Supporting Dual-Mode Forwarding in Content-Centric Network

Ravishankar Ravindran, Guoqiang Wang, Xinwen Zhang, and Asit Chakraborti
Huawei Research Center, Santa Clara, CA 95050, USA
{ravi.ravindran, gq.wang, xinwen.zhang, asit.chakraborti}@huawei.com

***Abstract*—Information-centric networking (ICN) aims to replace current host-centric IP architecture with one based on efficient, secure, and reliable dissemination of information. This paper extends content-centric networking (CCN), one of the ICN architectures, for dual-mode forwarding, wherein, the forwarding plane distinguishes sharable traffic from non-sharable traffic characterized by its global shareability. The proposed extensions allow sharable traffic to be forwarded in a *non-expedite* mode following the packet processing of CCN, while non-sharable traffic is forwarded in the proposed *expedite* mode subjected to fast-path forwarding only involving the FIB. We enable this dual-mode forwarding using a new header which optionally includes a source-object name. We support our proposal with simulation results where the benefit of separating these two types of traffic in the forwarding plane results in better utilization of content store, pending interest table, and improved user QoE for sharable applications in terms of content delivery latency.**

I. Introduction

Content-centric network (CCN) [4] is a future Internet architecture where the narrow waist is based on named chunks. CCN forwarding plane enables location independent data dissemination, multicasting, caching, and information security as integral part of the network layer addressing several issues encountered in IP today. Two types of packet primitives are defined in CCN: *Interest* packet generated by a content consumer which includes the name of the content object being requested; and *Data* packet which includes the name, content, and a producer generated signature which binds the content to its name. CCN uses three lookup tables in its forwarding plane: first, a content store (CS) to match the name in an Interest to the cached content; second, a pending interest table (PIT) that aggregates similar Interests, allowing dynamic construction of multicast tree for information dissemination; third, a forwarding information base (FIB) table, which helps to route Interest packets to content producers.

Problem: In CCN, each data chunk is named and can be retrieved by expressing an Interest to the network with the name. During Data forwarding, the content is cached in CS of routers along the forwarding path, such that it can satisfy subsequent Interests with the same name. With this in-network caching mechanism, current CCN provides high efficiency of data dissemination for globally sharable content such as user generated content and public web pages. However, this universal caching and in-network processing mechanism introduces inefficiencies for non-sharable content, i.e. content that is personalized by the producer for the consumer such as two or multiparty conversational voice/video session, financial transactions, or control plane interactions.

Following are CCN considerations for non-sharable content: 1) Processing each non-sharable content chunk through the CS and PIT introduces unnecessary inefficiency for both non-realtime and realtime classes of sharable applications. Here realtime refers to content generated upon producers's initiative requiring immediate sync with the receivers such as VoCCN [3], while non-realtime content is generated before any user requests for it. 2) Caching content related to personal transactions and realtime applications can violate privacy requirement, since such content is not desired to be accessed by receivers other than the legitimate ones. 3) Non-sharable content may constitute low percentage of the overall Internet traffic [1], however CCN router maintains *per-packet* state for all Interests and Data due to PIT and CS processing, this shall result in usage of high-cost computational and memory resources in a router. This situation would be particularly worse in the network backbone, where routers aggregate high volume of Interests pertaining to non-sharable content. 4) From security perspective, when leveraged by malicious content producers and receivers, the non-sharable content can significantly affect the performance of applications with sharable content in the network; that is, caching and processing non-sharable content can aggravate denial of service (DoS) threat to the network. For these reasons, it is desirable to treat sharable and non-sharable content such that it would not affect each others performance objectives.

Dual-mode Forwarding: Considering these, we propose a dual-mode forwarding plane, as an extension to current CCN architecture, based on different types of traffic. Specifically:

We first classify traffic in ICN into two categories: globally sharable and non-sharable content, each of which can include realtime and non-realtime content.

We then propose augmented PDU headers for Interest and Data packets in CCN to distinguish the two types of traffic. Our main extension includes a header field to indicate the forwarding mode of a CCN PDU. For non-sharable content, a PDU includes both source- and destination-object names for a fast forward processing, where the names correspond to content requester and producer, respectively. Accordingly, we adopt two forwarding modes in CCN: for sharable content, a *non-expedite* mode is adopted, which follows the current CCN architecture; for non-sharable content, we adopt an IP-like *expedite* mode, where Interest and Data are forwarded based on source- and destination-object names in the PDUs using the FIB. Enabling such differentiation in the forwarding plane addresses or mitigates several afore mentioned issues resulting from CCN forwarding as proposed in [4].

Outline: The paper layout is as follows: Section II argues the need for the dual-mode forwarding plane. Section III proposes the CCN header to enable dual-mode forwarding, and Section IV discusses dual-mode packet processing. Section V presents simulation results to demonstrate the effectiveness of such traffic differentiation, and Sections VI and VII provides more discussion and conclude the paper.

II. Motivations

Following, we argue the need for dual-mode forwarding and considerations to mark a flow as expedite and non-expedite are highlighted.

Content shareability and CS efficiency: Content can be characterized as sharable or non-sharable based on its shareability property. While certain content such as public web pages or two party conversation can be easily classified as being sharable or non-sharable; many other applications such as multi-party voice conversation, enterprise video conference spanning multiple sites, video-on-demand stream with DRM privileges can be sharable or non-sharable depending on contextual factors including, content security and privacy issues of producer and consumer, and SLA requirements such as QoS and reliability requirements. As we see, shareability property of a content has several dimensions, and is driven by user's, application's, or service's requirements. Without delving more on modalities of classification, we focus this paper on supporting these two traffic types in CCN.

Classification based on shareability has implications on design of content router's forwarding plane. While, sharable applications can leverage in-network caching, non-sharable content is expected to pass through the router without retaining any memory of what is being requested or delivered. Contents under both these categories can be realtime or non-realtime. As sharable contents are driven by user popularity, they can leverage in-network caching to optimize the network bandwidth usage. Non-sharable content is personalized and expected to have very limited scope of interest and should be protected from exploitation, which motivates the need to avoid the PIT and the CS for such traffic. Also, caching non-sharable content consumes valuable CS and PIT resources at the detriment of sharable content applications. In CCN, the age of non-sharable content can be set to 0 to avoid caching, however this still consumes router's memory and computation facilities and does not prevent CCN forwarding plane to multicast such content to undesired users if the Interests were expressed with appropriate names.

PIT efficiency: CCN by default treats all content alike. A state is created in the PIT for every Interest and applies cache processing to all Data irrespective of the type of application. This mode of operation creates inefficiencies for both sharable and non-sharable applications. PIT table is a limited resource whose size increases, among other factors, with increasing load of non-sharable content and its affect shall be more severe at aggregation points of the Internet such as at aggregation routers of access networks or at point-of-presence (PoP) nodes of a WAN. PIT entries due to non-sharable content could also cause rejection of requests for sharable content, in turn affecting the QoE of such applications, particularly during high load scenarios. In the proposed dual-mode forwarding, expedite Interests are not subjected to PIT processing, resulting in more PIT resource for Interests of sharable applications. Though PIT in CCN provides several features including content multicasting, loop-free data forwarding, multi-path strategy forwarding, symmetrical Interest and Data forwarding, and flow balancing, these features may not serve the interest of non-sharable applications due to the nature of these applications.

Realtime application requirements: Just like queuing delay, any form of packet processing involving CS and PIT table lookup using the full name in Interest packets adds to end-to-end latency. Though this affects the performance of both non-sharable and sharable realtime applications, the latter benefits from caching and multicast efficiency. Forwarding traffic with stringent QoS requirements in expedite mode can minimize this undesirable affect on non-sharable content streams. Realtime applications such as voice conversation typically have one-way delay requirement of 150ms. This requirement is barely met in cellular networks today due to factors such as user mobility, resource constraints, and radio frequency impairments. Any further processing delay in the routing nodes would demand further optimization in these already resource constrained access networks.

To address the above issues, we propose dual-mode forwarding. In dual-mode, applications are allowed to tag its Interests and Data as either being expedite or non-expedite. Expedite or fast-path processing in CCN shall only use the FIB, while non-expedite or slow-path forwarding follows the normal CCN processing. Applications generating non-sharable content are candidates for expedite forwarding while applications generating sharable content are candidates for non-expedite forwarding.

III. CCN Header

Considering the above discussion, we extend the current CCN headers to achieve the following objectives: (1) To enable dual-mode forwarding operation, i.e., forwarding Interest and Data in expedite or non-expedite mode as desired by applications; (2) Use source-object name to forward packets to destination application generating Interests; (3) Allow CCN nodes to only store non-corrupt content objects;(4) Broaden the scope of CCN to potentially enable enhanced services by the network. Fig. 1 and Fig. 2 show our proposed CCN Interest and Data header, respectively.

Following we describe the role and function of each field in the Interest and Data PDUs in the context of two forwarding modes. A field definition holds for both types of PDUs unless otherwise specified.

Message Type: This mandatory field identifies the type of packet with value of either `Interest` or `Data` type.

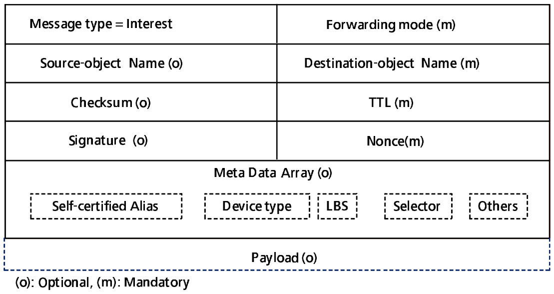

Fig. 1. **Interest packet in dual-mode CCN forwarding.**

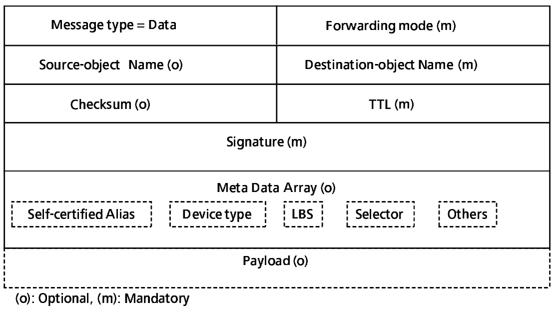

Fig. 2. **Data packet in dual-mode CCN forwarding.**

Forwarding Mode: This mandatory field is set by application with value of either `expedite` or `non-expedite` mode. This field is checked by router to differentiate expedite and non-expedite mode of forwarding. Considering the criticality of this field, a concern here may be related to support for device mobility wherein non-sharable applications could benefit from in-network caching. These features are still preserved in dual-mode forwarding considering the flexibility to shift from expedite to non-expedite mode by setting the *Forwarding Mode* bit appropriately through appropriate control plane intervention on an end-to-end basis or enabled selectively at service points through configuration based on certain contextual information, such as initiation of handoff operation.

Source-object Name: This optional field is required if the application elects the expedite mode of forwarding. Here the name identifies an object that could be an application, service, device, or user. The name is an aggregatable human name as proposed in [4]. Name aggregation is a requirement for scalable content routing. In terms of forwarding operation, if the forwarding is set to expedite mode, the source-object name is used to forward the Data PDU back to the entity expressing the Interest.

Destination-object Name: This is a mandatory field. In expedite mode this field is used for Interest forwarding using the FIB and bypassing both PIT and the CS. In non-expedite mode it follows the default CCN forwarding operation. The naming requirements in this case are similar to those discussed for the source-object name field.

Checksum: This is an optional field to prevent content routers from receiving corrupted Interests or Data. Enabling routers to identify corrupted Interest (or Data) allows an upstream (or downstream) router to re-request the Interest (or Data) from its downstream (or upstream) node. Checksum is computed over all fields of the header and the payload by the CCN protocol where the Interest or Data is being generated. One could argue the use of the signature in Interest or Data to validate the integrity of the content. There are a few reasons why this is not the best approach: First, a content producer may be different from a content distributor, therefore the content could be produced and signed by an entity different from the one distributing it; Second, a packet verification using signature requires access and use of a public key, and is computationally costlier compared to verifying the checksum. We note that the checksum validation is imposed at every hop of the network, and re-generated if any field in the header is updated.

Time to Live (TTL): This is a mandatory field representing the life time of an Interest or a content chunk. In expedite mode, it represents the maximum number of hops the Interest or Data PDU is allowed to be forwarded before it is dropped, and the field is updated at each content router. In non-expedite mode this field represents the lifetime of the content and is used by the content store to determine the amount of time the content can be cached before it is purged. Based on content policy, this field may be updated by a content router. In non-expedite mode, for the Interest PDU, in addition to expressing the life time of the Interest, this field can be used to express persistent Interests to support *push* based events as proposed in [6]. Alternatively, a TTL of 0 in a non-expedite Data PDU can be used to represent realtime notification event send to its subscribers.

Signature: This is a mandatory field for a Data PDU. The signature is generated as a function of name, content, and other metadata of interest, and associates provenance to the content. Upon receiving a Data PDU, the requester verifies the signature to validate the provenance. A content router could also optionally validate the signature for the same purpose.

Nonce: This is a mandatory field in the Interest PDU, and is used to detect duplicate Interests by a router in non-expedite mode as suggested in [4].

Metadata: The metadata fields are optional. This field includes user- or content-centric information that can be used for policy-based forwarding, storage, and processing operations. For example, consumer-specific self-certified alias (flat names) can be used to enable access control and support content VPN service. When certified alias exist in both Interest and Data PDUs, a content router can enforce access control by matching the ones included by the requester to the ones generated by the producer. Device type and GPS can be used to support location based services (LBS). Metadata such as a selector can be used to scope the content lookup in the content store. In addition, other value-added services can be enabled by specifying instructions in metadata to conduct a specific data processing, such as transcoding or content mash-up.

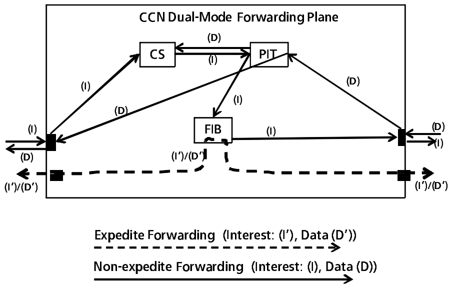

Fig. 3. **Dual-mode packet processing.**

IV. Packet Processing

Fig. 3 shows the packet processing in a CCN router supporting dual-mode forwarding. Interest and Data packets are marked as expedite or non-expedite by applications, based on the data's shareability requirement as determined by the service. At the incoming face of the CCN router, packets are classified based on the *Forwarding Mode* bit. If the bit is set to expedite mode, the forwarding follows the fast-path processing as shown by the dotted lines, where the corresponding Interest and Data packets are processed only through the FIB. If the *Forwarding Mode* is set to non-expedite mode, the forwarding follows the normal CCN processing involving the CS, PIT, and FIB.

V. Simulation Analysis

In this section we describe our experimental setup and results from simulation analysis. Our goal is to study the benefits of our proposed forwarding differentiation over a simulation framework, i.e., to compare the efficiency gained by operating CCN in the dual-mode by stressing the CS and PIT resource while network links are provisioned with sufficient capacity to avoid bandwidth bottleneck. We choose NS3-DCE [5] for our analysis. The following modification and extensions were made to CCN to conduct our study.

Changes to CCNx: In ccnx-0.4.0, the PIT and FIB are implemented as a tightly cross-referenced lookup table, making it practically difficult to implement an efficient fast-path forwarding implementation without significant changes. However, our changes ensure special handling of expedite traffic to achieve our goal of subjecting it only to FIB processing. CCN Interest and Data packets are modified to include a source-object name field, and a tag to identify it as expedite or non-expedite. In expedite case, during Interest processing, Interests are forwarded using the destination-object name to the next hop, while the source-object name in the Data PDU is used to forward it to the requesting application.

Performance metrics: Additional states are added to CCNx to collect performance metrics. The key performance metrics used to compare the efficiency of the forwarding schemes include maximum CS size, maximum PIT size, cache hit ratio, and average round-trip time (RTT). RTT is measured from when an Interest is issued to the time when the Data is received and is averaged over all the content objects received during an application session.

TABLE I
Simulation Setting

CCNx	Release Version-0.4.0
Sharable Content Request Rate	Poisson(mean=80$requests/s$)
Mean Sharable Content Size	Geometric(mean=100$Chunks$)
Number of Sharable Contents	2000
Popularity Distribution	Zipf(Exponent parameter=2)
Number of Popularity Classes	100
Content Store Size (default case)	4000$Chunks$
Sharable Content Chunk Size	$1024B$
Voice Call Rate	Poisson(mean:$[0.5-2.5]calls/s$)
Voice Call Duration	Constant(60s)
Voice Packet Size	$160B$

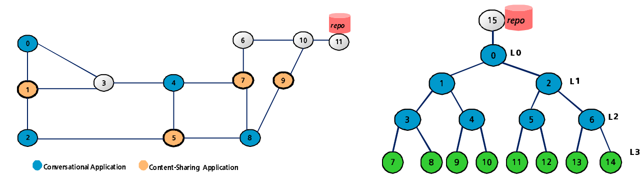

(a) Abilene (b) Tree

Fig. 4. **Simulation topologies.**

Simulation parameters are summarized in Table I. Internet2 abilene topology and a four level tree topology shown in Fig. 4(a) and Fig. 4(b) respectively are used for the study. The FIB is provisioned statically and Interest routing follows shortest path first logic. The traffic for the analysis is a mix of content-sharing and voice conversations. The traffic details are as follows.

Content sharing application: This application models traffic due to content sharing among users. With reference to Fig. 4(a), node 11 in the abilene topology and node 15 in the tree topology are the repository nodes for sharable contents. In the abilene case, nodes $\{1, 5, 7, 9\}$ are chosen to generate requests for shared contents, while all nodes at level-3 (L3) in the tree do the same. The content sharing application is built on *ccndsendchunks* and *ccncatchunks2* utilities included in CCNx release. We operate *ccncatchunks2* with window size of 1, and the content expiry is set to $16s$.

Conversational application: This application simulates peer-to-peer streamed conversational content. The application is modeled as constant bit rate voice application with an Interest packet generation rate of $50packets/s$ and voice payload of $160B$ in the Data. With reference to the two topologies, nodes $\{0, 2, 4, 8\}$ in abilene and L3 nodes in the tree are chosen to generate traffic for conversational content. In the non-expedite case, the expiry of voice packets is set to minimum allowable value of $1s$.

Next we present the results emphasizing on benefits achieved by dual-mode forwarding compared to the unmodified CCN case. Though statistics are collected in all the routers, the results are reported for selected ones based on maximum CS and PIT load, and impact on the performance metric due to dual-mode forwarding.

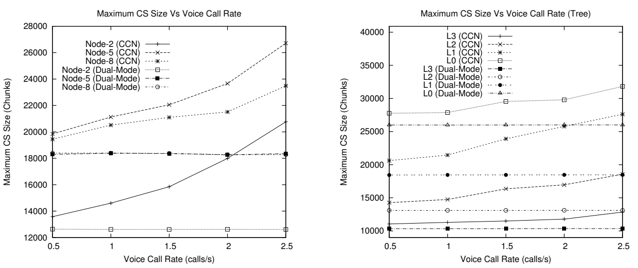

(a) Abilene (b) Tree

Fig. 5. **Maximum CS size with varying voice call rate.**

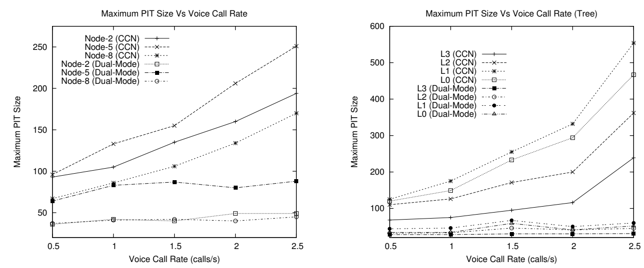

(a) Abilene (b) Tree

Fig. 6. **Maximum PIT size with varying voice call rate.**

A. Content Store Efficiency

For the abilene case, results shown here correspond to nodes 2, 5, and 8 as these nodes have the highest CS and PIT utilization, while for the tree case we average the performance of all nodes at each level, L3 being the leaf level and L0 the root. Fig. 5(a) and Fig. 5(b) compares the performance of maximum CS size for varying voice load under the two forwarding modes. Without dual-mode, the maximum CS size increases with increasing rate of the voice calls. This is expected, as increasing call rate causes more calls to be active per unit time, hence more voice content is cached in the CS, increasing its utilization in both edge as well as the transit routers. For the tree case, the maximum CS size at the root is almost 2.5 times the CS size at the leaf level, this is a direct consequence of aggregation of voice load as we traverse up the tree hierarchy. In the dual-mode case, for both the topologies, the maximum CS size is unaffected by the voice load. Here, the voice data packets bypass the CS processing, leaving no memory of it the CCN nodes, this causes the CS size to remain the same irrespective of the voice call rate.

B. Pending Interest Table Efficiency

Here we analyze the efficiency gained with respect to maximum PIT size. The PIT size increases when the requested content objects aren't available in the CS, making its behavior depend on the cache management policy. The results here are presented with respect to nodes 2, 5, and 8 for the abilene and for different levels in the tree case. In the abilene case shown in Fig. 6(a), with increasing voice call rate, while the PIT size grows proportionally under native CCN forwarding, the effect is non-linear in the tree case as shown in Fig. 6(b). Further, for the tree, the PIT size at the root is also about 2.5 times of that of the leaf level due to voice load aggregation. The PIT size at L1 is more than L0 because of higher call volume handled at L1 than L0. In contrast, in the dual-mode, the PIT size remains unaffected in the dual-mode case. This is because the expedite marked voice Interests are subjected to fast-path forwarding using the FIB, bypassing the PIT processing. Similar to above CS case, the reason for the difference in PIT size for three nodes is due to the Interest arrival pattern and the Interest routing logic in the network.

C. Content-Sharing: Cache Hit Ratio

The cache hit ratio for the content sharing application is presented with respect to node 5 for the abilene case and as average of the L3 nodes in case of the tree. Hit ratio performance depends on cache management policy. CCNx implements a least recently used (LRU) like cache management scheme where objects are tracked by accession number and content staleness metrics. When the number of content chunks exceeds the engineered cache bound, both the metrics are used to evict cache objects, but we observed certain anomalies with the implementation. With increasing voice call rate, the cache hit ratio in the unmodified CCN case is expected to decrease, as more voice content contends for limited CS resource, but the observation was not so for both the topologies. In Fig. 7(a), for the unmodified CCN case, we observe the performance is unaffected and in case of Fig. 7(b) very little deterioration is observed. To explain this, we refer to Fig. 8 which compare the performance of CS size for maximum and minimum voice call rates respectively. The sawtooth nature of the graphs is due to the cache eviction operation that clears 500 content chunks every $5ms$ over a cache clearing cycle period of $15s$. The observation of CS size is made every $200ms$. From Fig. 8 we see, for both the cases, when the voice call rate is increased from $0.5calls/s$ to $2.5calls/s$, the upper bound far exceed the CS bound threshold of $4000chunks$, and the lower bound of the CS increases almost 3 times. This increase in CS lower bound negates the expected degradation of content-sharing application's hit ratio. In case of the dual-mode, when the cache bound is set to the default case of $4000chunks$, no improvement was observed. This is because, in dual-mode, the CS operates at lower cache size compared to the unmodified CCN scenario where sharable content benefit from bigger CS size. To be fair to the dual-mode, we had two options, either to modify CCNx's LRU behavior to follow the strict cache bound or increase the cache bound in the dual-mode to the average CS size achieved in the unmodified CCN case for different call rates respectively. With the latter approach, significant improvement is observed in both abilene, Fig. 7(a), and the tree, Fig. 7(b), scenarios, respectively. We later verified the improvement even when the LRU behavior was modified to clean the cache to the engineered CS size.

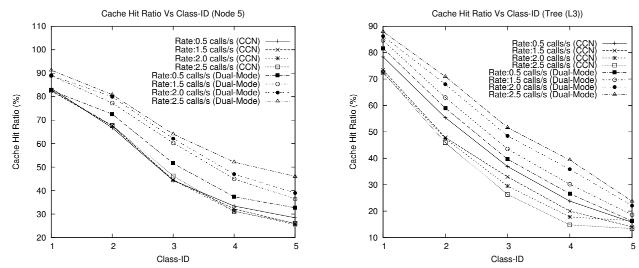

(a) Abilene (b) Tree

Fig. 7. **Cache hit ratio for varying class-id.**

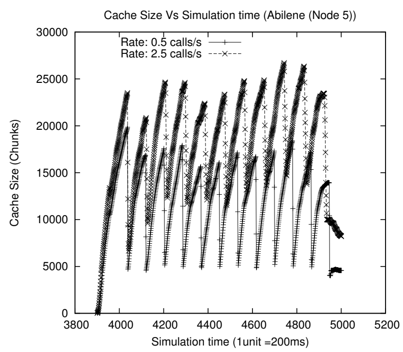

Fig. 8. **Cache size trace for varying voice call rate.**

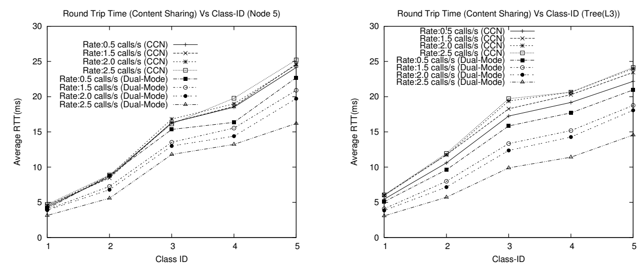

(a) Abilene (b) Tree

Fig. 9. **File sharing round trip time for varying class-id and voice call rate.**

D. Content-Sharing: Round-Trip Time

Fig. 9(a) and Fig. 9(b) show the average RTT performance for the content-sharing application for the two topologies. With modification discussed above, for both topologies, the RTT performance improves with increasing voice call rate, this is because the probability of content hit rate improves compared to that in the unmodified CCN case when voice traffic shares the content store. The average RTT improvement is better in case of the tree than the abilene case, because of efficient use of the CS at the aggregation points. For both the topologies we note that the percentage improvement of RTT for each content class increases with increasing voice load indicating the advantage gained by separating the expedite traffic from non-expedite traffic.

E. Voice Application: Round-Trip Time Performance

Comparing the unmodified to dual-mode CCN, the voice application did not show any improvement for both the topologies. It was observed to be as good as or little worse than the unmodified CCN case. The main reason for this is CCNx's implementation of the PIT and FIB as one logical look up data structure which prevents an efficient fast-path forwarding implementation. An implementation which separates the FIB from the CS and the PIT processing should result in improved latency performance for expedite applications.

VI. Discussion

Similar concern of traffic differentiation in ICN was raised in [6], which proposes extensions to CCN to deal with loss tolerant streaming content called *Channels* and loss intolerant content called *Documents*. The authors suggest a new type of Interest called *Persistent Interest* which reduces the overall control overhead while preserving the benefits of CCN. However, these extensions do not address the issues related to dealing with non-sharable content. The traffic differentiation proposed in [6] can be considered as sub-category under our proposed classification based on content sharability.

An important change we propose to the CCN protocol is to include a source-object name for expedite mode. This exposes the name or identity of the content consumer to the network and the producer. First of all, we believe that for conversational applications such as voice and financial transactions, the session is bi-directional and hence the content producer needs to know the name or identity of the content consumer, e.g., for authentication purpose. Therefore we believe this privacy violation against the content consumer is not a major issue for such applications. Furthermore, for privacy of content consumers, mechanisms such as ANDaNA [2] can be applied.

VII. Conclusion

This paper proposes to distinguish ICN traffic into two types for delivering sharable and non-sharable content, where the criteria for shareability depends on several contextual factors. To enable this dual-mode forwarding in CCN, we proposed extended Interest and Data headers addressing efficiency, reliability, and security issues. We demonstrated the benefits of our proposal through simulation analysis. We observed that forwarding non-sharable traffic using only FIB improves CS and PIT efficiency and round-trip time for sharable content, therefore achieving better QoE for end users.

References

[1] Cisco visual networking index: Forecast and methodology, 2009-2014: http://tinyurl.com/3p7v28,2010.

[2] S. DiBenedetto, P. Gasti, G. Tsudik, and E. Uzun. Andana: Anonymous named data networking application. In *Proc. of NDSS*, 2012.

[3] V. Jacobson, D. K. Smetters, N. Briggs, and et al. Voccn: Voice-over content-centric networks. In *Proceedings ACM ReArch*, 2009.

[4] V. Jacobson, D. K. Smetters, Thornton, and et al. Networking named content. In *Proceedings CoNEXT*, 2009.

[5] M. Lacage. Outils d'expérimentation pour la recherche en réseaux, *Ph.D. thesis*, *de l'Université de Nice-Sophia Antipolis*. November, 2010.

[6] C. Tsilopoulos and G. Xylomenos. Supporting diverse traffic types in information centric networks. In *Proceedings of ACM SIGGCOMM, ICN Workshop*, 2011.