

On Inter-Domain Name Resolution for Information-Centric Networks

Konstantinos V. Katsaros, Nikos Fotiou, Xenofon Vasilakos,
Christopher N. Ververidis, Christos Tsilopoulos,
George Xylomenos, and George C. Polyzos*

Mobile Multimedia Laboratory,
Department of Informatics,
Athens University of Economics and Business,
Patision 76, Athens 104 34, Greece
{ntinos,fotiou,xvas,chris,tsilochr,xgeorge,polyzos}@aueb.gr
<http://mm.aueb.gr>

Abstract. Information-centric networking (ICN) is a paradigm that aims to better reflect current Internet usage patterns by focusing on information, rather than on hosts. One of the most critical ICN functionalities is the efficient resolution/location of information objects i.e., name resolution. The vast size of the information object namespace calls for a highly scalable and efficient name resolution approach. Currently proposed solutions either rely on a DHT structure, thus ensuring load balancing and scalability at the cost of inefficient routing, or on hierarchical structures, thus preserving routing efficiency at the cost of limited scalability. In this paper, we study in detail the tradeoff between state/signaling overhead versus routing efficiency for a generic name-resolution system based on a novel DHT scheme with enhanced routing properties, and compare it to DONA, an ICN architecture based on hierarchical resolution and routing.

Keywords: content-centric, future internet, named data, rendezvous.

1 Introduction

The Internet has been transformed from an academic curiosity to a global infrastructure for the massive distribution of information, with over 1 billion of connected devices [6], 1 trillion of indexed web pages [12] and Exabytes of annually transferred data [6]. Unlike this evolution in usage, at its core the Internet continues to operate upon a host-centric communication model, even though obtaining and disseminating information is currently the primary interest of users. This tussle between the core Internet functionalities and its actual usage has led the networking community to investigate new architectures for the Future Internet, many of which revolve around information-centrism (e.g. [15],[17]).

* Work supported by the ICT PURSUIT project, under contract ICT-2010-257217.

Information-centric networking (ICN) places information dissemination at the heart of the architecture. *Information Objects* (IOs) constitute the fundamental entity in ICN. Users in ICN issue requests by naming the desired IOs in order to obtain information; it is the responsibility of the network to locate and deliver the desired IOs to the requester, thus decoupling the communicating parties: end-users need not be aware of each other's network location or, even, identity. Locating IOs is therefore a core function in the context of ICN. In effect, the network operates as a system in which producers of information announce the availability of IOs and requests for IOs are resolved through a name-based anycast scheme, resulting in a *name-resolution system* (NRS).

To face the scalability requirements imposed by the vast number of IOs, several research initiatives have considered the use of *Distributed Hash Table* (DHT) schemes (e.g., [7,16]), which exhibit significant scalability due to their inherent load balancing. However, this advantage comes at the cost of inefficient routing, which may overlook physical network proximity, administrative domain boundaries, routing policies, or a combination thereof. Other approaches, such as DONA [17], adapt to the underlying routing mechanism, requiring however extensive replication of routing information across the inter-domain topology.

In this paper, we examine the emerging tradeoff between state/signaling overhead versus routing efficiency in the context of ICN. To this end, we first design a *DHT-based name-resolution system* (DHT-NRS) based on H-Pastry, a hierarchical version of the Pastry DHT [22]; H-Pastry adapts to the underlying network topology, administrative structure and routing policies. We then present a detailed performance evaluation and comparison against the DONA scheme. Our purpose is to shed light on the expected performance of each approach, paying particular attention to the effect of the underlying inter-domain topology and the traffic workload. Our evaluation is based on a full-fledged simulation environment and considers a wide range of performance aspects, including routing efficiency and signaling overhead, as well as memory and processing load.

In the following, we describe name-resolution in the context of ICN and present our requirements for the envisioned name-resolution system (Section 2). Next, we discuss alternative approaches and position our work in the solution space (Section 3). In Section 4 we first describe the H-Pastry DHT scheme and then proceed with the description of DHT-NRS. We present and discuss the results of our performance evaluation in Section 5 and conclude in Section 6.

2 Name-Resolution in ICN

In ICN, the network focuses on information itself rather than on the endpoints participating in the communication. The network acts as a mediator that decouples content providers from content consumers, radically changing the model of interaction with the network. The network (a) accepts end-user requests for IOs with the purpose of locating the content and enabling its delivery and (b) interacts with content providers which supply information about the availability and location of content. The basic functionality of the NRS in ICN is to match

end-user requests with the accumulated information on available content and trigger the delivery of the corresponding data. This matching is based on the IO name, which thus becomes a first class entity in ICN. In this context, we outline below a set of requirements for an NRS, as derived by its central role in an ICN architecture, also identifying the key challenges in this new paradigm.

Flat Identifiers. A critical factor in the design of an NRS is the form of the IO *identifiers* (IDs) employed. Existing systems achieve scalability by applying extensive aggregation on hierarchically structured identifiers (e.g., DNS, IP routing). A critical constraint of hierarchical systems is that the nodes in the root(s) of the tree(s) must be aware of all available prefixes in the name space. The use of flat, semantic-free identifiers has been proposed to decouple location from identity and thus achieve persistence and increased security [11,17,25]. Recent developments in DNS [14] suggest that the increasing demand for naming information cannot be easily handled within the constraints imposed by hierarchical naming. Therefore, our NRS must operate on top of a flat namespace.

Scalability. An NRS must scale as the number of items and corresponding requests grows. Considering that (a) the current amount of unique web pages indexed by Google is greater than 1 trillion [12] and that (b) billions [6] of devices ranging from mobile phones to sensors and home appliances will be joining the network to offer additional content, we should plan for unique IOs in the order of 10^{13} ; other studies raise this estimate to 10^{15} [7]. This vast amount of IOs and related requests must be handled through the NRS.

Fault Tolerance and Fault Isolation. The criticality of the NRS calls for a non-central point of failure, which is the major vulnerability of centralized architectures. Furthermore, the architecture should provide fault isolation i.e., the failure of a part of this distributed system should only have a local impact.

Low Signaling Overhead. The expected amount of available information, coupled with the corresponding demand by end-users, is expected to yield an increased signalling load for the NRS, in terms of IO announcements and requests for IOs. Hence, the required information exchange between the nodes of the distributed NRS should be kept to a minimum.

Low Latency. Resolution should complete with minimum delay, yielding low response times for end-users. This calls for efficient routing and load distribution among system nodes.

Routing Policy Compliance. Finally, due to the high volumes of expected resolution traffic, routing should respect the policies of the *Autonomous Systems* (ASes). These policies reflect business relationships established between ASes, with policy violations having an economic impact on their operation.

3 Related Work

The current equivalent of an NRS for the Internet is DNS, therefore we must first explain why DNS is inadequate for our purposes. To begin with, DNS is susceptible to *Denial of Service* (DoS) attacks. This vulnerability stems not only

from the limited redundancy in name-servers, but also by the fact that many servers have a single point of attachment to the Internet [21]. In addition, load is not equally balanced between root servers due to the fact that names are not equally distributed among top level domains (e.g. `.com` names are far more than `.edu` names). Finally, DNS supports a hierarchical namespace, which constrasts with our requirement for flat identifiers.

The limitations of DNS have led the research community to look into DHT-based approaches for name resolution e.g., [21,25,20]; typical DHTs used in these approaches are Chord [24] and Pastry [22]. The main problem with DHTs is the logarithmic number of hops required (on average) for lookup/resolution. The common workaround for this problem is caching and/or replication, as well as incrementing the average node degree in the DHT. Furthermore, overlay routing in DHT approaches suffers from non-compliance to inter-domain routing policies i.e., the overlay routes followed when resolving a name may violate the routing policies of the underlying physical network [23]. To the best of our knowledge, no solution has been proposed for this problem in any of the available DHT overlays, except for [9], which was developed for the purposes of this work.

In MDHT [7] a multilevel DHT architecture is proposed for name resolution in the ICN context. MDHT aggregates IO registration entries at higher levels of the inter-domain hierarchy, which raises scalability concerns. The proposed solution is an indirection level which allows the NRS to resolve content provider names, with the resolution of the content itself taking place at a lower level. The description of the proposed system lacks details on the operation of the employed DHT, as well as a performance evaluation of the proposed architecture. A similar approach is followed in [16], where an inter-domain name-resolution architecture is presented. Again, the aggregation of routing information at the higher levels of the hierarchy poses significant scalability concerns. In a similar manner to [7], an indirection level is introduced to map *scopes* of information to lower level resolution nodes. The presented performance evaluation is based on significant abstractions e.g., intra-domain routing overhead and caching have been coarsely modeled based on observations made in different contexts which do not reflect the proposed architecture's intrinsic characteristics. In both approaches ([7] and [16]), the routing inefficiency of DHTs is circumvented by the aggregation of information at higher levels of the inter-domain structure.

The *Data-Oriented (and beyond) Network Architecture* (DONA) [17] follows a similar approach in that it also concentrates routing information at the higher levels of the inter-domain structure. DONA builds an information-centric layer over the Internet, based on a network of *Resolution Handlers* (RHs). The deployment of the RHs strictly follows the AS-level structure of the Internet, allowing DONA to directly adapt its structure and operation to the underlying network. Flat, self-certifying names are used to represent information. Content providers issue REGISTER messages to advertise their content to the closest RH which then propagates this information upwards in the hierarchy, also forwarding it across inter-domain peering links. Name resolution is based on a similar propagation of requests (FIND messages) upwards in the hierarchy, until a relevant entry is

found. At that point, requests follow the (downhill) reverse path of the corresponding REGISTER message. In this process, shortest path routing is followed and inter-AS routing policies are taken into account.

In the work most similar to ours [5], a performance comparison of content delivery between a DONA-like, hierarchical architecture and a DHT based alternative is presented. The authors study several aspects such as transfer latency and robustness, as well as the impact of in-network caching. However, this investigation does not focus on name resolution and is based on oversimplified network topologies: the DONA-like architecture uses a tree overlay consisting of randomly selected nodes of a two-level, hierarchical inter-domain topology. Our work is far more accurate in that (i) we consider a more realistic topology model that takes into account both multihoming and peering relationships (see Section 5), (ii) we fully implement DONA, allowing the structure of the RH network to mimic the underlying inter-domain topology and (iii) we consider a DHT-scheme that better adapts to the underlying topology (see Section 4.1).

4 An Enhanced DHT-Based NRS

As already discussed, the load balancing and robustness characteristics of DHTs can directly address the scalability concerns posed for an NRS in the context of ICN. However, these features also imply several important disadvantages i.e., name resolution follows longer paths than those offered by the underlying shortest path routing fabric, often unnecessarily crossing administrative domain boundaries and/or violating the established inter-domain routing policies.

In our work we have investigated the extent to which DHT routing can be improved, without sacrificing its scalability advantages. Our purpose is to assess the ability of a DHT-based approach to support an NRS in the context of ICN. To this end, we have designed DHT-NRS, an NRS based on an enhanced DHT design named H-Pastry [9]. In the following, we first provide an overview of the H-Pastry scheme and then we proceed with the description of DHT-NRS. We do not discuss security issues, as they are outside the scope of this paper. However, existing solutions for securing DHTs are also applicable to our work.

4.1 H-Pastry

Based on Pastry [22] and the Canon paradigm [10], H-Pastry is a multi-level DHT scheme that improves routing by taking physical network proximity, administrative domain boundaries and inter-domain routing policies into account. In order to adapt to the multi-level structure of an inter-network, H-Pastry employs a corresponding multi-level structure to partition the information identifier space, and the corresponding routing state. Distinct routing tables are maintained for each level of the inter-domain topology. At each level, H-Pastry nodes maintain routing information about nodes that are numerically closer to certain points in the identifier space (i.e., their own identifier and the identifiers required to fill their Routing table) than any other node at that particular level. In this manner,

the participating H-Pastry nodes recursively create H-Pastry rings that adapt to the AS-level topology of the network. Multihoming and peering agreements are also taken into account by appropriately introducing virtual nodes in the inter-domain topology graph [9]. H-Pastry presents good fault resilience properties: as described in [9], H-Pastry routing will fail in the event of a maximum of $H \cdot \lfloor L/2 \rfloor$ concurrent failures of nodes grouped in H sets of $\lfloor L/2 \rfloor$ adjacent IDs each (the corresponding number of failures for Pastry is $\lfloor L/2 \rfloor$ where H is the height of the domain level hierarchy and L the size of the leaf set).

By taking into account the aforementioned routing aspects, H-Pastry considerably improves DHT-based routing. As shown in [9], H-Pastry results in shorter routes, lowering path stretch by 55% and 47% compared to Chord [24] and hierarchical Chord [10], respectively, being comparable to regular Pastry. At the same time, H-Pastry also manages to confine traffic within administrative boundaries resulting in 27% less inter-domain hops and 55% shorter intra-domain paths than regular Pastry. Moreover, by taking routing policies into account, H-Pastry reduces valley-free policy violations per routing path by 56%, 31% and 36% compared to Chord, Pastry and hierarchical Chord, respectively.

4.2 DHT-NRS

Basic Functionality. The name resolution function in DHT-NRS follows the Publish/Subscribe paradigm i.e., publishers (content providers) and subscribers (content consumers) interact with DHT-NRS through a set of *brokers* responsible for matching publications and subscriptions. DHT-NRS is realized by the deployment of these brokers, named *Rendezvous Nodes* (RNs), across the inter-domain topology. Each RN is an H-Pastry node with a unique ID. We envision at least one RN per AS; more RNs will increase system scalability. For each IO a statistically unique ID is created (e.g., with a secure hash function). Publishers and subscribers interact with DHT-NRS via their local domain RNs, designated during network attachment. Name-resolution is then based on the exchange of the following control messages (an example is given in Figure 1):

- **Adv:** An advertisement message sent by the publisher(s) of an IO to register the IO with DHT-NRS (step I in Figure 1). This message contains information that maps the ID of the advertized IO to the network location of the issuing publisher. For each unique ID-publisher mapping, a corresponding *IO entry* (IOE) is maintained by an RN designated by H-Pastry as responsible for that ID (denoted as the *Rendezvous Point* (RVP) for the ID). Advertisement messages are routed towards the RVP through H-Pastry.
- **Sub:** A subscription message is issued by a subscriber to request a specific IO (step II in Figure 1). Subscription messages contain the ID of the requested IO along with information on the subscriber’s network location. Subscription messages reach the appropriate RVP through H-Pastry routing. Upon reception, an RVP searches its IOEs for the indicated IO ID. If the subscription message refers to an unknown IO, the RVP discards the subscription.
- **Ntf:** A notification message is sent by an RVP to notify the publisher(s) of an IO to start the delivery of the corresponding data (step IV in Figure 1).

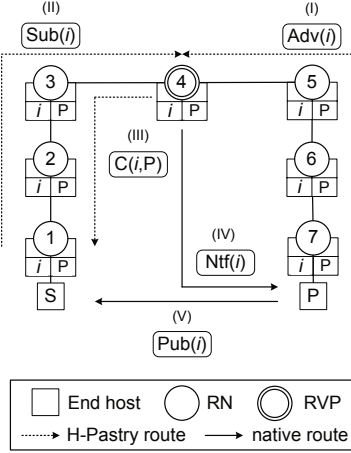


Fig. 1. DHT-NRS Example; the (i, P) records are stored at RNs 4, 5, 6 and 7 during step I, and at RNs 1, 2 and 3 during step III

- **Pub:** Upon notification from the RVP, a publisher sends the IO to the subscriber(s) (step V in Figure 1) via an underlying forwarding mechanism.

Caching. Even though DHT-NRS is based on a novel DHT scheme that adapts to the underlying network topology, routing is still based on DHT primitives which are characterized by their logarithmic relation to the number of RNs in the inter-network. Adopting established practice (see Section 3), we employ caching to shorten the resolution paths. More specifically, we cache the route towards the publisher of an IO. During advertisement of an IO, a cache entry is created in all intermediate RNs encountered by the **Adv** message (step I in Figure 1). Moreover, when a **Sub** message reaches an RVP, the RVP actively pushes a copy of the respective IOE to all RNs in the path followed by the subscription (step III in Figure 1). Further subscriptions for an IO trigger **Ntf** messages directly towards the publisher upon a cache hit, reducing the part of the resolution path that follows H-Pastry’s routing. Each cache entry contains a *time to live* (TTL) value that allows for its invalidation and is subject to application level criteria. Cache replacement follows a *Least Recently Used* (LRU) mechanism.

Node Failures. Due to the vital role of the NRS in locating information in ICN, we expect the deployment process to follow a path similar to DNS i.e., rely on highly provisioned RN servers. However, the multitude of DoS attacks against DNS signifies the critical role of system resilience, while the vast size of the expected workload further urges for a design able to cope with node failures. The use of a highly distributed DHT-based solution is a design choice aligned with this desire. Based on the uniform distribution of IO and node IDs in the identifier space, a DHT based solution accomplishes an even distribution of IOEs to nodes, so that a node’s failure has a modest impact to the entire system, proportional

to deployment density. The fault resilience properties of H-Pastry add to the DHT-NRS robustness (see Section 4.1). In contrast, hierarchical designs (e.g., DONA) allow even a single node failure to affect the routing of messages to/from the entire sub-tree rooted at the failed node. Obviously, the impact of a node failure is isolated to that part of the inter-domain topology, but the consequences become more severe for nodes at higher levels. Multiple physical incarnations of resolution nodes ameliorate this problem, but only at an additional cost.

5 Performance Evaluation

In order to investigate the extent to which the DHT-NRS proposed fulfills our requirements, we focus on two major aspects: (i) system load, including memory, signaling and processing overheads, and (ii) routing performance. Our investigation offers a detailed comparison with DONA, with the purpose of revealing, as well as quantifying, the inherent (dis)advantages of both architectures. Particular attention is paid to (i) the effect of the underlying inter-network structure on perceived performance and its relation to the structural characteristics of each architecture and (ii) the effect of the popularity characteristics of content on the effectiveness of DHT-NRS’s caching scheme.

Topologies. Understanding the Internet topology is a crucial factor in designing new inter-domain protocols. Typically, the AS-level topology is described as a multi-tier hierarchy of interconnected ASes, mostly using transit customer-provider links. However, several measurements and studies argue that the Internet topology is evolving into a mesh dominated by multi-homing and peering relationships [3,18,19]. These studies converge on the number of the ASes, which is estimated to be around 35.000. Evaluating a protocol using a realistic topology of 35.000 ASes (and approximately 200.000 annotated links) is obviously a technical challenge on its own. To overcome this constraint, while also preserving memory and processing resources for the evaluation of non trivial DHT sizes, we used inter-domain topologies generated by the algorithm presented in [8]. The size of these topologies is manageable for evaluation and they maintain the same characteristics (i.e., business relationships) as in the measured graphs.

Workload. Our evaluation considers a detailed model of inter-domain Internet traffic. Based on the measurements presented in [18], we generate a mixture of various traffic types (e.g., Web, Video, P2P)¹. The resulting traffic mix does not distinguish between user and control plane traffic i.e., the signaling required to locate and start content transmission, except for DNS which only constitutes 0.17% of all traffic. However, our evaluation focuses on the control plane requests made to the NRS to support user plane traffic. Therefore, we translated the user plane traffic mix into a control plane equivalent. To this end, we derived the actual number of IOs for each traffic type by dividing the corresponding data volume with the median IO size of that traffic type as measured in relevant

¹ We did not use DNS trace data as they reflect the current Internet architecture. For example, they omit requests sent directly to content owners (e.g., HTTP requests).

studies [2,1,4]. We used the median instead of the mean object size, so as to avoid skewing the results due to the long-tail characteristics of some distributions (e.g., the Pareto tail of Web object size distribution). The set of subscription messages is then shaped based on studies for each traffic type which model its popularity characteristics, including its temporal evolution [2,4,13].

5.1 Results

Our evaluation employs a 400 domain topology, which follows a hierarchical model with six levels, but also contains multi-homing and peering links between the domains. We deploy a population of 4400 RNs/RHs uniformly across the domains. The cache size in DHT-NRS is expressed as a proportion of the median number (m) of registration entries per RH in DONA. We choose the median value, since the distribution of state in DONA is considerably skewed, as we show below. Moreover, we examine scenarios with *infinite cache size* (ICS) i.e., no limitation on the available cache size; this reflects the upper limit for the performance of caching in DHT-NRS. In each scenario, we use a workload corresponding to 25 GBs of traffic, resulting in an average² of 2430379 subscription messages for 1032030 IOs. This size limit was imposed by the limitations of the simulation environment. Since end hosts typically reside in access networks i.e., networking domains that have no customer domains and thus do not act as transit domains, all Adv/Sub messages in DHT-NRS (REGISTER/FIND in DONA) are injected into the network from a randomly chosen RN (or RH respectively) residing in an access network. Finally, we considered a single publisher per IO.

Routing. The routing performance of the name-resolution system affects both the latency perceived by the end-users, as well as the traffic load on the network. We express the routing performance of DHT-NRS with the *stretch* metric, defined as the ratio of the number of inter-domain hops required for a Sub message to reach the RVP and the corresponding Ntf message to reach the RN serving the publisher of the desired IO, over the hop count required by an identical FIND message to reach the same target in DONA i.e., the RH that issued the corresponding REGISTRATION.

Figure 2(a) shows the stretch values derived for several scenarios with different cache sizes. Stretch ranges from 2.84 without caching to 1.95 in the ICS scenario. Under current traffic patterns, DONA significantly outperforms DHT-NRS in routing efficiency. This is only possible however because DONA extensively replicates IOEs throughout the hierarchy, thus guarantying the existence of the desired registration on the shortest, policy-compliant path towards the publisher. DONA is superior to DHT-NRS regardless of cache size, since caching is less aggressive than replication, reactively adapting to, and therefore highly depending on, the request patterns. Hence, non popular IOEs are evicted from DHT-NRS caches, or even never requested again, yielding low cache hit ratios. As shown in Figure 2(b), the cache hit ratio is 27.17, 31.75, 37.76 and 53% for


² We used a different workload instance in each experiment to increase randomness.

the 50% m , 100% m , 150% m and ICS scenarios respectively. For each scenario, the corresponding savings in the resolution path lengths are 13, 15.39, 18.24 and 31%, respectively. These values are subject to the popularity characteristics of the workload, and show that with current workload patterns route caching in DHT-NRS cannot exceed a cache hit ratio value of 53%, even with unrealistically large cache sizes, resulting in low routing improvements.


In order to provide an insight of the potential benefits of caching in scenarios where item popularity does allow higher cache hit ratios, we also consider the stretch value of resolution paths in which the Sub messages have hit a cache. As shown in Figure 2(a), stretch ranges from 1.34 in the hypothetical, best case scenario (ICS) to 1.67 in the 50% m scenario. These results show that caching achieves a substantial reduction of resolution path lengths in the case of popular IOs, which are still however at least 34% longer compared to DONA. This is due to the indirection mechanism in DHT-NRS, as the RVP may not reside in the shortest path connecting a subscriber and a publisher. However, in cases of multiple publishers per IO (e.g., replication points), we expect the routing properties of H-Pastry and the caching mechanism to further shorten resolution paths by selecting the closest publisher.

State. In our evaluation, we use the term “state” to refer to the IOEs maintained at each node of the name-resolution system. The state size is related to the total number of IOs and determines the amount of resources required to support the operation of the architecture *wrt* memory and lookup processing load. Figure 2(c) shows the cumulative distribution function of the state size for both DHT-NRS and DONA. Note that the x axis is in log scale. The difference between the two architectures is significant, with the size of the state maintained by RN nodes in DHT-NRS being considerably lower than the corresponding load imposed on DONA’s RHs. More importantly, we see that DHT-NRS achieves a more uniform distribution of state across the participating nodes compared to DONA. For instance, 95% of the RNs in DHT-NRS maintain less than 550 IOEs (0% m scenario), while 95% of the RHs in DONA maintain up to almost 33000 IOEs. The highly skewed state distribution in DONA is due to the accumulation of registrations at higher levels of the inter-domain hierarchy, and poses a significant challenge for network operators, who would have to resort to dense deployments of RHs in order to cope with the associated resource requirements.


Moreover, it is important to point out the relation of the observed state distribution skewness to the structure of the inter-domain topology. In the considered topologies, slightly less than 50% of access networks reside at level 2, meaning that a major part of the registrations is stored only at the first two levels of the hierarchy. This is demonstrated in Figure 2(d) which shows the average state size per node at each level of the hierarchy. Note that the y axis is presented in log-scale. This figure suggests that the inter-domain topology structure causes the concentration of excessive state in DONA at the top-most levels of the hierarchy. On the one hand, these top-level domains face disproportionate overhead compared to lower level domains, incurring the corresponding CAPEX/OPEX overheads. On the other hand, however, it also suggests that the scalability




(a) Inter-domain stretch of DHT-NRS




(b) Cache hit ratio




(c) Cumulative state overhead



(d) Distribution of state across hierarchy levels



(e) Cumulative lookup overhead



(f) Distribution of lookup overhead across hierarchy levels

Fig. 2. Performance evaluation results: DHT-NRS vs. DONA

challenge in DONA is mostly concentrated in a limited sub-area of the inter-domain topology, making cloud-based solutions a compelling option.

Figures 2(c) and 2(d) also show the amount and distribution of state required to avoid cache replacement (i.e., the ICS scenario). This hypothetical scenario requires a considerable amount of memory which is highly unlikely to be available in practice (see Section 2). However, even in this scenario, almost 8% of RHs in DONA have higher memory requirements than any RN in DHT-NRS. Moreover,

we see that in DHT-NRS state is distributed more evenly across the hierarchy levels, with the exception of level 1 RNs in the ICS scenario where the routing scheme causes a concentration of cached IOEs (see also the next section).

Processing Overhead. The state size per node determines the processing overhead per IOE lookup at the RVP and the overall processing overhead is determined by the actual number of lookups performed by a node. For each architecture, we define the *lookup overhead* (LO) metric as the sum of the total number of **Sub**/**Find** messages forwarded by each RN/RH respectively, and the total number of messages *terminating* at each node i.e., triggering communication with the publisher. Figure 2(e) shows the cumulative distribution function of LO for both architectures. A major fraction of DONA RHs is subject to less LO than RNs in DHT-NRS, due to the considerably worse routing performance of DHT-NRS that results in **Sub** messages traversing longer networking distances and thus consuming resources at more intermediate RNs. However, a closer look at the LO distribution across the inter-domain hierarchy reveals a disproportionate overhead for the top-level domains in DONA (Figure 2(f)). Evidently, the fact that a major part of the access domains resides at level 2 of the hierarchy results in a corresponding resolution overhead for the top-level domains, which again calls for the use of considerable processing (as well as memory) resources.

Interestingly, Figure 2(f) also shows that RNs at the top-most domains of the hierarchy in DHT-NRS are subject to considerably higher processing overhead compared to lower level RNs. The additional overhead is due to the forwarding of **Sub** and **Ntf** messages (ID ownership is evenly distributed across the RNs) and is attributed to the structure of the inter-domain topology and the design of H-Pastry: to adapt to the inter-domain hierarchy structure, H-Pastry causes the top level domains to be included in the first non-local ring of each of the access networks at level 2. Therefore, paths towards non-local RVPs are most likely to pass through these domains. As a substantial fraction of access networks resides at level 2, a proportional number of IOs are served from publishers at these domains, resulting in a significant part of the overall **Ntf** messages reaching RNs at these domains via the top-level domains. Similarly, H-Pastry causes the **Sub** messages originating from level 2 access domains to first traverse the top-level domains before taking a downhill direction towards the RVP.

Advertisement/Registration Overhead. In addition to subscriptions, signaling overhead is also generated by IO registrations. We characterize this overhead as the total number of single inter-domain hop transmissions required for the registration of a single IO to complete. Our measurements show that DHT-NRS requires on average 6.34 inter-domain transmissions per IO in the 0%*m* scenario, while DONA requires on average 35.56 transmissions. These numbers indicate an excessive inter-domain traffic load in the case of DONA, attributed to the (limited) flooding method used to disseminate registration messages to the upper levels of the inter-domain hierarchy (as well as to peering domains). Multihoming plays an important role in this as it results in registration messages being transmitted to multiple domains at higher levels. In the employed

topologies 56.75% of all domains are multihomed, with 2.4 providers on average. Unlike DONA, by carefully partitioning the identifier space DHT-NRS forms a structured overlay that only requires the targeted routing of Adv messages.

6 Conclusions

In this paper we investigated name-resolution in the context of ICNs, focusing on the emerging tradeoff between routing state overhead and routing efficiency. Motivated by the significant scalability characteristics of DHTs, but also concerned about their routing inefficiency, we engaged in the design of a name resolution system based on an enhanced DHT scheme, aiming to improve routing performance. Then we engaged in a detailed performance evaluation and comparison against DONA, an alternative that yields efficient routing performance but raises significant scalability concerns. Our findings reveal the effect of the inter-domain topology structure and traffic patterns. With current traffic patterns, route caching in DHT-NRS cannot compete with the extensive replication of routing information in DONA, yielding stretch values ranging from 1.95 to 2.84. However, this comes at the cost of a heavily skewed distribution of memory and processing overhead, as well as significant signaling overhead for the registration of content in DONA. The replication mechanism of DONA along with the structural characteristics of the inter-domain topology result in disproportionate resource requirements for a limited area of the internetwork, at the top-most level domains in the hierarchy, indicating the need for large-scale centralized solutions (e.g. cloud computing).

References

1. Bellissimo, A., Levine, B.N., Shenoy, P.: Exploring the use of BitTorrent as the basis for a large trace repository. Tech. rep., University of Massachusetts Amherst (2004)
2. Busari, M., Williamson, C.: ProWGen: a synthetic workload generation tool for simulation evaluation of web proxy caches. *Computer Networks* 38(6), 779–794 (2002)
3. CAIDA (2011), <http://www.caida.org>
4. Cheng, X., Dale, C., Liu, J.: Understanding the Characteristics of Internet Short Video Sharing: YouTube as a Case Study. CoRR abs/0707.3670 (2007)
5. Choi, J., Han, J., Cho, E., Kim, H., Kwon, T., Choi, Y.: Performance comparison of content-oriented networking alternatives: A tree versus a distributed hash table. In: Proc. of the IEEE 34th Conference on Local Computer Networks (LCN), pp. 253–256 (2009)
6. Cisco: Cisco Visual Networking Index 2010-2015 (June 2011)
7. D’Ambrosio, M., Dannewitz, C., Karl, H., Vercellone, V.: MDHT: a hierarchical name resolution service for information-centric networks. In: Proc. of the 2011 ACM SIGCOMM Workshop on ICN, pp. 7–12. ACM, New York (2011)
8. Dimitropoulos, X., Krioukov, D., Vahdat, A., Riley, G.: Graph annotations in modeling complex network topologies. *ACM Transactions on Modeling and Computer Simulation* 19, 17:1–17:29(2009)

9. Fotiou, N., Katsaros, K.V., Vasilakos, X., Tsilopoulos, C., Ververidis, C.N., Xylomenos, G., Polyzos, G.C.: H-Pastry: An adaptive multi-level overlay inter-network. Tech. Rep. 2011-MMLAB-TR-002, Athens University of Economics and Business (2011), <http://mm.aueb.gr/technicalreports/2011-MMLAB-TR-003.pdf>
10. Ganesan, P., Gummadi, K., Garcia-Molina, H.: Canon in G Major: Designing DHTs with Hierarchical Structure. In: Proc. of the 2004 ICDCS, pp. 263–272 (2004)
11. Ghodsi, A., Koponen, T., Rajahalme, J., Sarolahti, P., Shenker, S.: Naming in content-oriented architectures. In: Proc. of the ACM SIGCOMM ICN Workshop, New York, NY, USA, pp. 1–6 (2011)
12. Google: We knew the web was big (July 2008), <http://googleblog.blogspot.com/2008/07/we-knew-web-was-big.html>
13. Guo, L., Chen, S., Xiao, Z., Tan, E., Ding, X., Zhang, X.: A performance study of BitTorrent-like peer-to-peer systems. IEEE Journal on Selected Areas in Communication 25(1), 155–169 (2007)
14. ICANN: ICANN Approves Historic Change to Internet’s Domain Name System (2011), <http://www.icann.org/en/announcements/announcement-20jun11-en.htm>
15. Jacobson, V., Smetters, D.K., Thornton, J.D., Plass, M.F., Briggs, N.H., Braynard, R.L.: Networking named content. In: Proc. of the 2009 ACM CoNEXT, pp. 1–12. ACM, New York (2009)
16. Rajahalme, J., Särälä, M., Visala, K., Riihijärvi, J.: On name-based inter-domain routing. Computer Networks 55, 975–986 (2011)
17. Koponen, T., Chawla, M., Chun, B.-G., Ermolinskiy, A., Kim, K.H., Shenker, S., Stoica, I.: A data-oriented (and beyond) network architecture. In: Proc. of the 2007 ACM SIGCOMM, pp. 181–192. ACM, New York (2007)
18. Labovitz, C., Iekel-Johnson, S., McPherson, D., Oberheide, J., Jahanian, F.: Internet inter-domain traffic. In: Proc. of the 2010 ACM SIGCOMM, pp. 75–86. ACM, New York (2010)
19. Oliveira, R., Pei, D., Willinger, W., Zhang, B., Zhang, L.: The (in)completeness of the observed internet AS-level structure. IEEE/ACM Transactions on Networking 18, 109–122 (2010)
20. Pappas, V., Massey, D., Terzis, A., Zhang, L.: A Comparative Study of the DNS Design with DHT-Based Alternatives. In: Proc. of the 2006 IEEE INFOCOM, pp. 1–13 (2006)
21. Ramasubramanian, V., Sirer, E.G.: The Design and Implementation of a Next Generation Name Service for the Internet. In: Proc. of the 2004 ACM SIGCOMM, pp. 331–342. ACM, New York (2004)
22. Rowstron, A., Druschel, P.: Pastry: Scalable, Decentralized Object Location, and Routing for Large-Scale Peer-to-Peer Systems. In: Guerraoui, R. (ed.) Middleware 2001. LNCS, vol. 2218, pp. 329–350. Springer, Heidelberg (2001)
23. Seetharaman, S., Ammar, M.: Inter-domain policy violations in multi-hop overlay routes: Analysis and mitigation. Computer Networks 53, 60–80 (2009)
24. Stoica, I., Morris, R., Liben-Nowell, D., Karger, D.R., Kaashoek, M.F., Dabek, F., Balakrishnan, H.: Chord: a scalable peer-to-peer lookup protocol for internet applications. IEEE/ACM Transactions on Networking 11(1), 17–32 (2003)
25. Walfish, M., Balakrishnan, H., Shenker, S.: Untangling the web from DNS. In: Proc. of the 2004 USENIX NSDI, p. 17. USENIX Association, Berkeley (2004)