu, H. Wu, and D. Yang. Wire Speed Name Lookup: A GPU-based Approach. In *Proceedings of USENIX NSDI*, 2013.
- [31] C. Yi, A. Afanasyev, I. Moiseenko, L. Wang, B. Zhang, and L. Zhang. A Case for Stateful Forwarding Plane. *Computer Communications: ICN Special Issue*, 36(7):779–791, April 2013.
- [32] C. Yi, A. Afanasyev, L. Wang, B. Zhang, and L. Zhang. Adaptive Forwarding in Named Data Networking. *ACM SIGCOMM CCR*, 42(3), 2012.
- [33] H. Yuan, T. Song, and P. Crowley. Scalable NDN forwarding: Concepts, issues, and principles. In *Proc. of IEEE ICCCN*, 2012.
- [34] L. Zhang et al. Named Data Networking (NDN) Project. Technical Report NDN-0001, October 2010.