VDR: A Virtual Domain-based Routing Scheme for CCN

Jie Li†*, Jiachen Chen*, Mayutan Arumaithurai*, Xingwei Wang† and Xiaoming Fu*
†College of Information Science and Engineering, Northeastern University, China.
Email: jieli.neu.edu@gmail.com, wangxw@mail.neu.edu.cn
*Institute of Computer Science, University of Göttingen, Germany.
Email: {jie.li, jiachen, arumaithurai, fu}@cs.uni-goettingen.de

1. INTRODUCTION

The advent of Content-Centric Networks (CCN [1]) separates the data from its location and provides benefits like multi-source data retrieval, multicast, in-network cache, *etc.* However, the routing in CCN, especially the Forwarding Information Base (FIB) scalability becomes the Achilles' heel since the routing identities are no longer assigned by the network managers and traditional aggregation in IP cannot work efficiently.

Solutions like [2] propose ISP-based aggregation. The ISPs add (physical) domain-specific prefixes thereby renaming the data for the purpose of aggregation. *E.g.*, `/alice-blog` will be renamed to `/att/atlanta/alice/alice-blog` since Alice is a user of AT&T in Atlanta. Although the concept of aggregation can help with the FIB scalability, this particular solution faces a dilemma when the data is replicated or moved. *E.g.*, if Alice blog is replicated in Columbus, the solution should either *a*) rename the data to `/att/columnbus/alice/alice-blog` in order to maintain the aggregation, but loses the benefits of using CCN since they are treated as two different data pieces; or *b*) maintain the same name to get the benefit of CCN thereby losing out on aggregation since the routers have to maintain 2 entries: `/att/atlanta` and `/.../alice-blog`.

In [3], the authors propose hash-based solution similar to Distributed Hash Table (DHT), wherein each router is responsible for only a subset of the global FIB. The data name is hashed to identify the router that has the necessary information (FIB entries) to reach the sources and the query is then forwarded to this router. This solution achieves FIB scalability at the cost of: *a*) longest prefix matching: since the hash is performed on the full data name; *b*) aggregation: because each data has to have an entry; *c*) single point of failure: only one router is responsible for a particular subset of names/requests; and *d*) path stretch: all the requests have to go to the responsible router in order to be routed towards the intended source(s) even if it is next to the requester.

In this work, we propose VDR, a routing scheme based on *virtual domains* – domains that are not bound to physical routers – to exploit the benefits of aggregation and hashing while avoiding the issues from [2] and [3].

ICN'15, September 30 - October 02, 2015, San Francisco, CA, USA.
ACM 978-1-4503-3855-4/15/09.
http://dx.doi.org/10.1145/2810156.2812600.

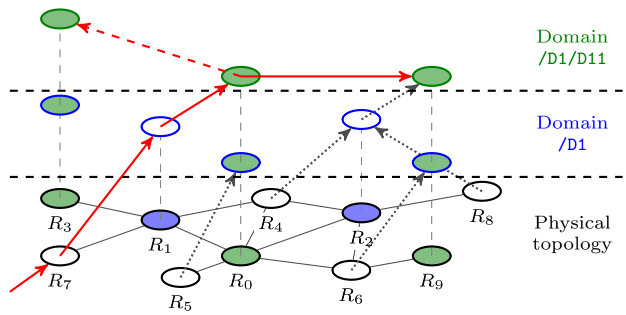

Figure 1: Example of VDR (FIB indicated with arrows).

2. VIRTUAL DOMAIN ROUTING IN CCN

In CCN, hierarchically structured names are used to simplify the name space management and aggregation similar to IP. However, IP addresses reveal the physical domain relationships while the hierarchical names cannot. Therefore, we introduce the concept of *virtual domains* to correspond these *virtual locations* (names). *E.g.*, when a consumer is requesting a data named `/facebook/simpson/gallery/...`, we see it as, the consumer needs to enter the (virtual) domain `facebook` and then the sub-domain `simpson`, and so on. These domains are created with the names instead of the physical locations.

To address the unbounded name space issue caused by the variable-length names, we use hash functions similar to [3]. But we calculate hashes for *each component* of the name and use the hash values for the inter-virtual-domain routing. A virtual domain `/D1/D11` represents all the names whose first component can be hashed to value `D1` and second component to `D11`. True that calculating a hash for each component increases the computation overhead on the routers, but this overhead can be easily reduced by precalculation techniques.

VDR then embeds these virtual domains onto physical routers similar to network virtualization [4]. Each router can belong to different domains at different levels (or even no domain at all). *E.g.*, in Fig. 1, routers R_0, R_3 and R_9 belong to domain `/D1/D11` (which means they also belong to domain `/D1`). The routers R_1 and R_2 belong to `/D1` while the others do not have a domain. The FIB in the routers either store how to reach a (sub-)domain D if it does not belong to D, or store the sub-domain information of D if it belongs to D. *E.g.*, Fig. 2 shows a possible FIB layout of R_0 in Fig. 1. In the level 1 domain table (left), since R_0 does not belong to

Dom.	Next Hop
/D1	→
/D2	R_1
/D3	R_5

Lv.1 Dom. Table

Dom.	Next Hop
/D1/D11	R_1
/D1/D12	→
/D1/D13	R_2

Lv.2 Dom. Table

Prefix	Next Hop
/facebook/...	R_9
/twitter/...	E_2
...	...

Intra Dom. Table

Figure 2: FIB implementation in VDR.

domain /D2, the next hop for /D2 is R_1. But R_0 belongs to /D1, therefore the next hop points to a second level domain table (middle) which describes the intra-/D1 information. If the network manager decides not to separate /D1/D11 further (which means /D1/D11 is a leaf domain), the intra-domain table becomes a prefix↦next_hop table (similar to the FIB in CCN, right table in Fig. 2). The router knows how to reach the data provider(s) for each name (or prefix). A longest-prefix match can be performed accordingly. The arrows in Fig. 1 represent the FIB for /D1 and /D1/D11.

When a request is issued in the network, the routers can forward it based on the virtual domain (name), level by level till it reaches a leaf domain. *E.g.*, when R_7 in Fig. 1 receives a request for /facebook/simpson/gallery/image1.png (red arrow), it would hash the first component and know that it does not belong to the domain /D1 (we assume *hash*("/facebook") =/D1). R_7 will then forward the request to R_0, which is the nearest node that belongs to the domain /D1. R_0 does not need to do further forwarding to reach domain /D1/D11 since it already belongs to the sub-domain. Instead, it performs intra-domain routing and forwards the packet to R_9 since R_9 serves prefix /facebook/simpson/gallery. We believe that it is feasible for routers to know how to reach each other since it is not unlike the routing in IP.

In VDR, aggregation can be easily achieved. The first components are hashed and grouped into domains. Therefore, the routers that are not in these domains only need one entry for each domain rather than one entry for every name in the domain. While performing routing in the leaf domains, the routers can pick either the best source or multicast a request to several potential providers (dashed arrow in Fig. 1) just like CCN and thereby exploit the dynamicity provided by CCN design. The cache replacement rule can also be modified based on the domains, *e.g.*, data in the domain to which the router belongs can have a longer life time since the router has higher probability to serve a request for that domain and thereby increase the cache hit rate. Since the routing table of each leaf domain is maintained on multiple routers, VDR can avoid the issue of single point of failure and balance the workload on the routers.

3. PRELIMINARY EVALUATION

We simulated VDR in a RocketFuel [5] topology (AS-1239, Sprintlink) to demonstrate the scalability and efficiency of our solution. The topology contains 315 routers and we place 945 end hosts on these routers, of which 500 are chosen as providers. A name space having 20,000 4-level random names is used in our simulation. We compare VDR with hash-based routing [3] and vary the level on which the leaf domains appear.

The traditional CCN routing which maintains all the provider information on each router is similar to VDR *level*=0 (represented as VDR Lv=0 in Fig. 3). It can achieve shortest-path forwarding (with lowest network load and average latency) but at the cost of more FIB entries maintained in the network.

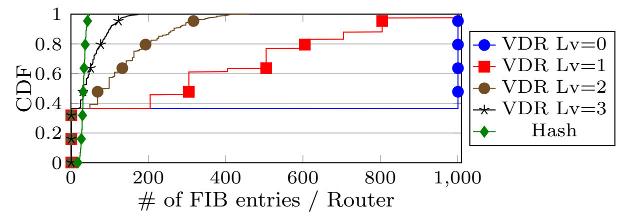

Figure 3: CDF of the states on each router.

With the increase of the domain levels, we can observe that the number of states stored in the network decreases. When *level*=3, the maximum number of entries per router drops to 175 and 90% of the routers store fewer than 100 entries (see VDR Lv=3 in Fig. 3). Because of the aggregation, the total number of FIB entries with VDR Lv=3 is approaching that with hash-based routing. We envision that VDR can get more benefit from aggregation when the name space is larger.

But we acknowledge that since the requests have to be forwarded based on the domains, they can be directed away from the data provider when there is no router belonging to the domain on the shortest path. The path stretch increases with the number of levels and cause (slightly) longer latency and more network traffic. This is a tradeoff between the scalability (space) and efficiency (time). We are working on finding the optimal balance in this tradeoff.

4. CONCLUSION

We propose VDR that uses the concept of virtual domains for routing in CCN. VDR addresses the FIB scalability issue in CCN by facilitating aggregation without compromising on the advantages of CCN such as ease of data replication and obtaining the data from a closer source. Preliminary evaluation shows that VDR lowers the number of FIB entries present at a router significantly without affecting much on path latency. We are currently working on a mechanism to allocate routers to different levels of the virtual domain, optimizing the solution and will perform extensive evaluations to compare the performance gain to the state of the art.

Acknowledgment

This research was partly funded from the National Science Foundation for Distinguished Young Scholars of China under Grant No.61225012 and No.71325002, the EU-JAPAN initiative by the EC Seventh Framework Programme Grant Agreement No.608518 (GreenICN), NICT under Contract No.167, the Volkswagen Foundation Project "Simulation Science Center" and the Chinese Scholarship Council.

5. REFERENCES

[1] V. Jacobson *et al.*, "Networking Named Content," in *CoNEXT*, 2009.
[2] L. Zhang *et al.*, "Named Data Networking (NDN) Project," PARC, Tech. Report NDN-0001, 2010.
[3] L. Saino *et al.*, "Hash-routing schemes for information centric networking," in *ICN*, 2013.
[4] "Network virtualization," https://en.wikipedia.org/wiki/Network_virtualization.
[5] R. Mahajan *et al.*, "Inferring Link Weights using End-to-End Measurements," in *IMW*, 2002.