

VSCCN: CCN with a Very Simple Control Plane

Insang Choi, Byungjoon Lee, Hongseok Jeon, Hoyoung Song, and Youhyeon Jeong

Smart Node Technology Research Team, ETRI

{ischoi, bjlee, jeonhs, hsong, yhjeong}@etri.re.kr

Abstract— The proliferation of user-generated content available over the Internet has introduced a new aspect for the design of the Future Internet, called Information Centric Networking (ICN), which deal with the need for unified information access and consumption. Content-Centric Networking (CCN) is one of the important ICN methodologies, which defines a new transport for content request and response packets. For the operation of CCN data plane, it is normally assumed that a control plane based on a routing protocol such as OSPF is used. While routing protocols such as OSPF are quite suitable for the interoperation between the legacy IP network and CCN network, such protocols are not suitable for supporting content mobility because of their slow convergence time. Therefore, in this paper, we suggest a control plane that supports content mobility with a very simple extension to the original CCN message formats.

Keywords— ICN, CCN, Routing, Control Plane, Mobility

I. INTRODUCTION

The proliferation of user-generated content available over the Internet has introduced a new aspect for the design of the Future Internet, which accelerates the transition from a host-centric to an information-oriented networking paradigm. Information Centric Networking (ICN) tries to deal with the lack of a global content naming and addressing scheme, as well as with the need for unified content access and consumption [1]. CCN (Content-Centric Networking) [2] is one of the important ICN methodologies, which is receiving a lot of attention from Future Internet research communities.

In ICN, all the information pieces are referenced by names, not by addresses. Normally, to handle the accesses to the information pieces by names, ICN methodologies use name-location mapping function $f:N \rightarrow L$ where N is the set of all information names, and L is the set of all possible locations. One of the most popular implementation options for f is to use a repository where the name-location mapping information is stored. However, CCN takes a more friendly approach that deploys Forwarding Information Base (FIB) entries which guides content-requesting packets to the actual content locations, which is quite similar to the legacy IP packet forwarding procedure that uses FIB entries as well. To create FIB entries, [2] have suggested using OSPF for disseminating name-location information throughout CCN routers. With the disseminated information, each CCN router is able to calculate FIB entries by itself. In this manner, the control plane of CCN can be implemented in a similar way to that of IP network.

The advantage of the CCN control plane that is based on IP routing protocols such as OSPF is that we are able to deploy CCN routers incrementally, by replacing edge IP routers with

CCN routers (Figure 1). As unknown LSA records in OSPF packets are safely ignored by the legacy IP routers, OSPF is suitable to deliver content name-location information.




Figure 1. IP-CCN interoperability [2]

The disadvantage of the control plane that is based on IP routing protocols is its convergence time that is not enough to support content mobility [3]. For example, let us assume that the content `/parc.com/media/art` is relocated from A and B to F. It takes time for a FIB entry (`/parc.com/media/art, F`) to be deployed on E. Until then, the content is simply not usable. Therefore, such control plane is not a good enough to support content mobility.

Therefore, in this paper, we suggest a novel CCN architecture, VSCCN, which has a very simple control plane. VSCCN is the same with CCN in every aspect except that VSCCN does not use routing protocols to disseminate name-location mapping information. Instead, VSCCN introduces REGISTER messages to correctly deploy FIB entries. Moreover, VSCCN does not heavily depend on FIB entries to forward INTEREST packets. Instead, VSCCN uses *core score* as a metric to forward INTEREST packets to the core of CCN network until to reach a CCN router that has matching FIB entries. With these two simple modifications to the basic CCN, we achieve content mobility with a reasonable handover time.

This paper is organized as follows. In section II, we briefly cover related research activities that have control planes which might be used as CCN control planes, and compare with our approach. In section III, we explain the VSCCN architecture in detail. Assumptions that we have made to design VSCCN are also covered. In section IV, we briefly cover the pros and cons of VSCCN with several research issues. In section V, we conclude this paper.

One thing to note is that this suggestion is in the work-in-progress stage. Therefore, we do not give extensive evaluation results in this paper. We only try to give a brief sketch of the proof that VSCCN is actually feasible, and we left the performance evaluation study as future work.

II. RELATED WORKS

A.W. Kazi [4] suggested a control plane based on bloom filter, which is originally targeted for reducing INTEREST packet flooding. The suggestion uses DATA packet to deliver content summaries represented by bloom filters to all CCN routers on the DATA path. Besides, the suggestion is based on the wrong observation that INTEREST packets are flooded to all CCN routers to find content locations (INTERESTs are not flooded but guided by FIB entries), the suggestion is not suitable for supporting content mobility because the content location information slowly converges.

R. Ravindran [5] demonstrated OSPF-based CCN control plane at CCNxCON 2011, and showed a chatting program to work without any static configuration. This contribution is actually an implementation of the control plane suggested in [2], thus it does not provide any solution for the content mobility issue.

S. Shanbhag et al [6] suggested a service-oriented routing architecture based on Ant Colony Optimization. Though this suggestion provides a routing optimization method, the service mobility performance is basically restricted by the underlying CCN routing mechanism because the method executes optimizations on already-deployed FIB entries.

Z. Ming et al [7] suggested a service and information oriented network architecture, SIONA, which is very similar to CCN and VSCCN. In contrast to CCN, SIONA uses *register* messages to implement its control plane. The *register* messages are for deploying *registration table* (RT) entries, which correspond to CCN FIB entries. However, [7] does not suggest any forwarding scheme for *register* messages. Moreover, clients for a specific content cannot access the content until all RTs correctly converge, which is no different from the approaches based on routing protocols.

III. VSCCN CONTROL PLANE

In designing VSCCN control plane, we assumed that each VSCCN node is configured to have *core score*. The core score is a value that indicates how the VSCCN router is close to the core of the VSCCN network. To guarantee the correctness of VSCCN, core scores should be properly assigned.

The core score can be calculated in various ways. For example, if a VSCCN router has two 10Gbps links that lead to the core of the VSCCN network and eight 1Gbps links that aggregate traffic coming from the network edge, we can use $2 * 10 + 1$ as its core score. As there is a tendency that routers which are close to the core of network have greater link bandwidths for upstream links, this might give a moderately good approximation for core scores. However, the detail of the core score calculation algorithm is beyond the scope of this paper. It needs further research to prove its correctness.

In this paper, to make explanation simple, we assume that the core scores are manually given by network administrators. We also assume that there is a link management protocol that exchanges core score values between two adjacent VSCCN routers. The implementation overhead of the link management protocol is negligible because of its simplicity.

A. Registration Message Processing for FIB deployment

To deploy FIB entries on VSCCN routers without the aid of routing protocols, we introduce a VSCCN-specific control message, *REGISTER*. Basically, register messages are forwarded to all VSCCN ‘upstream’ routers that have greater core scores than the sender of the register messages. Whenever a new content is registered, the router that hosts the content sends a new *REGISTER* message. The message format is provided in Figure 2.

Name
Location
Nonce

Figure 2. REGISTER message format

As depicted, a *REGISTER* message delivers the name of the content, its location, and a nonce value. The nonce value is random, and is used to identify a *REGISTER* message uniquely in a short period of time. The nonce value plays an important role because VSCCN routers do not know the VSCCN network topology due to the lack of global routing protocol. It aids VSCCN routers to suppress duplicate *REGISTER* messages are flooded.

On receiving an ‘upstream’ *REGISTER* message r from a face f , a VSCCN router v updates its FIB entries according to the following rules. We assume the FIB entries are tuples of content name and the list of faces which are tuples of face and nonce value.

- (1) If no matching FIB entry is found, adds $\langle r.Name, \langle f, r.Nonice \rangle \rangle$ in the table. After updating the table, forwards r to all *upstream* routers.
- (2) If there is a matching entry m for r , checks if the face list l of m if l has f , or whether there is a face that has the same nonce value. If true, drops r to suppress further processing. If false, sends r to each face in l , and sends m to f after converting as a new *REGISTER* message $\langle r.Name, v, r.Nonice \rangle$. After that, adds f to the list l with $r.Nonice$, and forwards r to all *upstream* routers.




Figure 3. REGISTER message flow

Summarizing above rules, all register messages are basically forwarded to ‘upstream’ direction to the core of the VSCCN networks. Therefore, all FIB entries created by the messages are basically facing the ‘downstream’ direction to the edge of the VSCCN network. As all the REGISTER messages eventually hit VSCCN core routers, core routers would have multiple FIB entries for the same content name.

However, placing multiple entries for the same name in the core routers is not enough to support VSCCN operation. For example, if there are the same content files `/.../a.mp4` in separate branches, such as router A and D of Figure 3, the router B does not know the existence of the same content within router D. Thus, the router B would send INTEREST messages to router A only. To prevent this problem, VSCCN nodes send duplicated REGISTER messages to other faces whenever a REGISTER message ‘hits’ the matching FIB entries that share the same name, or the name that could be aggregated. Therefore, in Figure 3, the two REGISTER messages (① and ②) for the same content are duplicated at the router C to the downstream directions (③ and ④), which guarantees that all the paths from the content to the core have FIB entries for the content.

On receiving a ‘downstream’ REGISTER message r coming from a VSCCN router that has the same or greater core score value, the receiving VSCCN router v updates its FIB entries in the same way with upstream REGISTER messages, except that the direction is completely reversed. Therefore, FIB updates in an upstream routers eventually lead to the FIB updates in all downstream routers (④, ⑤, ⑥ and ⑦).

Therefore, the REGISTER forwarding scheme described above makes all tree branches which share a common root router have consistent FIB entries, even though there are multiple locations of the same content within the same tree. It is because all related FIB entries are disseminated to downstream VSCCN routers via duplicated REGISTER messages. We use this characteristic for sketching the proof of VSCCN correctness in subsection III.D.

B. INTEREST packet processing for content request delivery

Contrary to the INTEREST forwarding scheme of the basic CCN [2], VSCCN routers forward INTEREST packets to upstream direction even though there is no matching FIB entry. This makes no global FIB convergence is required to support content mobility. The INTEREST forwarding stops when no further forwarding is possible because there is no upstream router. The undeliverable INTEREST packets are dropped. The VSCCN INTEREST forwarding rule is summarized as follows.

- (1) If no matching FIB entry is found, forwards an INTEREST packet i to all VSCCN routers which have the same or greater core store values only when there is no matching PIT entries, after leaving a PIT entry with a nonce value $i.Nonce$.
- (2) All the other cases are processed by following the forwarding rule of the basic CCN [2].




Figure 4. INTEREST message flow

Therefore, all the INTEREST messages eventually hit FIB entries, which will lead the messages to the actual content locations (Figure 4).

The problem of this scheme is that INTEREST packets travel longer than the basic CCN INTEREST packets. However, if there is matching content, INTEREST-DATA packets are correctly exchanged.

C. Content Mobility

To show that VSCCN handles content that moves very fast from one router to another, we use a scenario in Figure 5, where the content `/.../a.mp4` moves from router A to route F.




Figure 5. INTEREST message flow: nomad content

As the relocated content introduces a new REGISTER message that eventually makes all downstream routers of H to converge immediately, all the following INTEREST packets will be correctly processed. Thus, VSCCN is able to handle content mobility fast enough to support handover.

Besides, VSCCN similarly handles the case where router moves with content. For example, if router A is relocated to the downstream position of router N, the router A will purge all FIB entries that it has, and sends new REGISTER messages which will eventually make the tree rooted by the router M have consistent FIB entries with other VSCCN routers (Figure 6).




Figure 6. INTEREST message flow: nomad router

D. Sketch of the Proof – VSCCN feasibility

Before completing the proof, we first assume that all FIBs are always correctly deployed in the same tree that shares the same core VSCCN router. (*Assumption 1*)

As INTEREST packets would be correctly delivered to all core routers if core scores are correctly configured, all INTEREST packets will be eventually delivered to all content locations after hitting a set of matching FIB entries. Therefore, if the *assumption 1* is true, VSCCN operates correctly.

To prove the *assumption 1* is always satisfied, we use the fact that every upstream REGISTER packet deploys FIB entries on every VSCCN routers on the upstream path. If a matching FIB is found in one of the upstream routers u and the FIB is facing other VSCCN routers, all the FIB entries of u will be pushed to all downstream routers of u . Therefore, duplicated contents on different downstream branches of u make all downstream routers of u have consistent FIB entries. Thus, INTERESTs from one of the downstream routers of u will be correctly handled. INTEREST from all the other routers will be correctly forwarded by core routers, because the core routers are guaranteed to have all FIB entries for the requested content.

Even though there is no duplicated content in the downstream branches of u , INTERESTs from one of the downstream routers of u will be all eventually handled by u or one of the downstream routers of u . INTEREST from all the other routers will be correctly forwarded by core routers, because core routers are guaranteed to have all the FIB entries for the requested content. As any VSCCN routers can be u , the *assumption 1* is shown to be satisfied. $\square$

IV. BENEFITS AND DEMERITS

The first benefit of VSCCN is its convergence time which is good enough to support content mobility.

The second benefit is the locality of FIB deployment. For example, as shown in Figure 3, FIBs for a specific content are localized within the trees where the content files reside. Not until the content is completely transferred to other VSCCN router and registered again, FIBs are not transferred outside of the trees. This characteristic is good for attaining scalability

because it reduces the number of FIBs to be deployed on the network. It means that VSCCN saves FIB memory for content files which are popular only in a specific network area.

The first downside of VSCCN is that INTEREST packets are forwarded to content even though it does not exist. If a malicious user sends a lot of INTEREST packets for random names, the packets might incur massive traffic to the core of VSCCN network. To handle that issue, we are testing a method of putting a signature inside of INTEREST packets.

The second problem of VSCCN is the control plane based on REGISTER packet does not allow easy interoperability with legacy IP routers. However, interoperability with IP router is not a big issue currently because the current version of VSCCN operates in IP overlay mode only.

The third problem is that current VSCCN architecture does not have any means to remove stale FIB entries. This problem will be fixed in the later version of VSCCN.

V. CONCLUSIONS

VSCCN is a simple modification to the basic CCN, targeted to solve the content mobility problem. VSCCN also prevents FIB entries from being deployed outside a network tree where content resides, which is good for attaining scalability.

Currently, we are working on solutions for the issues listed in the section IV. Besides, we are elaborating an algorithm that automatically calculates core scores. As VSCCN architecture depends on core scores for its correct operation, it is important to have a scheme for their calculation.

The newer version of VSCCN which adopts above enhancements will be publicized in a near future with extensive analysis results of its performance.

ACKNOWLEDGEMENT

This research was supported by the KCC (Korea Communications Commission), Korea, under the 'R&D on Smart Node Technology for Cloud Networking and Contents Centric Networking' support program supervised by the KCA (Korea Communications Agency) (KCA-2011-11913-05003)

REFERENCES

- [1] B. Ahlgren, P. A. Aranda, P. Chemouil, S. Oueslati, L. M. Correia, H. Karl, M. Söllner, A. Welin, "Content, Connectivity, and Cloud: Ingredients for the Network of the Future", *IEEE Communications Magazine*, pp.62-70, July 2011.
- [2] V. Jacobson, D. K. Smetters, J. D. Thornton, M. F. Plass, N. H. Briggs, and R. L. Braynard, "Networking named content," *In Proc. of ACM CoNEXT'09*, 2009.
- [3] B. Ahlgren, C. Dannewitz, C. Imbrenda, D. Kutscher, and B. Ohlman, "A Survey of Information-Centric Networking (Draft)", Available: <http://drops.dagstuhl.de/opus/volltexte/2011/2941/pdf/10492.KutscherDirk.Paper.2941.pdf>
- [4] A. W. Kazi, "Prefetching Bloom Filters to Control Flooding in Content-Centric Networks", *In Proc. of ACM CoNEXT Student Workshop*, 2010.
- [5] R. Ravindran, "Dynamic Routing in CCN", *CCNxCON*, 2011.
- [6] S. Shanbhag, N. Schwan, I. Rimac, and M. Varvello, "SoCCeR: Services over Content-Centric Routing", *In Proc. of ACM SIGCOMM ICN Workshop*, 2011.
- [7] Z. Ming, D. Li, C. Xia, and M. Xu, "SIONA: A Service and Information Oriented Network Architecture",