Scalable Name-Based Inter-Domain Routing for Information-Centric Networks

Sangmun Kim
Florida State University
kim@cs.fsu.edu

Zhenhai Duan
Florida State University
duan@cs.fsu.edu

Fernando Sanchez
Universidad San Francisco de Quito
fsanchez@ufsq.edu.ec

***Abstract*—Given the large volume of content on the Internet, scalability is one of the critical challenges in the design and development of information-centric network (ICN) architectures. In this paper we develop a scalable name-based inter-domain routing (NIDR) system that can meet the demanding requirement of supporting the large volume of content on the Internet in ICNs. NIDR adopts two critical techniques to improve its scalability. First, NIDR assumes a URL-like hierarchical content naming structure, and routes content-request packets based on only the *domain name* of content instead of the *complete content name*. Second, given the large number of domain names on the Internet, only a small subset of Internet domain names are announced at the inter-domain level. Domain names unknown to the NIDR system are first mapped (and routed) to their corresponding attachment point (AP) networks that are in the NIDR system, which is supported by a name resolution service. In addition to presenting the design of the NIDR system, we also evaluate the performance of NIDR and compare it with both BGP and EPIC, an enhancement over BGP. Our simulation studies based on the simBGP simulator show that NIDR can perform comparably with EPIC.**

I. Introduction

Information-centric network (ICN) architectures have attracted a lot of research attention in recent years [1], [24]. Compared to the host-centric, point-to-point communications architecture of the current Internet, ICNs possess many potential advantages in improving the content availability, content response time, network security, among others [13]. In addition, ICNs decouple content from the location (IP address) of the content at the network layer. As a consequence, ICNs have the potential to completely remove the IP addresses used in the current Internet architecture [21], if an ICN with proper name-based routing and forwarding is universally deployed on the Internet.

Despite the potential advantages of ICNs over the current host-centric Internet architecture, ICNs also face a few critical challenges before it can be globally deployed on the Internet [1], [10], [24]. One of the critical challenges faced by ICNs is the scalability of such a system [2], [10], given the sheer-volume and ever-increasing number of data items on the Internet. (In this paper we use the terms *content* and *data* interchangably.) For example, considering webpages alone, there are at least 10^{12} unique URLs on the Internet [11], as reported in the year of 2008. It is dauntingly prohibitive to design any global routing schemes to support the request and delivery of all the data items on the Internet directly based on their content names.

In this paper we present a scalable name-based inter-domain routing (NIDR) system that can meet the demanding requirement of supporting the high volume of content on the Internet. NIDR adopts two critical techniques in order to improve its scalability. First, NIDR assumes a URL-like hierarchical content naming structure, where a content name consists of a domain name followed by the path name of the content, for example, example.com/movies/movie.mov. It is simpler to support content name aggregation as considered in this paper with a hierarchical naming structure. Importantly, an NIDR-based ICN node routes content-request packets based on only the *domain name* of the content instead of the *complete content name* on the global Internet.

However, as we will show later, the number of domain names on the Internet is still large; it is in general several orders of magnitude than the number of network IP prefixes on the current Internet. It will still be prohibitively expensive to handle the routing traffic associated with all the domain names of Internet content, given that Internet routers are already stretched in handling routing traffic associated with the (smaller number of) network IP prefixes. In order to further improve the scalability of NIDR, the second technique adopted by NIDR is routing indirection. More specifically, only the reachability of a small subset of domain names are announced on the global Internet. Content domain names unknown to the NIDR system are first mapped to the ISP networks that provide the Internet connectivity for the corresponding unknown domains in the NIDR. We assume that the domain names of such ISP networks are in the NIDR system, and we refer to such an ISP network as the attachment point (AP) for the corresponding unknown domain name.

The domain name of the AP network is included in a new *attachment point* field in the content-request packet, so that the packet can be routed in the NIDR system. After the content-request packet arrives at the AP network, it can be further forwarded using the original content name. We note that the introduction of the routing indirection techinque does not break the design principles of ICNs. As in traditional ICNs, content is decoupled from the location of the content in an NIDR-based ICNs. In particular, the AP field is only used to route a content-request packet when an ICN node does not

978-1-4673-8590-9/15/$31.00 ©2015 IEEE

have the routing information of the original domain name of the requested content. Other components of an ICN node are not affected by the AP field in a content-request packet. For example, cached content in a content store of an ICN node is only associated with the orginal content name. More detailed information is provided in Section III.

The mapping from an unknown domain name to the corresponding AP is carried out using a name resolution service (NRS), for example, similar to the current DNS system. The details of NRS are left out from the current paper. In addition to presenting the design of the NIDR system in this paper, we also conduct performance studies of NIDR and compare it with both BGP [19], the current *de-facto* Internet inter-domain routing protocol, and EPIC [5], an enhancement over BGP to improve the convergence property of BGP. Our simulation studies based on the simBGP simulator [20] show that NIDR can perform comparably with EPIC in terms of both the number of update messages and convergence time following a failure event.

The remainder of the paper is structured as follows. In Section II we will motivate the NIDR system and briefly describe the related work. In Section III we will present the design of the NIDR system. We will conduct performance studies in Section IV and discuss further improvement of NIDR in Section V. We conclude the paper in Section VI.

II. Motivation and Related Work

In this section, we will first motivate the design of NIDR, and then we will briefly discuss the related work, focusing on the research efforts on developing scalable routing on the global Internet in ICNs.

A. Motivation

In order to develop a scalable name-based inter-domain routing system for ICNs, we need to make a few decisions including the granularity of content reachability information to be propagated on the Internet, and how far such information should be propagated. Given the large volume of content on the Internet, it is infeasible (and indeed unnecessary) to announce the reachability information (names) of all content on the Internet. Coarser-grained content reachability information should be propagated on the Internet, for example, the domain-level names of content, as suggested in the original paper of content-centric networking [13]. In order to understand the feasibility of announcing the reachability of all domain names on the Internet, in Figure 1 we show the number of registered domain names (in millions) across all top-level domains from the year of 2008 to 2014 [23].

As we can see from the figure, the number of registered domain names is large and growing fast. To put these numbers in perspective, in Figure 2 we show the number (in thousands) of network IP prefixes and AS numbers (ASNs) in the BGP-based inter-domain routing system on the Internet [6], over the same period of Figure 1. As we can see from the figures the number of registered domain names is in general several orders of magnitude more than that of the network IP prefixes and ASNs in the BGP routing system over the same time. For example, at the end of 2014, there were above 280 *millions* of registered domain names. In contrast, there were only about 525 *thousands* of network IP prefixes in the BGP routing system at the same time.

Given that Internet routers are already stretched in handling the routing traffic associated with the (smaller number of) network prefixes, it is evident that propagating the reachability information of all content on the global Internet, even at the aggregated domain-name level, may not be feasible. In addition to the granularity of content reachability information to be propagated, we also need to consider how far content reachability information should be propagated, which will similarly affect the amount of routing traffic that an Internet router needs to handle. The decision can be affected by a number of factors, including, for example, multi-homing [22], a common practice adopted by many networks to improve the availability and user-perceived reliability of Internet connectivity.

Should multi-homing not be supported on the Internet, some simpler routing scheme can be conceived. (In the context of this discussion, multi-homing refers to the practice for a network to connect to more than one providers.) For example, ISP networks other than tier-1 networks do not need to maintain the complete routing table, instead, a default route can be used to forward any packets with an unknown destination to the corresponding provider network. Only tier-1 ISP networks need to maintain a default-free routing table. As a consequence, the reachability information only needs to be propagated uphill from the originating network to the corresponding tier-1 network in the Internet inter-domain hierarchy, and does not need to be propagated downhill in the hierarchy [9]. (All tier-1 networks need to exchange network domain reachability information to build default-free routing tables.) Given that the number of network domains for which a non-tier-1 ISP needs to maintain reachability information is relatively small, and assuming tier-1 ISP networks have sufficient resources, in this case, we may be able to simply announce the reachability information of all content at the aggregated domain-name level.

However, multi-homing is a common practice and it is likely to be a preferred technique by many networks to improve the reliability of their Internet connectivity. Consequently, in the design of scalable name-based inter-domain routing system, we cannot exclude the existence of multi-homing. (The aforementioned routing scheme can still be used in multi-homing environment; however, the AS-level path taken by a packet may not be optimal.) In the design of NIDR, reachability information of domain names in the NIDR system will be propagated to all the network domains in NIDR, not only to tier-1 networks.

B. Related Work

In recent years a lot of works have been carried out in the area of ICNs. In this subsection we only briefly discuss the works that are most relevant to the current paper. We refer interested readers to [1], [24] for broader surveys on

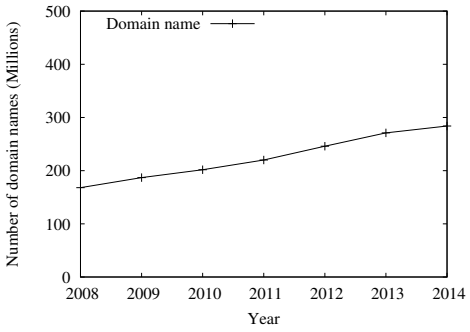

Fig. 1. Number of network domain names.

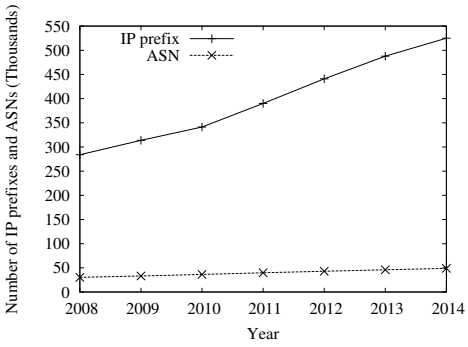

Fig. 2. Number of IP prefixes and ASNs

the works in ICNs, and [2] for a survey on the naming and routing in ICNs. Routing content at the granularity of domain names instead of complete name at the inter-domain level was proposed in the original CCN paper [13]; however, as we have shown above, given the large number of registered domain names, announcing the reachability information of all domain names on the Internet may not be feasible.

In order to scale named-data networking (NDN) [25] to handle the large volume of content on the Internet, an NDN FAQ [18] suggested to attach the ISP name of a content producer to an Interest packet, and the ISP name will serve as a Forwarding Hint. The Interest packet is forwarded based on the Forwarding Hint, if the content name is not recognized by a router. However, it did not discuss how a content requester obtains the ISP name information in the first place. Similarly, Lee *et al.* [15] also proposed to map the (unknown) domain name of an Interest packet to the hosting autonomous system (AS) of the content using a global mapping service such as DNS, and to route the Interest packet based on the AS number (ASN) of the hosting AS. However, it did not discuss any inter-domain routing protocol (for propagating domain name reachability information). In addition, it did not discuss the impact of the ASN-based routing on the ICN routers in terms of both components and packet processing. Compared to these two pieces of work, we develop a complete NIDR system.

In CONET [7], a name-based routing scheme called *lookup-and-cache* was developed. In this routing system, a router may not contain the routing information for all content names. When a router needs to route an Interest packet whose name is not in the routing table, the router will look up the routing information of the name in a DNS-like *name-system*, and the returned routing information is also cached by the local router. While the DNS-like name-system in CONET is used to return *routing* information for an unknown name, the name-resolution service (NRS) considered in this paper is used to return the *attachment point* information for an unknown (domain) name. In addition, we develop a scalable name-based inter-domain routing protocol, instead of relying on traditional routing protocols such as BGP to provide the reachability information of the small set of (popular domain) names as done in CONET.

III. Design of NIDR

In this section we present the design of the scalable name-based inter-domain routing (NIDR) system. We will first outline the assumptions made in the design of NIDR and provide a brief overview of NIDR and the operations of an NIDR-aware ICN. We will then describe the packet format and processing at an ICN router supporting NIDR. Next, we will develop the scalable NIDR system. Toward the end of this section, we will provide a discussion on other aspects of NIDR.

A. Assumptions and Overview

NIDR assumes a hierarchical inter-domain network architecture that is identical to that of the current Internet [9]. All networks have a domain name as in the current Internet. In addition, we assume that a name-based intra-domain routing protocol such as NLSR [12] is adopted by individual networks for propagating the reachability information of content originated inside their corresponding networks. In addition, although the basic ideas of NIDR can be applied to any ICNs with a hierarchical content naming structure, to make our discussions concrete, in this paper, we present NIDR in the framework of CCN [13] (or NDN [25]). To ease exposition, in the following we will simply refer to an NIDR-aware ICN as an ICN, as long as there is no confusion. We will continue referring to CCN and NDN as examples of traditional ICNs that do not support NIDR.

In NIDR, Internet domain names are classified into two categories based on if their reachability information is announced and propagated in the inter-domain routing system. The first category includes all the domain names that are announced in NIDR, and we refer to such domains as *routed domains* (RDs). In essence, all the networks with an AS number (ASN) in the current BGP-based Internet inter-domain system can be an RD, including all ISP networks and some large enterprise networks. The reachability information of RD names are announced and propagated in the NIDR system, and Interest packets destined to RD names can be directly routed in the NIDR system.

The second category includes all the domain names that are not announced in NIDR, and we refer to such domains as *lookup domains* (LDs). When an ICN node receives an Interest packet with an LD name and does not have routing

information for the LD name, a name-resolution service (NRS) lookup is performed and a new *Attachment Point* (AP) field is populated in the Interest packet, where the AP is an RD and knows how to route the packet based on the content (domain) name. After the AP field is populated, the Interest packet is then routed based on the AP field. After the packet arrives at the AP domain (assuming the Interest packet has not been satisfied in the previous ICN nodes on the way to the AP domain), the packet is further routed based on the content (domain) name.

As an example, consider content with name example.com/movies/movie.mov, and assume that example.com is an LD network whose domain name is not announced in the NIDR system. Furthermore, we assume that the corresponding AP network for example.com is isp.com. When a router (or the originating machine) receives an Interest packet for the content, it first determines the AP network of the domain name example.com using NRS (assuming the node does not have the routing information for domain name example.com), and then includes the AP domain name isp.com into the attachment point field of the Interest packet. The packet can then be forwarded by the intermediate network domains based on the AP domain name. When the Interest packet arrives at isp.com, it can be further forwarded by isp.com to example.com.

The mapping from an LD name to its corresponding AP is performed by a name-resolution service (NRS), for example, similar to the current DNS system. In this paper, we will only focus on the design of NIDR and will leave out the details of NRS as future work.

B. Packet Format and Processing

In order to carry the attachment point (AP) information in an Interest packet, we include a new field named *attachment point* in the Interest packet in an NIDR-aware ICN. Figure 3 shows the format of Interest packets in the ICN. For the simplicity of illustration, in the figure we show the new "Attachment Point" field in Interest packets based on the original packet format introduced in [13]. The new field can be included in the newer format of Interest packets in the same manner [16]. The new AP field will contain the domain name of the AP network for the corresponding LD name of the content.

When an ICN router needs to forward an Interest packet but does not have the routing information for the domain name of the content, it will issue an NRS lookup request packet to obtain the corresponding AP information of the domain name, if the AP field is empty. After the AP information is returned, the ICN router will insert it into the AP field of the Interest packet. We note that it is likely that this NRS lookup will be performed by the originating machine (or the first hop router) of the Interest packet. All the other fields in an Interest packet in ICN are identical to the fields in an Interest packet in the original CCN. Similarly, the format of Data packets in an NIDR-aware ICN is the same to the format of Data packets in CCN.

The packet processing engine in an ICN node is also very similar to that in the original CCN. Recall that a CCN node

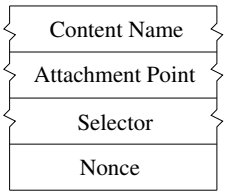

Fig. 3. Format of Interest packets in NIDR-aware ICN.

(or router) contains three main components: the forwarding information base (FIB), content store, and pending interest table (PIT). In an NIDR-aware ICN, the content store and PIT are not changed compared to those in the original CCN. We emphasize that, although a new AP field is included in Interest packets, content in ICN is only associated with its original content name, *not the AP field*, for both content caching and delivery. In particular, a data item in a content store at an ICN node maintains its original content name, and entries in PIT also only contain the original content name of Interest packets, without the AP field. This will minimize the impact of the new AP field on both the design of ICN routers and the management overhead at AP networks and the corresponding LD networks.

The AP field is only used when an ICN node needs to look up the FIB to continue forwarding an Interest packet (when the node does not have the routing information of the corresponding content name). The FIB is populated using the name-based routing algorithms such as NLSR and NIDR (see the next subsection on the details of NIDR). The AP field is only used during an FIB lookup in order to forward an Interest packet. All other aspects of an NIDR-aware ICN node is identical to that in CCN, including the basic packet forwarding procedure performed at an ICN node after receiving an Interest packet or a Data packet. We omit the details here and refer interested readers to CCN [13] or NDN [25].

We emphasize that despite the introduction of the new AP field in an Interest packet, NIDR does not violate the basic design principles of ICNs, in particular, the principle that content should be decoupled from the location of the content. An Interest packet can be satisfied by an intermediate NIDR-aware ICN node if its content store has the requested content. It is not necessary to always deliver an Interest packet with a populated AP field to the AP domain.

C. Name-based Inter-Domain Routing

In order to achieve higher scalability in NIDR, only the reachability information of a small subset of domain names are exchanged at the inter-domain level on the Internet, including all ISP networks and some large content provider networks and enterprise networks. Recall that we refer to these domains as routed domains (RDs). In essence, all network domains with an AS number (ASN) in the current BGP-based routing system can be an RD. The reachability information of RD names is exchanged using NIDR, which is a path-vector routing protocol at the Internet inter-domain level. Each network has a domain name, and the networks (autonomous systems, or

ASes) participating in the NIDR system also have an ASN as on the current Internet. All routers participating in NIDR have a domain name. We note that as a common practice, the domain name of a router in an ISP network normally contains some geographical or POP location information, which can facilitate the debugging of routing problems.

NIDR is a path-vector routing protocol, and to a large degree, it is similar to the BGP routing protocol on the current Internet. However, a number of inherent features of ICNs simplifies and improves the design of NIDR. For example, the security of NIDR can be greatly improved compared to the BGP, due to the inherent security feature of ICNs [13]. Indeed, regarding the security support, NIDR is closer to S-BGP [14] than BGP. NIDR adopts the basic message format of BGP, with two principal changes. First, while BGP messages use IP addresses or IP prefixes, NIDR messages use domain names or name prefixes. Second, NIDR messages may contain additional information for improving the security or convergence properties of NIDR, which we will describe later, after presenting the basic messages of NIDR.

Like in BGP, the reachability information of domain names is carried in an UPDATE message [19], [22], which is in turn carried in an ICN Data packet [13]. The content name of a Data (and Interest) packet related to an UPDATE message is router_domain_name/nidr/update. We discuss the exchange of Interest and Data packets between two neighboring NIDR routers after we describe the NIDR protocol. In an UPDATE message, NLRI (network layer reachability information) refers to the domain name of an AS instead of IP prefix as in BGP. Similarly, the NEXT-HOP attribute refers to the domain name of a NIDR router instead of its IP address. The AS-PATH attribute contains the sequence of ASNs of the networks that an UPDATE message has traversed. Before an AS sends an UPDATE message to a neighboring AS, it will prepend its own ASN onto the AS-PATH.

In order to improve the security situation of NIDR, a route attestation is also carried in an UPDATE message as in S-BGP, where each AS signs the part of AS-PATH up to and including the ASN of the neighboring AS to which the message is sent (the signature covers the ASN of downstream neighboring AS to prevent some network reachability hijacking attack on the Internet [17]). This security feature can also be achieved by relying on the signature carried in each ICN Data packet. However, encapsulating the entire Data packet received from a neighboring AS in an UPDATE message could greatly increase the overhead of NIDR. Consequently, we let each AS sign the concerned portion of AS-PATH in an UPDATE message, instead of encapsulating all previous Data packets related to the UPDATE message. In this paper, we do not discuss the security key management, and assume its existence as part of ICN deployment.

Moreover, in order to improve the convergence performance of NIDR, all UPDATE messages in NIDR carry a list of sequence numbers so that a receiving NIDR router can distinguish an old UPDATE message from a new one and can eliminate obsolete routing information from the routing table [5]. The sequence numbers are maintained and inserted by border routers in the NIDR system. Other route attributes developed in BGP, in addition to NEXT-HOP and AS-PATH, can be similarly extended to NIDR and we omit them here.

D. Discussion

As we have discussed, NIDR messages are exchanged using the ICN Interest/Data packets. However, network events affecting network reachability can occur at any time, and need to be propagated to other nodes as early as possible so that the routing system can converge to another stable state. On the other hand, it is not possible for a node to promptly issue an Interest packet in response to a network event in order to get the Data packet (UPDATE message). Given these observations, we adopt a simple strategy for NIDR border routers to exchange reachability information, where long-term Interest packets are used in the context of NIDR [4], [18]. In particular, when an NIDR border router is up, an Interest packet concerning NIDR messages will be sent to each of its neighbors. These NIDR Interest packets are long-term in the sense that, a neighbor will not delete an Interest packet from the PIT after forwarding a Data packet (NIDR messages) to the node.

So far we only have discussed NIDR for supporting single path routing. However, we note that existing multi-path routing techniques such as the ones discussed in [8] can be easily adopted in NIDR. Importantly, due to inherent support of multicast and multi-path routing in ICNs. we believe that multi-path routing in NIDR could be potentially simplified and improved compared to the multi-path routing in BGP.

One shortcoming of NIDR is that it slightly increases the overhead of forwarding a packet as an AP domain name is included in an Interest packet. However, we note that the overhead in general is relatively small, given that the domain names are normally short (which is particularly the case for ISP networks) and the speed of modern routers are normally high. In addition, since only a small subset of domain names are propagated in the NIDR system, it is likely that more efficient lookup algorithms can be developed to speed up the forwarding of packets. Moreover, the performance of the NIDR-based inter-domain routing system could also be improved with a small number of domain names in the routing system.

IV. Performance Studies

In this section we perform simulation studies to illustrate the performance of NIDR, and contrast it with BGP and EPIC [5]. EPIC is an enhancement over BGP to improve its convergence property using *forward edge sequence numbers* (*fesn*). We implement NIDR in the simBGP simulator [20], which has implemented BGP and EPIC.

A. Simulation Set-Up

In the simulation studies, we use two different topology families—Clique (i.e., complete graph) and Waxman random topologies. The Waxman topologies are generated using the

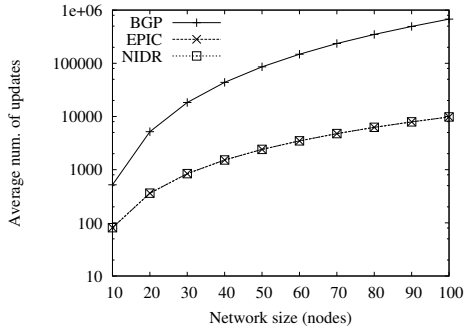

Fig. 4. Number of messages (Clique, fail-down).

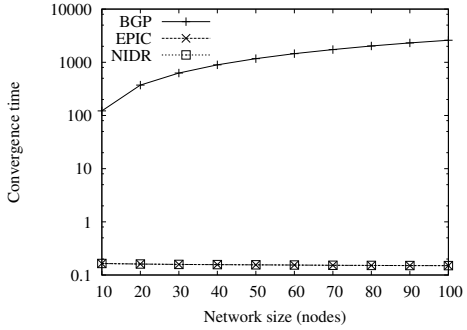

Fig. 5. Convergence time (Clique, fail-down).

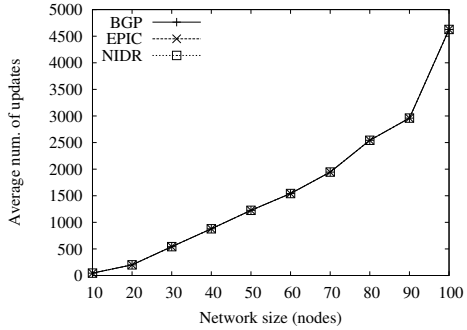

Fig. 6. Number of messages (Clique, fail-over).

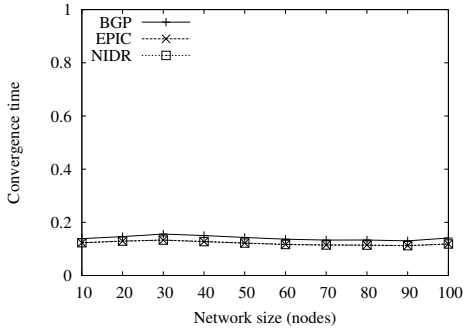

Fig. 7. Convergence time (Clique, fail-over).

Brite topology generator [3] with both α and β set to 0.5. The propagation delay on each link is chosen randomly between 0.01 and 0.1 seconds. The processing delay on each node is chosen randomly between 0.001 and 0.01 seconds. The values of other parameters are set to their default values, for example, the MRAI timer is set to 30 seconds.

We simulate both link fail-down and fail-over events. To create a fail-down event, we attach a dummy node to a randomly chosen node in the network topology. We fail this link during the simulation. To create a fail-over event, we attach a dummy node to *two* randomly chosen nodes in the topology. We randomly fail one of the two links between the dummy node and the topology. To simplify the simulation set-up, only the dummy node announces a network IP prefix (for BGP and EPIC) or domain name (for NIDR). All other nodes do not announce prefixes or domain names. We repeat the simulation 30 times, each with different attach points and random seeds.

For each simulation run, we ensure that the routing system is stable before the failure event occurs. We summarize the total number of update messages (including both withdrawals and announcements) sent after the failure event during the simulation, and the time it takes for the routing system to converge to a stable state (i.e., the convergence time). We then compute and report the average number of update messages and the average convergence time over the 30 simulation runs.

We note that in these simulation studies, we focus on the convergence performance of NIDR. The performance gain of NIDR in announcing a small subset of all domain names is illustrated in Figures 1 and 2. In general the adoption of NIDR can result in several orders of magnitude reduction in the number of domain names announced at the inter-domain level on the Internet, compared to the traditional ICN designs. In this section we evaluate the convergence property of NIDR and compare it with BGP and EPIC. In this comparison, we can consider that all the ASes in the current BGP-based Internet inter-domain routing system participate in the NIDR system and announce their domain names. Domains not in the BGP routing system are lookup domains in the NIDR system and their names are not announced.

In addition, we do not evaluate the performance impact of the security features of NIDR, which could increase the processing time of update messages in NIDR. These processing times are ignored in the simulation studies and left as future work (but see [14] for a general understanding of the processing overhead caused by the added security features).

B. Simulation Results

Figures 4 and 5 show the average number of update messages and convergence time (in seconds) following a link fail-down event in the Clique topologies, respectively. We first note that, in essence, NIDR performs identically with EPIC, because both of them utilize the technique of sequence numbers to limit the path exploration problem in BGP. In addition, both of them outperform BGP in terms of both the number of update messages and convergence time. We

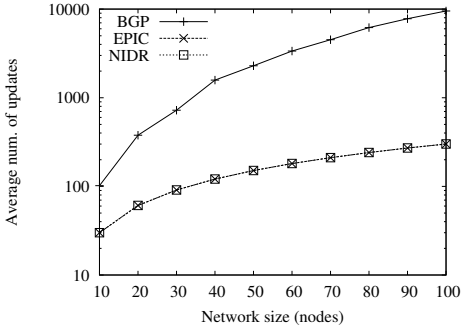

Fig. 8. Number of messages (Waxman, fail-down).

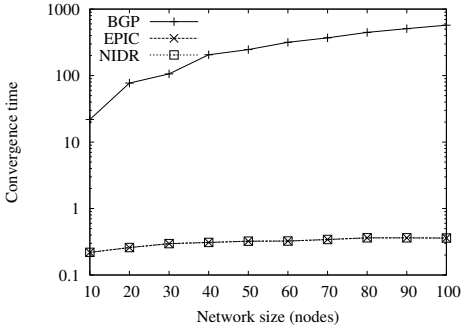

Fig. 9. Convergence time (Waxman, fail-down).

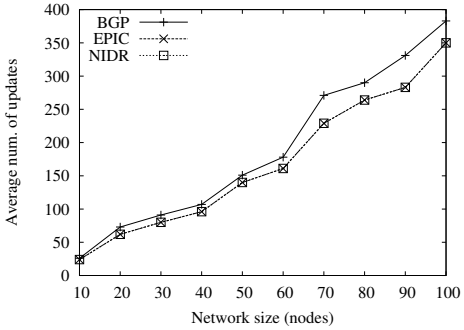

Fig. 10. Number of messages (Waxman, fail-over).

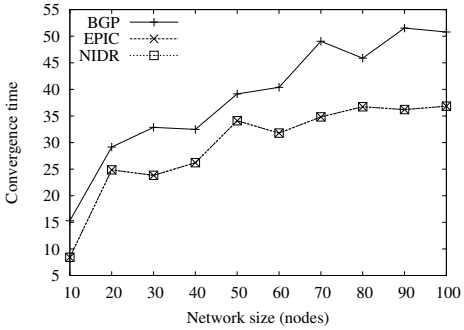

Fig. 11. Convergence time (Waxman, fail-over).

note that in this implementation we use long-term Interest packets for routing messages in NIDR, which are delivered at the beginning of the simulation run, therefore, they are not counted in the number of update messages following a failure event. Otherwise (if an Interest packet is delivered following a failure event), the performance of NIDR could be slightly worse than EPIC. We adopt this behavior for the simulation runs for NIDR.

Figures 6 and 7 show the average number of update messages and convergence time following a fail-over event in the Clique topologies, respectively. Based on Figure 6 we can see that following a link fail-over event, the three routing protocols—BGP, EPIC, and NIDR—generate the same number of updates messages. Note that, in a *clique* network topology, a node will choose the *valid* alternate route to the dummy node, no matter which of the three protocols is used. However, we can see from Figure 7 that EPIC and NIDR have slightly smaller convergence time compared to BGP, although they generate the same number of update messages. In essence, this difference is caused by the faster convergence due to the adoption of sequence numbers in EPIC and NIDR. Because of the introduction of sequence numbers, obsolete routing information can be removed earlier in EPIC and NIDR than in BGP. Figure 12 shows a simple clique network topology with four nodes A, B, E, and F, with a dummy node D connected to two nodes A and E.

Consider the routing activities at node F. Assume that node F prefers the route via A to reach destination node D, i.e.,

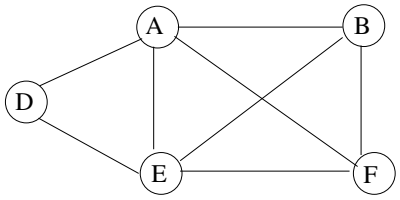

Fig. 12. A simple clique network topology.

the AS_PATH is [FAD] at the beginning. Let $d_{i,j}$ denote the propagation delay on the link between node i and node j. We consider a case where $d_{A,F} > d_{A,B} + d_{B,F}$ so that when node A sends a routing message to node B and node F, the routing message arrives at node F via node B first (after some processing at node B), ahead of the routing message directly delivered from node A to node F.

Now assume the link between nodes D and A fails, and all nodes should switch to the alternate route to reach node D via E. In particular, after node A chooses alternate path [ED], it will send the new path [AED] to node B and node F, respectively. However, given the assumption discussed above, the update message containing AS_PATH [BAED] arrives at node F (sent by node B) ahead of the update message containing [AED] (sent by node A). Now let us see what happens in BGP and EPIC (and NIDR). In BGP, node F will continue using the old best route via node A, i.e., [AD], because it is more preferred over the new path sent by node B, i.e., [BAED]. Only after receiving the alternate path [AED]

will node F switch to a new best path [ED]. In contrast, in EPIC (and NIDR), after receiving the alternate path [BAED], node F will eliminate all the obsolete paths, including the current best path [FAD], and switch to the new best path [ED]. Based on this example, we can see that although BGP, EPIC, and NIDR may generate the same number of update messages, EPIC and NIDR may still converge faster than BGP in terms of time.

Figures 8 and 9 show the average number of update messages and convergence time following a fail-down event in the Waxman random topologies, respectively. Figures 10 and 11 show the performance of the three routing protocol in the fail-over event for the Waxman random topologies. From the figures we can make the same observations as we have made in the failure events for the Clique topologies. In particular, NIDR and EPIC essentially have the same performance in terms of both number of update messages and convergence time, and both of them outperform BGP.

In summary, we can see that NIDR essentially has the same performance as EPIC, and outperforms BGP. NIDR and EPIC have the same performance because they both adopt the same technique of sequence numbers to limit the path exploration problem in BGP.

V. Discussions

In this paper we have only considered NIDR as a routing protocol *between* network domains, in the same way as the external BGP (E-BGP) between network domains. NIDR can be used to exchange routing information among border routers in the same network domain, in the same way as the internal BGP (I-BGP) [19], [22]. Given the design of NIDR, it is easy to see that the techniques developed for I-BGP can be similarly applied to NIDR as a routing protocol between border routers in the same AS (I-NIDR), such as route reflection. However, we note that, ICNs have a few built-in features, including, for example, routing loop prevention and native support of multicast, which may greatly simplify the design of I-NIDR and advanced techniques for deploying I-BGP may not be needed including route reflection. We will investigate the details of I-NIDR in a separate work.

VI. Summary and Future Work

In this paper we have presented and evaluated a scalable name-based inter-domain routing (NIDR) system for information-centric networks. NIDR adopted two critical techniques to improve its scalability—hierarchical content naming structure and routing indirection. In NIDR, only the reachability information of a small subset of all domain names are announced and propagated on the global Internet at the inter-domain level. Other domain names are looked up via a name-resolution service (NRS) and routed based on their corresponding attachment points carried in a new field in the content-request packets. NIDR resulted in several orders of magnitude reduction in the number of domain names announced at the inter-domain level. In addition, our simulation studies also showed that NIDR can perform comparably with EPIC and outperform BGP. As future work, we will fully investigate the details of the security support of NIDR and other aspects of NIDR, including I-NIDR and multi-path routing.

References

[1] B. Ahlgren, C. Dannewitz, C. Imbrenda, D. Kutscher, and B. Ohlman. A survey of information-centric networking. *IEEE Communications Magazine*, 50(7):26–36, 2012.

[2] M. Bari, S. Chowdhury, R. Ahmed, R. Boutaba, and B. Mathieu. A survey of naming and routing in information-centric networks. *Communications Magazine, IEEE*, 50(12):44–53, 2012.

[3] BRITE. Boston university Representative Internet Topology gEnerator. http://www.cs.bu.edu/brite/.

[4] A. Carzaniga, M. Papalini, and A. L. Wolf. Content-based publish/subscribe networking and information-centric networking. In *Proceedings of the ACM SIGCOMM workshop on Information-centric networking*, pages 56–61. ACM, 2011.

[5] J. Chandrashekar, Z. Duan, Z.-L. Zhang, and J. Krasky. Limiting path exploration in BGP. In *Proc. IEEE INFOCOM*, Miami, FL, Mar. 2005.

[6] CIDR report. http://www.cidr-report.org.

[7] A. Detti, N. Blefari Melazzi, S. Salsano, and M. Pomposini. CONET: a content centric inter-networking architecture. In *Proceedings of the ACM SIGCOMM workshop on Information-centric networking*, pages 50–55. ACM, 2011.

[8] C. Filsfils, P. Mohapatra, J. Bettink, P. Dharwadkar, P. De Vriendt, Y. Tsier, V. Van Den Schrieck, O. Bonaventure, and P. Francois. BGP prefix independent convergence (PIC) technical report. Technical report, Cisco, 2011.

[9] L. Gao. On inferring autonomous system relationships in the internet. *IEEE/ACM Transactions on Networking*, 9(6), Dec. 2001.

[10] A. Ghodsi, S. Shenker, T. Koponen, A. Singla, B. Raghavan, and J. Wilcox. Information-centric networking: seeing the forest for the trees. In *Proceedings of the 10th ACM Workshop on Hot Topics in Networks*, page 1. ACM, 2011.

[11] Google. We knew the web was big... http://googleblog.blogspot.com/2008/07/we-knew-web-was-big.html.

[12] A. Hoque, S. O. Amin, A. Alyyan, B. Zhang, L. Zhang, and L. Wang. NLSR: named-data link state routing protocol. In *Proceedings of the 3rd ACM SIGCOMM workshop on Information-centric networking (ICN)*, pages 15–20. ACM, 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 the 5th international conference on Emerging networking experiments and technologies (coNEXT)*, pages 1–12. ACM, 2009.

[14] S. Kent, C. Lynn, and K. Seo. Secure border gateway protocol (S-BGP). *IEEE Journal on Selected Areas in Communications*, 18(4):582–592, 2000.

[15] J. Lee, T. K. Hoon-gyu Choi, and Y. Choi. Scalale routing for content centric network. In *Proceedings of AsiaFI 2012 summer school*, Kyoto, Japan, Aug. 2012.

[16] M. Mosko. Ccnx 1.0 protocol introduction. Technical report, Palo Alto Research Center, Apr. 2014.

[17] S. Murphy. BGP security vulnerabilities analysis. RFC 4272, Jan. 2006.

[18] NDN frequently asked questions. http://named-data.net/project/faq/.

[19] Y. Rekhter, T. Li, and S. Hares. A border gateway protocol 4 (BGP-4). RFC 4271, Jan. 2006.

[20] simBGP. simbgp: a simple bgp simulator. http://www.bgpvista.com/simbgp.php.

[21] W. So, A. Narayanan, D. Oran, and M. Stapp. Named data networking on a router: forwarding 20gbps and beyond. In *Proceedings of ACM SIGCOMM Demonstration*, Aug. 2013.

[22] J. Stewart. *BGP4: Inter-Domain Routing In the Internet*. Addison-Wesley, 1999.

[23] Verisign. Domain name industry brief. http://www.verisigninc.com/en_US/innovation/dnib/index.xhtml.

[24] G. Xylomenos, C. Ververidis, V. Siris, N. Fotiou, C. Tsilopoulos, X. Vasilakos, K. Katsaros, and G. Polyzos. A survey of information-centric networking research. *IEEE Communications Survey and Tutorial*, 16(2):1024–1049, 2014.

[25] L. Zhang, kc claffy, P. Crowley, C. Papadopoulos, L. Wang, and B. Zhang. Named data networking. Technical Report NDN-0019, NDN, June 2014.