



Towards Information-Centric Wireless Multi-hop Communication

Carlos Anastasiades, Torsten Braun

► To cite this version:

Carlos Anastasiades, Torsten Braun. Towards Information-Centric Wireless Multi-hop Communication. 13th International Conference on Wired/Wireless Internet Communication (WWIC), May 2015, Malaga, Spain. pp.367-380, 10.1007/978-3-319-22572-2_27 . hal-01728808

HAL Id: hal-01728808

<https://inria.hal.science/hal-01728808v1>

Submitted on 12 Mar 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Towards Information-Centric Wireless Multi-hop Communication

Carlos Anastasiades and Torsten Braun

Institute of Computer Science and Applied Mathematics
University of Bern,
3012 Bern, Switzerland,
{anastasiades, braun}@iam.unibe.ch

Abstract. Information-centric networking (ICN) addresses drawbacks of the Internet protocol, namely scalability and security. ICN is a promising approach for wireless communication because it enables seamless mobile communication, where intermediate or source nodes may change, as well as quick recovery from collisions. In this work, we study wireless multi-hop communication in Content-Centric Networking (CCN), which is a popular ICN architecture. We propose to use two broadcast faces that can be used in alternating order along the path to support multi-hop communication between any nodes in the network. By slightly modifying CCN, we can reduce the number of duplicate Interests by 93.4% and the number of collisions by 61.4%. Furthermore, we describe and evaluate different strategies for prefix registration based on overhearing. Strategies that configure prefixes only on one of the two faces can result in at least 27.3% faster data transmissions.

1 Introduction

With emerging community networks and social networking services, communication networks have shifted from resource interconnection to content sharing networks. In this context, information-centric networking (ICN) has gained much attention in recent years, because it addresses shortcomings of current host-based communication, namely scalability, flexibility to changing routes and security. In ICN, requests for content do not need to be sent to specific servers but nodes can express Interests and receive matching Data from any nearby node in response. The exchanged messages do not contain any source or destination addresses, which supports caching in any node. This is beneficial in multi-hop communication and communications with multiple requesters, because content does not need to be (re-)transmitted over the entire path but can be retrieved from the nearest cache. Because content is signed, integrity and authenticity of retrieved content is ensured and it is not important which node provided the content copy.

To transmit Interests towards content sources, forwarding tables need to be configured. In fixed ICN Internet communication, prefixes can be configured with the help of routing protocols [15], [7] and nearby copies can be found via cache synchronization [17] or redundant random searches [4]. In wireless networks,

topologies may change and periodic exchange of routing information may overload the network. Existing works [10], [1], therefore, exploit the broadcast nature of the wireless medium. Requests are flooded and content prefix are overheard to learn distances to available content sources similar to existing routing protocols like AODV [12]. To learn distances to content sources, endpoint identifiers and hop distances are added to all messages. However, when adding endpoint identifiers to messages, communication is not strictly information-centric anymore, which means that it loses its flexibility because it matters from which node content is retrieved (endpoint specified in messages). As a consequence, complicated handoff mechanisms are required as in current host-based communication. Current works in wireless information-centric communication do not consider forwarding tables to route Interests and either consider only one-hop communication to wired infrastructure [16], [6] or consider routing on a high-level [14] with endpoint identifiers [10], [1].

In this work, we investigate a different approach without endpoint identifiers. Nodes overhear broadcast content transmissions to configure prefixes in the forwarding tables. Routing can then be based on the same forwarding structures as for wired information-centric networks, which facilitates interoperability. We describe and evaluate different strategies for prefix registration and evaluate optimizations to improve multi-hop communication in case of multiple potential forwarders.

The remainder of this paper is organized as follows. Related work is described in Section 2. In Section 3 we describe overhearing and forwarding strategies to support opportunistic content-centric multi-hop communication. Evaluation results are shown in Section 4 followed by a discussion in Section 5. Finally, in Section 6, we conclude our work and give an outlook on future work.

2 Content-Centric Networking

2.1 Basic CCN Concepts

We base our investigations on the Content-Centric Networking (CCN) architecture [8]. It is based on two messages: *Interests* to request content and *Data* to deliver content. Files are composed of multiple segments, which are included in a Data message, and users need to express Interests in every segment to retrieve a complete file. CCNx [3] provides an open source reference implementation of CCN. The core element of the implementation is the CCN daemon (CCND), which performs message processing and forwarding decisions. Links from the CCND to applications or other hosts are called *faces*. A CCND has the following three memory components:

1. The Forwarding Information Base (FIB) contains forwarding entries to direct Interests towards potential content sources. Every entry contains a prefix, the face where to forward Interests and a lifetime value. To avoid loops, an Interest is never forwarded on the same face from where it was received.

2. The Pending Interest Table (PIT) stores pending forwarded Interests together with the face on which they were received. If Data is received in return, it can be forwarded based on face information in the PIT. Existing PIT entries prevent forwarding of similar Interests. Duplicate Interests are identified by a Nonce in the Interest header. Therefore, PIT entries form a multicast tree such that each Interest is only forwarded once at a time over a certain path.
3. The Content Store (CS) is used as cache in a CCN router storing received Data packets temporarily. It is used as short-term storage that can be used to reduce delays for retransmissions in multi-hop communication.

Prior to transmission, content is included and scheduled for transmission in a queue. To avoid duplicate content transmissions, an additive random delay (`data_pause`) is considered during content scheduling. If a node overhears the transmission of the same content from another node, it removes the scheduled content from the content queue (duplicate suppression). The *freshnessSeconds* field in content headers specifies how long a segment stays in the CS after its reception. Content can also be stored persistently in repositories of content sources. The *AnswerOriginKind* in Interest headers specifies from where the content should be retrieved, e.g., content store or repository.

2.2 CCN Forwarding and Routing

OSPFN [15] and NLSR [7] are routing protocols for large multi-hop networks. Link state advertisements (LSAs) are propagated throughout the entire network and each router builds a complete network topology. These protocols, therefore, target static networks without frequent topology and content availability changes to avoid many LSAs. To detect routes to content sources closer to the requester, which may not be on a configured path to the original content source, optimizations are required. Bloom filters [9], [17] have been proposed to exchange caching information among neighbors. However, nodes would need to guarantee minimal cache validity times in order to be useful. Other optimizations [5], [4] use an exploration phase, where Interests are flooded into the network to find new content sources closer to the requester. In [18] the concept of adaptive NDN Interest forwarding is investigated. Based on successful Data retrievals, forwarding nodes can select which FIB face to use in the next transmission. Routers inform downstream nodes if they cannot forward Interests due to missing FIB entries by Interest NACKs. Therefore, the approach targets unicast networks and requires the configuration of FIB entries in the first place.

Earlier works [10], [1] in wireless ICN communication use flooding to learn content reachability via additional information in message headers. The Listen First, Broadcast Later (LFBL) [10] algorithm adds additional header fields to both Interest and Data messages including the communication end-points and hop distances. Nodes perform flooding if they do not know endpoints and learn distance information via overhearing. Forwarding is based on hop distance to the endpoints. E-CHANET [1] is a similar approach, which also adds endpoint

(providerIDs) and distance (hop count) information to messages. To avoid multiple content sources responding to Interests, not only content names but also providerIDs in Interest and Data messages need to match. Mobility handoffs and provider selection are proposed to enable communication to the “closest” content source. However, both approaches circumvent one of ICN’s fundamental design goals of avoiding endpoint identifiers. Aggregating requests for the same content becomes much more challenging in case of multiple redundant content sources. Additionally, Data messages, which are returned from cache, may hold imprecise provider and distance information. Furthermore, forwarding based on hop distance, i.e., using only the distance as path selection criterion, may not work efficiently in dense environments if multiple nodes may have the same “hop distance” to the endpoints. Both approaches are similar to AODV to learn distances to content sources, but they do not investigate name-based routing. In this work, we describe a different approach without endpoint identifiers, where content prefixes are configured in the FIB. This enables Interest aggregation to support multiple concurrent requests.

3 Multi-hop Communication with Overhearing

Information-centric multi-hop communication requires a mechanism to configure forwarding prefixes in the FIB and selecting forwarders out of multiple potential options to avoid duplicate transmissions. We consider the extraction of content names in overheard broadcast Data packets and include it temporarily into the FIB such that Interests can be forwarded towards the content source. To enable overhearing without configured FIB entries (and in the absence of other nodes requesting content), Interests from local applications can always be forwarded to a broadcast face to probe for content in the environment. We call this *pass-through* of local Interests. Interests from other nodes are only forwarded if a matching FIB entry exists. Interests forwarded via pass-through always ask for content from repositories but not from caches. Additionally, we add a hop counter to Interest messages to limit the propagation of Interests. A hop counter in Interest messages is also compliant with the new CCNx 1.0 protocol [11].

3.1 Prefix Registration and Forwarding Strategies

Since Interests can not be forwarded on the face from where they have been received, multi-hop communication requires at least two (broadcast) faces for communication: one for receiving and one for transmitting / forwarding messages. To receive messages on either face, both broadcast faces need to be configured. In addition, to forward Interests towards content sources, eligible prefixes need to be configured to those faces in the FIB. If a node receives Interests on one face, they should be forwarded on the other face.

We describe three different strategies for prefix registration below. All forwarding strategies require the configuration of two broadcast faces to receive content. The difference between the strategies is how prefixes are associated to those faces.

Two Static Forwarding Faces (2SF): Every prefix needs to be configured on both broadcast faces such that Interests can be forwarded on the opposite face on which they were received. In this strategy, we do not apply any overhearing mechanisms. All prefixes need to be configured to both broadcast faces in an initial phase because no feedback can be received from the network. This may result in large FIB tables to route all potential prefixes (even if the corresponding content is never requested). In addition, it is more difficult to adapt to changing topologies since routing protocols are required for periodic updates. Please note that we use 2SF only for reference purposes.

Two Dynamic Forwarding Faces (2DFo): In this strategy, overheard prefixes are added to both broadcast faces similar to *2SF*. However, in contrast to *2SF*, prefixes are only added dynamically to the FIB if the content is available, i.e., it has been overheard. The approach uses pass-through and hop counters. Pass-through works as follows: if a request on face 1 has not been successful, the Interest is automatically transmitted via face 2. To avoid that forwarders discard it as duplicate, the requester needs to change the nonce in the retransmitted Interest.

One Dynamic Forwarding Face (1DFo): In contrast to *2DFo*, overheard prefixes are only registered on the face on which content transmission has been overheard, but not on both faces. Similar to *2DFo*, pass-through and a hop counter is used.

We explain forwarding with the help of Figure 1, which shows four nodes and their transmission ranges (dashed circles). Nodes in transmission range of the content source can broadcast an Interest on face 1 (pass-through) and all nodes in range of the content source can overhear Data transmission to configure the prefix in the FIB, i.e., they can configure the prefix on face 1 denoted by the solid grey circle in the figure. Nodes further away need to send Interests on face 2 so that nodes near the content source can forward it via face 1. Since they receive Data on face 2, they can configure the prefix on face 2 denoted by the white circular ring. As shown in the figure, only two broadcast faces are required to forward Interests from a requester multiple hops away by alternating the forwarding faces in the nodes along the path.

3.2 Configure Prefixes via Overhearing

Prefixes are configured by extracting the content name from overheard (broadcast) content transmissions. Only new content is registered in the FIB, i.e., only if the content is not already stored in the content store or the local repository. To limit message processing, only every n -th content message can be processed. By applying modulo operations on the received segment number, no additional state information is required.




Fig. 1: Forwarding via alternating faces.

If the 2DFo strategy is used, the prefix is registered to both broadcast faces. If the 1DFo strategy is used, the prefixes are only registered to the face from which the content has been received and to 1 face at maximum. If no prefixes are configured, Interests from local applications can be forwarded via pass-through. If an Interest needs to be forwarded via pass-through, the header field *AnswerOriginKind* is set to 0, which means that no cached answer from content store is accepted.

3.3 Interest and Data Forwarding

To avoid duplicate Interest (and Data) transmissions, the propagation of Interest messages needs to be controlled. The basic idea is the following: every node delays Interest transmissions randomly by an *Interest Forwarding Delay (IFD)*. Based on overheard Interest and Data messages, the IFD is increased or decreased. If a forwarder continuously receives Data to its transmitted Interests, it assumes to be a *preferred forwarder*.

Interest Forwarding Delay: To avoid duplicate Interest transmissions, every Interest is forwarded with a delay, which is randomly selected within a specified interval $[0, IFD_{max}]$. Based on overheard Interest or Data messages, IFD_{max} is modified. If Interests are transmitted and Data is received in return, IFD_{max} is decreased. By this, we can minimize the influence of large forwarding delays and reduce the time until the next Interest can be sent resulting in higher throughput. If the same Interest is overheard from another node before it has been transmitted, the Interest transmission is delayed once, i.e., *delayed send*, and the interval is slightly increased. If another Interest is overheard during that time, the Interest is discarded. The values are configurable but in our current implementation, we set IFD_{max} initially to 100ms. We halve IFD_{max} if a Data message is received in return and no other Interest has been overheard. If an Interest has been overheard IFD_{max} is slightly increased by 1ms.

Preferred Forwarders: Interests are forwarded based on the IFD. At some point, a forwarder may receive a Data message for which no Interest has been forwarded locally. Then, the forwarder knows that there is another node that forwarded the Interest faster. It assumes that it is *non-preferred* and adds a fixed delay (in our implementation 100ms) to every Interest transmission. Due to the fixed delay, non-preferred forwarders will only attempt to forward Interests if preferred forwarders have not retrieved content already. If preferred forwarders have moved away or their Interests collide, non-preferred forwarders can forward their Interests (*delayed send*). To avoid unnecessary switching, e.g., due to occasional collisions, a non-preferred forwarder can only become preferred, if it has performed a delayed send for N (in our implementation: N=3) subsequent times.

Data Transmission: Data is transmitted similarly to CCN with an additive random delay and Interests can be satisfied from cache. However, there is one modification: Data messages are only forwarded if the corresponding Interest has been forwarded and not only scheduled for transmission.

3.4 Data Structures

Our proposed approach uses the existing CCN data structures, namely the FIB, the PIT and the CS. For efficient multi-hop communication, we use the following additional hash tables:

Interest Table (IT): The IT points to the corresponding PIT entry and includes additional information for Interest forwarding. The information includes the number of times the Interest has been overheard, a list of faces where the Interest has been forwarded and multiple flags to specify whether: 1) a pass-through was required, 2) the Interest has been sent or only scheduled for transmission and 3) a delayed send was required due to an overheard Interest. An IT table entry is created whenever an Interest is received and forwarded. Cleanup can be performed at the same time as for the PIT table, i.e., when an Interest has been consumed or expired depending on the Interest lifetime.

Content Flow Table (CFT): The CFT is used to maintain the current IFD_{max} value and to remember whether it is a preferred node or not. It is created the first time when an Interest has been forwarded for an existing FIB entry. Similar to the FIB entries, CFTs have an expiration time, which can be regularly updated and checked at the same time as the FIB expiration. In our implementation, we set the expiration time to 500s but also larger values are possible. CFT entries keep track of existing forwarding faces and unsuccessfully forwarded Interests: if S (in our implementation: S=16) subsequent Interest transmissions have been detected over a specific face, the corresponding dynamically created FIB entries can be unregistered and removed from the content flow. If the content flow does not specify another outgoing face, the content flow is deleted.

4 Evaluation

We performed evaluations by simulations using our CCN framework based on OMNeT++ [13]. The simulation framework is based on CCNx and provides a complete CCN implementation including all memory components such as the Content Store (CS), Pending Interest Table (PIT), Forwarding Information Base (FIB) as well as content queues with the same delays and flags like CCNx. The framework allows us to process and store information as in CCNx (except signature calculations and verifications) but facilitates network setup, MAC layer observations and enables high scalability.

4.1 Simulation Scenario

The evaluation topology is shown in Figure 2. The small solid circles represent network nodes and the dashed circles denote different network partitions. All network nodes inside a network partition can directly communicate with each other. The white solid circles are only part of one partition and can not directly reach nodes in other partitions. Black solid circles in the intersections between partitions are potential forwarding nodes that can relay between white solid nodes. In total, there are 150 nodes: 25 nodes within each partition (white circles) and 10 potential forwarders (black circles) in each intersection A, B, C, D, M of neighboring partitions. The content source is always placed in partition 1 and a requester (white node) in partition 1 requests content such that forwarders (black nodes) in the intersections A, B and M can overhear it. This step is only used to populate the FIB in partition 1. The common evaluation parameters are listed in Table 1.

All network nodes have one 802.11g network interface and two broadcast faces configured to enable multi-hop communication. The broadcast data rate is set to 2Mbps. The segment size corresponds to the effective CCN payload without CCN headers and the content lifetime defines how long content stays




Fig. 2: Network with four partitions: all nodes in the same partition can see each other but not nodes from other partitions.


Parameter	Value
interface	1x 802.11g
data rate	broadcast: 2Mbps
segment size	4096 Bytes (RTS/CTS disabled)
content size	4 MB (1000 segments)
pipeline size	16
content lifetime CL	600s
data_pause	100ms
request interval RI	2000s, 30s
Interest lifetime	4s (default)
prefix validity time in FIB	40000s

Table 1: Common simulation parameters.

valid in the content store, i.e., corresponding to *freshnessSeconds* in CCNx. We set data_pause, i.e., a random CCN delay for every Data transmission, to 100ms. We evaluate multiple requests from different partitions and the request interval denotes the time interval between the requests. For a request interval of 2000s, content cannot be found in the nodes’ caches anymore. To evaluate competing concurrent flows, where part of the content can be found in the cache, we also evaluate a request interval of 30s. Requesters are always selected among the ”white nodes” within the corresponding partition to enforce multi-hop communication. We use the default CCNx Interest lifetime of 4 seconds. For the 1DFo, 2DFo and 2SF strategies, message processing is implemented as described in Section 3. The validity time of overheard prefixes in the FIB can be much larger than content lifetimes in the cache because name prefixes are significantly smaller than the content objects (multiple segments with payloads). Therefore, we set the validity time of overheard prefixes in the FIB to 40000s but delete entries after 16 (consecutive) unsuccessful Interest forwardings. For every configuration, 100 simulation runs have been performed.

4.2 Multi-hop Communication with Interest Table

In this subsection, we evaluate the performance gain of IT/CFT compared to regular CCN message processing. To populate the FIB entries for the 1DFo and 2DFo strategies, a requester in partition 1 requests content from a content source in the same partition. After 2000s, a requester in partition 2 is requesting the same content. Although the content does not exist in caches anymore, the FIB entries in intersections A and M are still valid. In Figure 4a, we compare 1DFo, 2DFo and 2SF in terms of transmission times and transmitted Interests. The file size is 4MB, which corresponds to 1000 segments. The left y-axis shows the requester’s transmission time in seconds and the right y-axis the transmitted Interest messages by listeners, i.e., nodes that do not actively request content, in partition 2. Figure 4a shows that the 1DFo strategy results in 6.4% shorter transmission times than 2DFo and 8% shorter transmission times than 2SF with



(a) Requester's Transmission Time and Listeners' transmitted Interests in partition 2 (b) Duplicate Interests and Collisions at Content Source in partition 1

Fig. 3: Comparing CCN to CCN with IT/CFT.

regular CCN. However, when applying IT/CFT, the transmission times of 1DFo can be drastically reduced by 47%. The differences compared to other strategies become larger with applied IT/CFT, i.e., 1DFo performs 27.3% faster than 2DFo and even 38.2% faster than 2SF.

1DFo performs better than 2DFo and 2SF because all "white" nodes in partition 2 have only one face configured. Because Interests can not be forwarded on the face of reception, Interest forwarding in partition 2 could be effectively avoided with 1DFo. As Figure 4a shows for the 1DFo strategy, listener nodes in partition 2 that do not have a direct connection to the content source do not forward Interests but only forwarders (black nodes in Figure 2) do. When applying IT/CFT, transmitted Interests by listeners could be reduced for 2DFo and 2SF because listeners realize that they are non-preferred nodes. However, because forwarding via two faces results in more forwarded Interests and overheard Interests, Interest forwarding delays IFD_{max} of 2DFo and 2SF are higher than with 1DFo, which explains the longer transmission times.

Figure 3b shows the number of received duplicate Interests and collisions at the content source in partition 1. With regular CCN, the number of duplicate Interests is high because all forwarders in intersections A and M receive the Interests at the same time and forward it almost immediately before overhearing forwardings of other nodes. In the 1DFo strategy, every forwarder transmits on average 636 Interests, while with the 2DFo and 2SF strategy 934 Interests are forwarded by forwarders and 440 Interests by listener nodes in partition 1. As a result, a content source receives 48.2% fewer duplicate Interests and experiences and 61.4% fewer collisions with the 1DFo strategy compared to 2DFo and 2SF. With IT/CFT, the number of duplicate Interests can be drastically reduced by at least 93.4% because non-preferred forwarders transmit only a few Interests, but duplicate transmissions cannot be avoided completely.

We would like to emphasize that multiple forwarders are beneficial for multi-hop forwarding with IT/CFT. Recall that forwarders delay Interest transmissions if they overhear the same Interest. Because these forwarders can transmit Interests via *delayed send* if they do not receive content, the communication can recover faster from collisions on the path. For example, in the current scenario with 20 potential forwarders in intersections A and M, content can be retrieved 20.7% faster with 1DFo compared to a scenario with only 1 forwarder.

4.3 Multi-hop Communication with Multiple Requesters

In this subsection, we evaluate the influence of multiple concurrent requesters to multi-hop communication. We only show evaluation results with IT/CFT because it performed better than regular CCN. In addition, we only show evaluation results with concurrent requesters in different partitions, i.e., partitions 2, 3, 4. Concurrent requesters in the same partition would receive the content without any efforts due to broadcast communication and caching. Similar to the last subsection, a requester in partition 1 requests content such that FIB tables at forwarders are filled. The concurrent requests start after a delay of 2000s to ensure that no content is cached. Then, requests in partitions 2, 3, 4 (in this order) start with an interval of 30s, which means that some content can be retrieved from cache but requesters quickly catch up resulting in concurrent request streams.




Fig. 4: Performance of Concurrent Requests.

Figure 4a shows the transmission time of a requester in partition 2 if there is one concurrent stream from partition 2, two concurrent streams from partition 2 and 3 or three concurrent streams from partitions 2, 3 and 4. Figure 4a shows that a requester in partition 2 needs 22.8% more time if there is a concurrent stream in partition 3. Please note that only nodes in intersection M can receive Interests from both partitions 2 and 3, while nodes in intersections A only receive Interests from partition 2 and nodes in intersection B only from partition 2. Two concurrent requesters in partition 2 and 3 can, therefore, result in two competing streams, which may result in slightly longer Interest forwarding delays IFD_{max} and consequently, slightly longer transmission times. Another requester

in partition 4 does not further increase the transmission time with 1DFo. Thus, the increase in transmission time does not grow linearly with increasing number of requesters. All concurrent requesters finish their transmission at the same time, which means that effective transmission times of requester 2 and 3 are approximately 30s and 60s shorter as Figure 4b shows.

Figure 4c shows the Interests transmitted by the forwarders in intersections A, B, M on average when using the 1DFo, 2DFo and 2SF strategy with three requesters in partition 2, 3 and 4. The bars on the left side show the transmitted messages if the requests are subsequent with an interval of 2000s such that requesters cannot profit from each other. The bars on the right side show the transmitted messages if requesters concurrently retrieve content in an interval of 30s. Figure 4c shows that with 1DFo, the transmitted messages can be reduced by 64.0%, i.e., to almost a third of the traffic of three requesters. This shows that 1DFo can efficiently handle concurrent multi-stream communication. For 2DFo and 2SF, the reduction is smaller because forwarders can forward Interests on two faces. Forwarding in CCN is continuously updated based on past experience and every node forwards Interests first over the face that it considers best. Different nodes may have different views on which face is "best". If requesters in partition 2 and 3 transmit Interests on different faces, they are occasionally both forwarded by intermediate nodes. Therefore, the reduction is only 46.74.5% with 2DFo and only 54.73% with 2SF. Please note that the absolute number of transmitted Interests is a slightly larger with 2DFo compared to 2SF because IDF_{max} is slightly larger with 2SF.

Figure 4d shows the transmitted Data messages by the content source in partition 1 for subsequent and concurrent requests. The number of transmitted Data messages is proportional to the number of received Interest messages at the content source. Figure 4d shows that all three strategies can limit the number of transmitted Data messages at an approximately constant level. The 2DFo strategy has the highest increase of transmitted Data messages of 9.1% from 1 to 3 requesters. For 2SF, the increase is 5.6% and for 1DFo, it is only 3.3%.

5 Discussion

It would be possible to use only one forwarding face and enable FIB forwarding on the same face of reception. However, this would have implications on PIT processing because Interests would be forwarded on the same face from where they were received. To avoid inefficiencies, the number of Interest forwardings would need to be limited. If each received Interest could only be forwarded once (minimum for multi-hop), it would be similar to our 2DFo strategy. Our evaluations have shown that 1DFo outperforms 2DFo, which justifies the use of two forwarding faces, which may be artificially created on the same wireless interface.

One may argue that nodes need to overhear content (for prefix registration) before Interests can be forwarded over multiple hops towards a content source. However, the same applies to existing approaches that use endpoint identifiers

[10], [1], i.e., each intermediate node needs to know the content source in order to forward content. To address this issue, we see three approaches. First, one could advertise prefixes, but in contrast to current routing protocols [15],[7], nodes should not learn the entire topology but only the immediate neighborhood via broadcast. Second, one could add an additional message flag to enable flooding over a limited number of hops (pass-through) if no forwarding entry is configured. Third, agent-based content retrieval [2] may be used for content retrieval in delay-tolerant networks over multiple nodes.

The validity time of overheard prefixes in the FIB is still subject to more investigations and may depend on network dynamics. We believe that negative acknowledgements (NACKs) as proposed in [18] to indicate unavailability of content would not be a good option in broadcast environments: although one node can not forward the Interest, another might be able to do so. Listening to NACKs from single nodes may, therefore, not be meaningful. Therefore, in this work, we tracked content responses for transmitted Interests and deleted forwarding entries if a consecutive number of Interests has timed out.

6 Conclusions and Future Work

In this work, we have described a new approach for information-centric wireless multi-hop communication based on overhearing. In contrast to existing approaches, no endpoint identifiers are used but overheard prefixes are registered in the FIB. Evaluations have shown that the current CCN architecture does not work well for wireless multi-hop communication if multiple paths (forwarders) are available. Every node that receives a message and has a forwarding entry, will forward it immediately resulting in many duplicate transmissions. Additional data structures to maintain Interest forwarding delays can reduce the number of duplicate Interest transmissions by 93.4% and reduce the number of collisions by 61.4%.

We have implemented two dynamic prefix registration and forwarding strategies based on overhearing, i.e. 1DFo and 2DFo, and compared it to a static strategy, i.e., 2SF. Evaluations have shown that 1DFo outperforms both 2DFo and 2SF in terms of both transmission times and transmitted messages. The 1DFo strategy results in 27.3% faster transmission times than 2DFo and even 38.2% faster transmission times than 2SF. Interest aggregation and caching works efficiently even in case of multiple concurrent streams. However, Interest aggregation is more efficient with 1DFo than with 2DFo or 2SF because Interests are forwarded over the same faces. As a consequence, the number of transmitted Data messages stay approximately constant for an increasing number of requesters.

The main challenge of the described approach is the fact that content needs to be overheard to configure forwarding prefixes. While this works well for popular content, it may be more difficult for unpopular content that is requested infrequently. As part of our future work, we will investigate options to address unpopular content in wireless multi-hop environments.

References

1. Amadeo, M., Molinaro, A., Ruggeri, G.: E-CHANET: Routing, forwarding and transport in Information-Centric multihop wireless networks. *Computer Communications* 36, 792–803 (2013)
2. Anastasiades, C., El Maudni El Alami, W., Braun, T.: Agent-based Content Retrieval for Opportunistic Content-Centric Networks. In: 12th WWIC. pp. 175–188. LNCS 8458, Paris, France (May 2014)
3. CCNx: <http://www.ccnx.org/> (April 2015)
4. Chiocchetti, R., Perino, D., Carofiglio, G., Rossi, D., Rossini, G.: INFORM: a dynamic interest forwarding mechanism for information centric networking. In: 3rd ACM SIGCOMM ICN. pp. 9–14. Hong Kong (August 2013)
5. Chiocchetti, R., Rossi, D., Rossini, G., Carofiglio, G., Perino, D.: Exploit the known or explore the unknown? hamlet-like doubts in icn. In: ACM SIGCOMM ICN. pp. 7–12. Helsinki, Finland (August 2012)
6. Grassi, G., Pesavento, D., Pau, G., Vuyyuru, R., Wakikawa, R., Zhang, L.: VANET via Named Data Networking. In: IEEE INFOCOM NOM. pp. 410–415. Toronto, Canada (April-May 2014)
7. Hoque, M., Amin, S.O., Alyyan, A., Zhang, B., Zhang, L., Wang, L.: NLSR: Named-data Link State Routing Protocol. In: 3rd ACM SIGCOMM ICN. pp. 15–20. Hong Kong (August 2013)
8. Jacobson, V., Smetters, D.K., Thornton, J.D., Plass, M.F., Briggs, N.H., Braynard, R.L.: Network Named Content. In: 5th ACM CoNEXT. pp. 1–12. Rome, Italy (December 2009)
9. Lee, M., Cho, K., Park, K., Kwon, T., Choi, Y.: SCAN: Scalable Content Routing for Content-Aware Networking. In: IEEE ICC. pp. 1–5. Kyoto, Japan (June 2011)
10. Meisel, M., Pappas, V., Zhang, L.: Listen First, Broadcast Later: Topology-Agnostic Forwarding under High Dynamics. In: ACITA. pp. 1–8. London, UK (September 2010)
11. Mosko, M.: Ccnx messages in tlv format (January 2015), <http://tools.ietf.org/html/draft-mosko-icnrg-ccnxmessages-00>, Internet Draft
12. Perkins, C.E., Royer, E.M.: Ad-hoc On-Demand Distance Vector Routing. In: 2nd IEEE WMCSA. pp. 90–100. New Orleans, LA, USA (February 1999)
13. Varga, A.: The OMNeT++ Discrete Event Simulation. In: ESM. Prague, Czech Republic (June 2001)
14. Varvello, M., Rimac, I., Lee, U., Greenwald, L., Hilt, V.: On the Design of Content-Centric MANETs. In: 8th WONS. pp. 1–8. Bardonecchia, Italy (January 2011)
15. Wang, L., Hoque, M., Yi, C., Alyyan, A., Zhang, B.: OSPFN: An OSPF Based Routing Protocol for Named Data Networking. Tech. rep., NDN Technical Report (July 2012)
16. Wang, L., Wakikawa, R., Kuntz, R., Vuyyuru, R., Zhang, L.: Rapid traffic information dissemination using named data. In: ACM NOM. pp. 7–12. Hilton Head, SC, USA (June 2012)
17. Wang, Y., Lee, K., Venkataraman, B., Shamanna, R., Rhee, I., Yang, S.: Advertising Cached Contents in the Control Plane: Necessity and Feasibility. In: IEEE INFOCOM NOMEN. pp. 286–291. Orlando, FL, USA (March 2012)
18. Yi, C., Afanasyev, A., Wang, L., Zhang, B., Zhang, L.: Adaptive forwarding in named data networking. *ACM SIGCOMM Computer Communication Review* 42(3), 62–67 (July 2012)