

Cloud computing for global name-resolution in information-centric networks

Xenofon Vasilakos, Konstantinos Katsaros and George Xylomenos
 Mobile Multimedia Laboratory, Department of Informatics
 Athens University of Economics and Business,
 Patision 76, Athens 104 34, Greece
 E-mail: xvas@aueb.gr, ntimos@aueb.gr and xgeorge@aueb.gr

Abstract—Information-Centric Networking (ICN) is a novel paradigm for future Internet architectures. It exploits the current trend in Internet usage which mostly involves information dissemination. ICN architectures based on the publish/subscribe model use names for information in order to route requests and data, as well as to facilitate in-network caching, anycasting and multicasting for efficient content delivery. However, the number of named information objects is expected to be huge in the future Internet, raising serious concerns with respect to a global-scale deployment of ICN. Routing and forwarding will require vast amounts of state, which pushes storage, maintenance and processing demands to the limit. In this paper we discuss the feasibility of deploying the Data Oriented Networking Architecture (DONA) by leveraging cloud computing facilities. We identify the exact scalability concerns for DONA based on simulations over a realistic model of the current Internet topology and find that registrations for information objects lead to a state explosion. For this reason, we then discuss how cloud facilities can assist DONA deployment, focusing on various options for deploying DONA in the cloud and their suitability for different areas of the inter-network.

I. INTRODUCTION

The *Information-Centric Networking* (ICN) paradigm was proposed during the past few years in an effort to reflect the current user needs for information on the Internet. A considerable amount of research is devoted to the design of alternative Internet architectures [1], [2], [3], [4], [5] along the lines of ICN, as today's Internet usage is increasingly turning towards disseminating information rather than offering pairwise communication between end hosts. Information in ICN is organized in named *Information Objects* (IOs)¹ which have a central role in the overall functionality. Routing and forwarding decisions are guided by the identity of IOs, rather than location addresses identifying end-hosts. Moreover, ICN adopts Publish/Subscribe as its service model, which results in inherently *decoupling information in space and time*. This means that requests resolve to the desired IOs rather than their hosting location, and that the model is asynchronous as entities requesting (subscribers) or offering (publishers) information do not have to be concurrently online. This decoupling is ideal for supporting mobile entities, i.e. applications running on mobile devices or mobile (software) agents. Apart from mobility support, ICN facilitates efficient communication models such

as *multicast* or *anycast*, while it inherently benefits traditional acceleration mechanisms such as caching.

A major challenge faced by all ICN efforts relates to *scalability* in global scale deployment. In ICN, IOs are named in order to guide routing and forwarding decisions in the network through a name resolution process i.e., IO names are used to determine the paths followed by subscription requests and the data published in response. However, the size of the namespace used for IOs is expected to be enormous. Considering that (a) the current number of unique web pages indexed by Google is greater than 1 trillion [6] and that (b) billions [7] of devices (mobile phones, sensors, home appliances etc.) will be offering additional content to future networks, one should expect that any name resolution approach will have to handle unique IOs in the order of 10^{13} . Some studies raise this estimate even further to 10^{15} [8]. Even though hierarchical naming schemes can reduce this number to 10^8 [9] and aggregation techniques designed for ICN [10] can mitigate the problem to some extent, *scalability concerns remain prevalent*.

In this context, several approaches have been proposed for name-resolution, focusing on handling the huge load for name resolution and routing state. The *Data-Oriented Network Architecture* (DONA) [5] is a name-resolution architecture designed to *perfectly adapt* to the underlying topology and policies, at the cost of accumulating vast amounts of routing state going upwards in the hierarchy of *Autonomous Systems* (ASes), with Tier-1 providers managing the entire namespace, in order to guarantee the global resolution of available information. Despite its susceptibility to excessive routing state, DONA remains an important approach, as its adaptation to the network topology is unique among ICN architectures, influencing other approaches such as [3]. The design of DONA involves a network of *Resolution Handlers* (RHs) responsible for maintaining information about the availability and the location of IOs in the network. RHs are deployed following the inter-domain topology architecture, with routing state being forwarded from customer to provider ASes and between peering ASes. Intuitively, the aggregation of routing information at the higher levels of the AS-level topology raises concerns on the ability of DONA to scale, especially when considering the huge size of the routing name space. In order to address these concerns, other approaches have taken advantage of the scalability properties of distributed hash tables (DHTs) [8], [11], [10], [12]. However, it has been

¹For brevity, we use the term IO to refer to both actual pieces of information and services.

shown that a significant penalty is paid for this scalability, expressed in highly stretched name-resolution paths, often also violating the underlying routing policies, even in the presence of extensive caching [12]. This is in direct contrast to DONA (and similar approaches [13]) where name-resolution paths always follow shortest-path, BGP compliant routes.

In this paper we take a detailed look at the scalability concerns for global name-resolution in ICNs, from the perspective of advances in cloud computing. Motivated by (a) the ability of the cloud to aggregate and offer on demand large volumes of resources, and (b) the routing efficiency of DONA, we revisit the related scalability concerns and attempt to assess the ability of cloud computing technology to address them. Our target is to investigate the extent to which cloud computing can assist in the realization of the DONA architecture, thus enabling an Internet scale, ICN name-resolution system with efficient routing properties. In this effort, we first carefully explore these scalability concerns. To this end, it is important to note that the exact magnitude of the routing state accumulated at each AS is primarily determined by the structure of the inter-domain topology and the routing policies established. Hence, we engage in the study of the impact of the inter-domain topology structure on the distribution of routing state in DONA. Our purpose is to offer a precise look on the resource requirements at each level of this hierarchy. Our observations then constitute the necessary input for assessing the suitability of the cloud paradigm for the support of the desired functionality and the different ways that cloud facilities can be exploited by the DONA architecture.

The paper is organized as follows. In Section II we provide the reader with background on the Internet topology landscape and the core design decisions in DONA. We proceed with our evaluation results of DONA and a discussion triggered by our findings on the impact of the Internet topology to DONA's scalability in Section III. Then, we discuss the potential of globally deploying DONA by using the services and capabilities offered by cloud computing, before discussing related work in Section V and concluding in Section VI.

II. BACKGROUND

In this section we first provide an overview of the AS-level routing landscape on the Internet, describing the structural characteristics of the inter-domain graph and the routing relationships between ASes. Then, we describe DONA, focusing on the routing state that it establishes across the topology.

A. Inter-domain routing

Inter-domain routing on the Internet follows a hierarchical structure, imposed by the routing relationships established through business agreements between ASes. In general, ASes at a certain level of the hierarchy act as providers of transit traffic services to lower level customers ASes. As a result, a hierarchy of *provider-to-customer* relationships is formed recursively, in which customers are charged by their direct providers for connectivity and traffic transit services; direct providers are themselves customers of other provider ASes and so forth, up until the highest level of the hierarchy composed

by *Tier-1* providers. Tier-1 ASes communicate either directly with each other or via other Tier-1 ASes, thus guaranteeing global connectivity and transit services between any two ASes on the Internet. Given that billing between providers and customers is mostly subject to the amount of transit traffic, ASes tend to establish connectivity with multiple providers for load and cost balancing reasons, as well as for redundancy; this practice is known as *multihoming*.

Apart from provider-customer relationships, ASes also establish *peering* relationships for the direct exchange of traffic.² The main incentive for establishing peering agreements is to avoid using the expensive transit services offered by providers. Charging between peering ASes is subject to the exact terms of the peering agreements: for symmetric traffic exchange patterns it may be zero, while for asymmetric ones it may depend on the balance of traffic exchanged between the peers.

Through the years, the Internet topology has dramatically evolved in scale and structure. *Multihoming* and the establishment of *multiple peering relationships* between ASes has resulted in an AS-level topology graph that is far from purely hierarchical, making the classification of ASes across the topology difficult. In this paper, we adopt the methodology proposed in [15] and also used in [14], according to which the set of ASes is classified into four tiers based on the number of their downstream customers (this number is called the *customer cone* or simply the *cone* of an AS). The four tiers are: (i) *Stub* networks, i.e. networks that have no more than 4 customer networks (all access networks belong to this category), (ii) *Small ISPs*, i.e. ASes that have a cone size between 5 and 50, (iii) *Tier-1* ASes, i.e. ASes at the highest level of the hierarchy which do not act as customers for another AS, (iv) *Large ISPs* i.e., the remaining ASes which have a larger cone than small ISPs but are not tier-1. The cone size limits used in this classification were derived by inspecting the shape of the cumulative distribution function (CDF) of cone sizes across the entire topology as in [15], with the results shown in Fig. 1 based on the latest CAIDA traces³; Table I shows the number of ASes in each class, based on these traces [16]. It is clear that the vast majority of ASes belong to the *Stub* category.

TABLE I
AS CLASSIFICATION.

Type	ASes	%
Tier-1	33	0.09
Large ISP	825	2.24
Small ISP	1717	4.66
Stub	34303	93.02
Total	36878	100

B. Data-Oriented Name Architecture

²We omit relationships between ASes which belong to the same organization, as these have been reported to be very rare [14].

³We do not show the CDF for the entire range of cone size values (up to 35753 ASes) for visibility reasons.




Fig. 1. Distribution of AS cone sizes.

The *Data Oriented Network Architecture* (DONA) [5] employs a two-level hierarchical name-space for *Information Objects* (IOs). Each IO is associated with a *principal*. The identifier of an IO is composed by the cryptographic hash of its principal's public key and a label unique to that principal, thus forming a two-level self-certifying unique name.

Name resolution is performed by *Resolution Handlers* (RHs) with at least one logical RH at each AS, possibly realized by one or more physical incarnations. RHs interconnect along the lines of the hierarchical interconnection of different ASes in the Internet, thus forming a corresponding RH hierarchy. Peering agreements between ASes are also included in DONA's resolution layer through the interconnection of the respective RHs. The role of an RH is to communicate and maintain registrations for IOs, as well as to propagate requests for name resolution to the corresponding IOs.

Publishers of IOs issue a REGISTER message towards their local RH for each IO they wish to make available to the network. The RH propagates the REGISTER message upwards in the inter-AS hierarchy as well as to RHs residing at peering ASes. Propagation stops at the top-most layer of the hierarchy i.e., Tier-1 RHs. The reception of a REGISTER message informs an RH of the availability of the corresponding object. The state kept at each RH for a received registration is a mapping between the name of advertised information and the IP address of the RH that forwarded the REGISTER message to the current RH. Therefore, resolution state consists of (name, next hop RH) pairs. Resolution requests are realized by FIND messages which are similarly forwarded upwards in the hierarchy, until some RH along the resolution path matches the request to a registration entry. Upon a match, FIND messages are propagated through the reverse path of the corresponding REGISTER messages until they reach the IO. Data transfer is triggered by the delivery of the FIND message to the publisher of the corresponding REGISTRATION message and takes place through regular IP routing/forwarding.

DONA also supports the extension of RH functionality with IO caching. In this case an RH can change the source IP address of an incoming FIND message with their own,

before further forwarding the message. As a consequence, the corresponding IO will be delivered to the RH which can cache them before forwarding them to the original requestor.

Since DONA's RH layer closely follows the inter-AS routing graph, name resolution takes place along the corresponding BGP paths. However, it is crucial to note that this is only enabled by the creation and maintenance of a registration entry at each RH traversed by a REGISTRATION message. This results in multiple RHs in the overall architecture maintaining name-resolution information for each IO, leading to an increased total amount of state that has to be handled, especially at the topmost layers. We make an attempt to quantify this overhead in the following section.

III. PERFORMANCE EVALUATION

In this section we first present our simulation setup and methodology and then proceed with the presentation and discussion of our findings.

A. Simulation setup and methodology

For the purposes of our study, we developed a custom, simplified DONA simulator, which simulates the registration of IOs along with the name resolution functionality. We employed the AS-level Internet graph inferred by the CAIDA BGP traces [16] in an effort to gain a realistic view of the distribution of routing state in DONA. We considered a simple simulation setup that comprises only one logical RH per AS, with REGISTER messages generated only by stub networks. The latter is commonly the case on the Internet, where content providers usually reside at the lowest level of the inter-domain hierarchy. Unless otherwise stated, each stub AS generates *one* distinct registration message that is propagated following the established customer-provider and peering links. In cases of multihoming, RHs forward registrations to the RH of their default provider AS which is randomly selected. Incoming REGISTER messages from RHs at peering domains are not further propagated upwards in the RH hierarchy. Based on this setup we investigate the process of routing state aggregation at each AS, measuring the number of registration entries maintained by each RH once all registration messages have been propagated throughout the network.

B. State distribution

With DONA multiple registration entries must be maintained across the hierarchy of ASes for each IO made available to the network. In the following, we express the size of the aggregated state at each AS as a percentage of the overall number of unique registration entries across the network of RHs, in an effort to quantify the degree to which per IO state is replicated throughout the network. We define the *multiplier* m , as the ratio of the total number of registration entries maintained throughout the inter-domain topology (RE_{total}) to the total number of unique information objects (RE_{unique}):

$$m = \frac{RE_{total}}{RE_{unique}}$$

Our measurements provide a value of $m = 210.70$ meaning that DONA state is heavily reproduced across the internetwork, as a result of both the overall architecture design and the structure of the inter-domain graph. This is the price paid in DONA for establishing a BGP compliant routing scheme for name resolution i.e., registration information must be heavily reproduced on the higher tiers of the topology in order for name resolution paths to fit their BGP counterparts.

Based on this observation, we next examine how this state is distributed across the network. Figure 2 shows the cumulative distribution function (CDF) of the size of the state maintained by ASes across the inter-domain topology. The x axis is in log scale and expresses this size as a percentage of the total number of IO entries in the network. The *All* curve refers to the state size across all tiers. The distribution is heavily skewed, both across and within the tiers of the inter-domain topology. More than 89% of all ASes are only burdened with one registration entry (their own), with the remainder (3809 ASes in total) being disproportionately loaded with state, even reaching the entire available state in the network. This is a direct consequence of the structure of the inter-domain topology. The vast majority of ASes belong to the Stub tier (93.02%, see Table I), out of which only a small fraction (3.61% or 1237 ASes) has more than one (and up to 4) customers that contribute with registration entries.

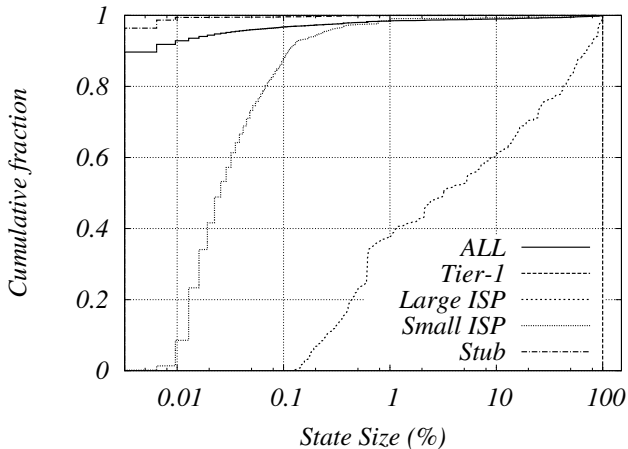


Fig. 2. Distribution of state size across the AS-level topology. Note that the state size is by default 100% for all Tier-1 ASes.

The heavily skewed distribution of state overhead becomes clearer by inspecting the average and median values of state size values per tier in Tables II and III. In Table II the *Average*, *Median* and *99% Confidence interval* columns present the corresponding statistical metrics for the number of registration entries per AS in each tier. Table III presents these values as a percentage of the total number of IOs in the network.

On closer inspection, each tier is defined by the cone size of the participating ASes, which in turn determines the size of the accumulated state i.e., all (distinct) registrations created by the members of a cone are eventually stored at the AS for which the cone is defined. This state is further augmented by the entries received from peering ASes. This

TABLE II
STATE SIZE ACROSS THE AS-LEVEL TOPOLOGY: ABSOLUTE VALUES

Type	Average	Median	99% Confidence Interval
Tier-1	31193.00	31193	-
Large ISP	6143.24	998	798.37
Small ISP	142.14	8	99.09
Stub	6.73	1	4.92

TABLE III
STATE SIZE ACROSS THE AS-LEVEL TOPOLOGY: RELATIVE VALUES (%)

Type	Average	Median	99% Confidence Interval
Tier-1	100.00	100.00	-
Large ISP	19.69	3.199	0.02559
Small ISP	0.46	0.025	0.00318
Stub	0.02	0.003	0.00016

is clearly demonstrated in Figure 3 which shows the relation between the cone size and the state size, this time expressed in absolute values (both axes are in log scale). The large concentration of points in the diagonal denotes that the amount of state accumulated by DONA RHs is mainly affected by the size of the cone of the current AS. At the same time, the scattered points above the diagonal show the impact of peering relationships on state size. This is most evident at the upper right part of the graph which depicts the full mesh connectivity at Tier-1. It is important however to note the existence of several cases where ASes with a relatively small cone size maintain state of disproportionate size (upper left area). In these cases, the establishment of peering relationships with ASes at higher tiers (and therefore large cone sizes) results in the reception of large amounts of routing state. Therefore, peering with higher tier ASes becomes quite unfavorable in DONA. Finally, we note that points below the diagonal correspond to ASes whose cone also contains non Stub ASes, which do not generate registrations.




Fig. 3. Relation between the cone size of ASes and the accumulated registration state.

C. Resource requirements

As demonstrated by these results, DONA state distribution across the RH network is heavily concentrated in some ASes,

especially those at tier-1, due to DONA's design and the structure of the inter-AS topology graph. It is therefore anticipated that a realization of DONA would present various technical and financial challenges and requirements for the different types of ASes across the Internet. At a technical level, DONA operation poses requirements in the following domains:

- **Memory:** this refers to both secondary storage and RAM for the maintenance of the registration entries. Efficient data structures are further required to facilitate indexing this information.
- **Processing:** this relates to the processing of FIND messages, including the processing power required for lookup operation on the indexed data structures.
- **Bandwidth:** this relates to the networking interface capabilities of each RH for REGISTER and FIND messages, which obviously affect RH throughput.

A set of back of the envelope calculations related to these aspects were provided in [5] in an effort to assess the applicability of the proposed architecture. In terms of memory requirements it was concluded that roughly 500 PCs with 8GB of RAM each would be required to allow all registration entries to be maintained in RAM, thus avoiding the delays of secondary storage access. The same ensemble of PCs was calculated to be able to handle an aggregate load of 1Tbps (due to incoming FIND messages). As registrations will inevitably also be maintained in secondary storage, it was concluded that a minimum of 4 TB is required at Tier-1.

Though providing a clearer view of the resource requirements for the support of DONA, we must note that these calculations are based on information related to the current use of the Internet. As the Internet rapidly evolves, we expect these requirements to significantly vary in the future. Apart from the constant increase of the available content in the Internet [7], the advent of the *Internet of Things* (IoT) is expected to further bring a multitude of IO that should be represented in a potential DONA realization. This will obviously increase the burden for the infrastructure, thus making these projections very conservative.

IV. THE POTENTIAL ROLE OF THE CLOUD

It is more than obvious that the requirements imposed by DONA can only be met by exploiting cloud techniques. Modern cloud data centers scale to even 100,000 servers, each one with enough resources to host several tens of virtual machines. At the same time, storage costs have rapidly decreased allowing modern data centers to store huge volumes of data. In this context, the aforementioned, conservatively calculated, DONA resource requirements could be well supported by current cloud data center capacities. However, this is only a rough indication of the technical feasibility of a DONA deployment. A series of issues emerge when attempting to gain a more concrete view on the realization of such an architecture in the context of cloud computing.

A. The impact of AS topology

As demonstrated by the results presented in section III-B, the distribution of state overhead, and the associated resource

requirements, are very skewed across and within tiers. With the notable exception of tier-1 ASes, a diverse set of requirements exists for the various ASes. This results in a discrepancy in the size of the corresponding investment for the deployment of RHs in private or public cloud data centers and, hence, an imbalance in the CAPEX/OPEX⁴ of ASes for the realization of DONA. As seen in section III-B, this imbalance is in principle determined by the structure of the AS level graph and in particular the cone size, as well as the number of peering links, of each AS. Revisiting these results, we see that for the vast majority of ASes, DONA deployment would not require a very large investment. In this respect, modern cloud data center capabilities seem to primarily provide the required solution for tier-1 ASes. However, particular attention must be paid to the effect of peering links. As we saw in Figure 3, several ASes residing at lower levels of the hierarchy have significant resource requirements for the support of DONA, due to the existence of peering links with higher level ASes. Furthermore, according to [14], a significant fraction of the established peering links is not present in public BGP data, especially in lower tier ASes, where this fraction reaches up to 90% for large content providers. This means that the cloud will be critical in enabling DONA even at lower hierarchy levels. At the same time, it further shows that the corresponding CAPEX/OPEX is far from proportional to the size of an AS (as expressed by the customer cone size).

In a similar vein, it is important to note the potential imbalance of resource requirements when considering the distribution of content providers (publishers) in the inter-network. It has been recently shown that no more than 150 large content providers are responsible for producing 50% (or perhaps even more) of the Internet traffic (e.g., YouTube [17]). Although we acknowledge that it is not straightforward to map measured traffic volumes to the number of unique information items published in the Internet, yet we can argue that it is indicative of the uneven distribution of content throughout the Internet. We plan to engage in a detailed quantification of this imbalance and the corresponding DONA dimensioning issues in our future work.

B. The role of virtualization

Current experience from large scale indexing and search services shows that cloud services can satisfy high requirements for computing and storage resources. However, the role of the cloud in the realization of an architecture like DONA cannot be restricted to the provision of large volumes of resources. By enabling multi-tenancy, virtualization in cloud computing allows a significant reduction of the associated CAPEX/OPEX of ASes. This means that the economic benefits of cloud computing could actually ease the realization of future Internet architectures like DONA. At the same time, virtualization further enables the *efficient up/down scaling* of the allocated resources in cases of significant workload variations. Such

⁴Capital Expenditure (CAPEX) includes business expenditure on assets such as upgrading existing facilities. Operational Expenditure (OPEX) covers the expenses for running a service or the depreciation of capital (facilities, etc.) with time.

variations can be caused by the establishment of an additional peering link between ASes, the interconnection of a new content provider or the emergence of particularly popular IOs (flash crowds).

Nevertheless, deploying DONA in the cloud comes with the cost of *routing stretch*. Recall that DONA is designed and tailored to the hierarchical inter-network substrate, and that RHs propagate resolution messages along the lines of established BGP-compliant routes. As a result, resolution in DONA does not suffer from stretch, an advantage which can cease to be valid when (partially or fully) offloading DONA's functionality to public cloud facilities such as Amazon or Rackspace, to name a few [18]. Pushing incoming requests from lower level ASes to public data centers would unavoidably result in added stretch during resolution. There is therefore a need to consider different modes of deploying DONA in the cloud, in addition to the default option of relying on public cloud providers.

One alternative to public cloud services is to invest in private facilities for data centers. While public cloud services seem to be a reasonable option for small scale ASes, private facilities seem more appropriate for large ASes and Tier-1 providers which, unlike small ASes, have the ability to invest in data centers so as to avoid the costs implied by unnecessary stretch. Also, note that such ASes need to deal with massive amounts of state and requests for resolution. Nonetheless, even small scale ASes may face huge demands after deploying DONA. As we discussed in section III-B, some small scale ASes are peers with either ASes residing higher in the hierarchy or big content providers such as YouTube. A hybrid cloud approach could work as a remedy for such ASes, based on small scale investments on private data center facilities to directly resolve (or accelerate resolution through replication or caching) "privileged" IOs. By *privileged*, we refer to any sort of content or service that has to be "specially treated" such as content from big content providers who are willing to pay for fast (i.e., stretch-less) resolution, or popular information for which caching data locally can reduce the cost of data transport through ISPs. For the remaining IOs, resolution can be offloaded to public cloud providers.

C. Enabling content distribution services through caching

On an another front, cloud capabilities are seen as an excellent means to facilitate DONA's caching functionality. By offering large volumes of storage, the cloud can therefore further contribute to the enhancement of the overall DONA service model, enabling network operators to couple name resolution with content distribution services. Taking a step further, these services could be enhanced by considering other features such as streaming capabilities, content pre-fetching mechanisms, etc. At the same time, the already established communication paths between different ASes, could be used for the gradual formation of inter-domain content delivery platform. This is in alignment with the currently observed market trends where network operators are getting involved in the *Content Delivery Network* (CDN) market [19].

V. RELATED WORK

There has been recently a lot of effort put on multi-level DHTs [12], [8], [11] for name resolution approaches. The motivation for such solutions is the scalability of hierarchical schemes and the load distribution properties of DHTs, with [10] also using Bloom-filter based aggregation techniques that fit the context of ICN. A common point for such solutions is the organization of the hierarchy based on inter-domain relations, enhanced by DHT influenced characteristics with respect to the routing of requests for resolution. For instance, [12] leverages the inherent ability of the Pastry DHT to respect locality through the use of a proximity metric, resulting in short routes and route convergence. However, all DHT-inspired solutions suffer from violating BGP routing policies as overlay hops go back and forth across the topological hierarchy with solutions such as [12] trying to mitigate that particular issue. Moreover, DHT-like routing increases hop stretch for resolution and, depending on data path formation, even for data routing as well. DONA on the other hand does not scale well, nor does it distribute the load equally among its resources across domain, as we saw, but it is BGP-compliant and it adapts to the current topological structure, which implies no stretch for resolution and data exchange. In practice however, some modes of deploying DONA in the cloud, as discussed in the previous section, will lead to variable levels of routing stretch for resolution messages.

VI. CONCLUSION

Routing and forwarding decisions within the context of Information-Centric Networking are dictated by the names of information objects and services made available to the network. This paper investigated DONA's feasibility as a future Internet name resolution architecture, by identifying and quantifying anticipated scalability concerns. Our simulation findings over realistic Internet topologies validate the susceptibility of DONA towards excessive routing state, with registrations for information or services being heavily replicated by a factor of 210.70. Moreover, the vast majority of ASes maintain minimum state, which means that the load for routing state is unevenly distributed across the hierarchy, with some examples even reaching to the size of the entire available state in the network. Nevertheless, we also find that some small-scale ASes maintain state of disproportionate size due to established peering relationships with bigger ASes. This is an important conclusion as peering relationships seem to be the trend in the inter-AS topology, a practice that is strongly unfavorable to DONA.

Based on our observations, we also discussed the feasibility of deploying DONA in a global scale through leveraging virtualization techniques and scalable storage abilities offered by data centers in the cloud. Virtualization of resources (storage, processing power) can deal with flash crowds, e.g. a sudden increase for popular items, or the establishment of peering relationships, by scaling the amount of allocated resources accordingly. We identified three modes of deploying DONA in the cloud, which trade-off the amount of investment on private cloud facilities against the routing stretch brought by reliance

on public cloud facilities. Based on our simulation findings and the properties of DONA, we conclude that investing in private data centers can facilitate tier-1 ASes, as well as ASes which maintain peering links with ASes residing at higher levels or big content providers. Smaller ASes facing high demand but lacking the ability to make large investments in private data centers can either turn to public cloud providers, or choose a hybrid approach that combines a small investment in private facilities along with offloading part of the resolution functionality to the public cloud.

ACKNOWLEDGMENT

The work in this paper was supported by the FP7 ICT project PURSUIT, under contract ICT-2010-257217.

REFERENCES

- [1] G. Xylomenos, X. Vasilakos, C. Tsilopoulos, V. Siris, and G. Polyzos, "Caching and mobility support in a publish-subscribe internet architecture," *IEEE Communications Magazine*, vol. 50, no. 7, pp. 52–58, 2012.
- [2] V. Jacobson, M. Mosko, D. Smetters, and J. Garcia-Luna-Aceves, "Content-centric networking," *Whitepaper, Palo Alto Research Center*, pp. 2–4, 2007.
- [3] W. Chai, N. Wang, I. Psaras, G. Pavlou, C. Wang, G. de Blas, F. Ramon-Salguero, L. Liang, S. Spirou, A. Beben *et al.*, "Curling: Content-ubiquitous resolution and delivery infrastructure for next-generation services," *IEEE Communications Magazine*, vol. 49, no. 3, pp. 112–120, 2011.
- [4] SAIL, *SAIL Project*. [Online]. Available: <http://www.sail-project.eu/>
- [5] T. Koponen, M. Chawla, B.-G. Chun, A. Ermolinskiy, K. H. Kim, S. Shenker, and I. Stoica, "A data-oriented (and beyond) network architecture," in *Proc. of the 2007 ACM SIGCOMM*. New York, NY, USA: ACM, 2007, pp. 181–192.
- [6] Google, *We knew the web was big, July, 2008*. [Online]. Available: <http://googleblog.blogspot.com/2008/07/we-knew-web-was-big.html>
- [7] Cisco, *Cisco Visual Networking Index 2010-2015, June, 2011*.
- [8] M. D'Ambrosio, C. Dannewitz, H. Karl, and V. Vercellone, "MDHT: a hierarchical name resolution service for information-centric networks," in *Proc. of the 2011 ACM SIGCOMM Workshop on Information-centric networking (ICN'11)*. New York, NY, USA: ACM, 2011, pp. 7–12.
- [9] A. Narayanan and D. Oran, "NDN and IP Routing: Can it scale?" *Proposed Information-Centric Networking Research Group (ICNRG), Side meeting at IETF-82, Taipei, November 2011*.
- [10] H. Liu, X. De Foy, and D. Zhang, "A multi-level dht routing framework with aggregation," in *Proc. of the 2012 ACM SIGCOMM Workshop on Information-centric networking (ICN'12)*. ACM, 2012, pp. 43–48.
- [11] J. Rajahalme, M. Särelä, K. Visala, and J. Riihijärvi, "On name-based inter-domain routing," *Computer Networks, Elsevier*, vol. 55, pp. 975–986, March 2011.
- [12] K. V. Katsaros, N. Fotiou, X. Vasilakos, C. N. Ververidis, C. Tsilopoulos, G. Xylomenos, and G. C. Polyzos, "On Inter-domain Name Resolution for Information-Centric Networks," in *Proc. of the 2012 IFIP TC6 International Conference on Networking, 2012*.
- [13] W. K. Chai, N. Wang, I. Psaras, G. Pavlou, C. Wang, G. de Blas, F. Ramon-Salguero, L. Liang, S. Spirou, A. Beben, and E. Hadjoannou, "Curling: Content-ubiquitous resolution and delivery infrastructure for next-generation services," *IEEE Communications Magazine*, vol. 49, no. 3, pp. 112–120, march 2011.
- [14] R. Oliveira, D. Pei, W. Willinger, B. Zhang, and L. Zhang, "The (in)completeness of the observed internet AS-level structure," *IEEE/ACM Transactions on Networking*, vol. 18, pp. 109–122, February 2010.
- [15] —, "Quantifying the completeness of the observed internet as-level structure," UCLA, Tech. Rep. 080026, 2008.
- [16] C. C. A. for Internet Data Analysis, July 2011. [Online]. Available: <http://www.caida.org>
- [17] C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian, "Internet inter-domain traffic," in *Proc. of the 2010 ACM SIGCOMM*. New York, NY, USA: ACM, 2010, pp. 75–86.
- [18] *Top 10 cloud computing providers of 2012*. [Online]. Available: <http://searchcloudcomputing.techtarget.com/photostory/2240149038/Top-10-cloud-providers-of-2012/1/Introduction/>
- [19] J. Wulf, R. Zarnekow, T. Hau, and W. Brenner, "Carrier activities in the cdn market - an exploratory analysis and strategic implications," in *Proc. of the 2010 International Conference on Intelligence in Next Generation Networks (ICIN)*, oct. 2010, pp. 1–6.