

NetCodCCN: a Network Coding approach for Content-Centric Networks

Jonnahtan Saltarin*, Eirina Bourtsoulatze*, Nikolaos Thomos[†] and Torsten Braun*

*University of Bern, Bern, Switzerland

Email: {saltarin,braun}@inf.unibe.ch, eirina.bourtsoulatze@gmail.com

[†]University of Essex, Colchester, United Kingdom

Email: nthomos@essex.ac.uk

Abstract—Content-Centric Networking (CCN) naturally supports multi-path communication, as it allows the simultaneous use of multiple interfaces (e.g. LTE and WiFi). When multiple sources and multiple clients are considered, the optimal set of distribution trees should be determined in order to optimally use all the available interfaces. This is not a trivial task, as it is a computationally intense procedure that should be done centrally. The need for central coordination can be removed by employing network coding, which also offers improved resiliency to errors and large throughput gains. In this paper, we propose NetCodCCN, a protocol for integrating network coding in CCN. In comparison to previous works proposing to enable network coding in CCN, NetCodCCN permit Interest aggregation and Interest pipelining, which reduce the data retrieval times. The experimental evaluation shows that the proposed protocol leads to significant improvements in terms of content retrieval delay compared to the original CCN. Our results demonstrate that the use of network coding adds robustness to losses and permits to exploit more efficiently the available network resources. The performance gains are verified for content retrieval in various network scenarios.

I. INTRODUCTION

In the IP protocol, core of the current Internet architecture, each packet is routed based on the location of the host to which it is addressed. However, nowadays Internet users care more about the content they want to obtain rather than where this content is stored. To address this mismatch, Jacobson *et al.* [1] proposed Content-Centric Networking (CCN), a new communication paradigm in which the importance is shifted from *where* the content is located, to *what* the content is. In the CCN model, the content is described by its *name* and the users demand content with the help of Interest messages that contain the name of the requested content. These Interests are transmitted over the network until they reach a node holding a copy of the content whose name matches that of the Interest message. Once such a node is reached, a copy of the requested content is encapsulated in a Data message and it is sent back to the requester following the reverse path of that followed by the Interest message. As the Data message is transmitted backwards to the requester, intermediate nodes can store copies of it, so they can reply to future Interests for the same content.

One of the advantages of CCN is that it allows clients to exploit multiple paths in a native way, as clients can simultaneously transmit Interests over all their network interfaces (e.g., LTE and WiFi) to retrieve the content segments that

comprise the requested content. This leads to a better use of the network resources and reduces the time needed to collect all the content segments. However, when multiple clients are interested in the same content (e.g., a popular video stream), and/or when the content is distributed across multiple sources (e.g., in a distributed storage system), the optimal content delivery rate is only attained if the segments are delivered over the optimal set of multicast trees [2]. This means that the Data and Interest messages should be transmitted over these multicast trees, thus, the nodes need to know where they need to forward each Interest to follow these multicast trees, which does not scale for large and dynamic topologies. Furthermore, the computation of the optimal set of multicast trees needs a central entity that is aware of the network topology, which is hard to be done in dynamic networks. An alternative solution to the computation of the optimal multicast trees is to use network coding [3]. With network coding all the network nodes perform coding operations on the received packets instead of just replicating and forwarding them as in traditional networks. The receivers decode the information when they receive a decodable set of packets, *i.e.*, as many linearly independent coded packets as the number of source data segments.

The application of network coding in CCN has been explored in [4] where the NC3N architecture has been introduced. In this approach, Interests contain information about the content segments available at the client, based on the approach proposed in [5]. Nodes holding content segments that matches the name prefix of the Interest, reply only if they can provide a network coded content segment that increases the content available at the client. However, in the presence of multiple clients, (i) the aggregation of Interests is problematic, since Interests for the same content segment from different clients contain different content availability information; and (ii), when a client sends multiple Interests in parallel to receive multiple content segments, it will include the same information about the content that it already has. This is undesirable as a node that has a matching content segment will reply to multiple Interests with the same content segment, that will be duplicated for the client. Inspired by [4], CodingCache has been proposed in [6] which uses network coding to replace the content segments in the cache of the network nodes. Due to the increased content segment diversity in the network, the

cache hit rate is improved. However, this approach suffers from the same drawbacks as the architecture in [4]. In [7], the multicast delivery of network coded content in Information-Centric Networks is optimized by finding the evolution of the content segments that are stored in the network. The drawback of this approach is that it does not scale well with the number of network nodes, because it needs a central entity that is aware of the network topology and the clients' requests.

In this paper we propose NetCodCCN, a protocol for integrating network coding in CCN. Our proposed solution solves the shortcomings of the approaches presented in [4] and [6]. Specifically, (i) we eliminate the need to include in the Interests the information about the content available at the client, thus, simplifying Interest aggregation; (ii) we allow nodes to keep information about the content segments they have sent on each face, reducing the number of duplicate segments; and (iii) we allow clients to send multiple Interests in parallel, by modifying the way in which the nodes process the Interest messages.

We have implemented NetCodCCN by making the necessary changes to CCNx [8], and performed experiments to compare it to unmodified CCNx. Our results demonstrate that NetCodCCN offers large gains in terms of the time needed to retrieve the original content object. Moreover, it adds robustness to losses and permits to exploit more efficiently the available network resources in multi-source multicast scenarios. To the best of our knowledge, this is the first practical implementation of network coding integrated into CCNx.

II. DATA RETRIEVAL IN CCN

We focus on content communication over wired networks represented by directed acyclic graphs $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where \mathcal{V} and \mathcal{E} denote the set of network nodes and the set of links connecting them, respectively. Each network consists of a set of source nodes \mathcal{S} that generate and/or store content objects, a set of clients \mathcal{U} that demand content objects and a set of intermediate nodes \mathcal{R} through which the content objects are requested and transmitted. Hence, we have $\mathcal{V} = \mathcal{S} \cup \mathcal{U} \cup \mathcal{R}$, where every node $v \in \mathcal{V}$ is connected with its neighboring nodes through a set of faces \mathcal{F}_v .

In CCN, content objects are split into smaller segments that fit into Data messages. Each segment is uniquely identified by a name. We denote a content object as $C_p = \{c_{p,1}, \dots, c_{p,N}\}$ where N is the number of segments in C_p and p is the name of the content object, which serves as a name prefix for the segments. The name of each segment $c_{p,n} \in C_p$ is generated by appending the name of the content object p with the segment identifier n . For instance, the name of the segment $c_{p,1}$ is `/provider/videos/largevideo.h264/1`, where `/provider/videos/largevideo.h264` is the name prefix, p and 1 is the segment id.

Each source $s \in \mathcal{S}$ stores content objects that can be requested by the clients. A client $u \in \mathcal{U}$ that is interested in a content object $C_p = \{c_{p,1}, \dots, c_{p,N}\}$ should send a set of Interest messages $I_p = \{i_{p,1}, \dots, i_{p,N}\}$, one for each segment. These interests are sent over a set of faces \mathcal{F}_u^p that are

configured to forward Interests for content with name prefix p . The information about which faces a node can use to send Interests for specific name prefixes is stored in the *Forwarding Information Base (FIB)* table.

In CCN, each node $v \in \mathcal{V}$ has a cache, or *Content Store (CS)* in CCN terminology, where segments that pass through the node can be stored. These segments can be used later to reply to Interests for segments with a matching name. Therefore, a node $v \in \mathcal{I} \cup \mathcal{S}$ holding a copy of the segment $c_{p,n}$ in its CS will reply to any Interest $i_{p,n}$. If the CS of node v does not contain a segment matching the name of the Interest $i_{p,n}$, the node v first checks its *Pending Interest Table (PIT)*, that keeps track of the Interests forwarded by the node and all the faces over which those Interests have arrived. If the node v finds in its PIT an entry that matches the name in the Interest, it knows that it has already forwarded $i_{p,n}$ and hence the segment $c_{p,n}$ is expected. In this case v does not forward $i_{p,n}$ again, but only adds the face f over which the Interest has arrived to the respective PIT entry. When the PIT does not have any entry that matches the Interest $i_{p,n}$, the node v forwards the Interest to its neighboring nodes over the set of faces \mathcal{F}_u^p configured in its FIB.

Once the requested segment is found in the CS of an intermediate node or in a source node, it is transmitted to the client in a Data message over the reverse path of that followed by the Interest. When a node v receives a Data message with the segment $c_{p,n}$ over a face $f \in \mathcal{F}_v$, it first checks its CS. If a segment with the same name exists, the arrived segment $c_{p,n}$ is considered duplicated and is not transmitted further. If there is no matching segment in the CS, the node checks the PIT for an entry that matches the name of the segment $c_{p,n}$. If there is no matching PIT entry, the segment $c_{p,n}$ is considered unsolicited and it is discarded. If a matching PIT entry is found, the segment is forwarded to all the faces specified in the corresponding PIT entry. Additionally, the segment $c_{p,n}$ may be added to the CS, according to the caching policy.

III. TOWARDS NETWORK CODING ENABLED CCN

In this section we describe the benefits that network coding can bring to CCN. First, we motivate the use of network coding by presenting three scenarios in which CCN does not perform efficiently. Then, we show how network coding can alleviate the drawbacks of CCN in the mentioned scenarios, while also bringing additional benefits.

A. Motivation

Nowadays, communication devices usually come with multiple network interfaces that can be used to gather content, e.g., smart-phones usually have WiFi and 3G/LTE interfaces. However, in the traditional host-centric networking, using multiple interfaces in parallel to retrieve content is a difficult task, as end-to-end connections need to be established for each interface. In CCN, multipath content retrieval is naturally supported, as the clients can distribute all the Interest messages needed to retrieve a content object over all its available faces.

However, there are some scenarios where CCN does not provide efficient support for multipath content retrieval:

- *Multi-source unicast*: Let us consider the case illustrated in Fig. 1a, where a client u is interested in a content object C_p composed of N segments that are distributed across multiple sources \mathcal{S} . In this case, the client and the intermediate nodes need to select properly the face over which they send the Interest for each segment, such that it reaches the right source. This is done using the information stored in the FIB table. However, keeping the FIB tables of all the nodes updated for each segment of C_p does not scale well, in particular in large networks and in the presence of unreliable sources that can become available/unavailable at any moment.
- *Single-source multicast*: Let us now consider the case where a single source stores the N segments that compose the content object C_p , but more than one client is interested in C_p , as illustrated in Fig. 1b. In order to minimize the time needed for each client to receive the complete set of segments that compose C_p , while also minimizing the number of duplicated transmissions of the same segment in the network, the segments need to travel over cost-efficient multicast distribution trees [2]. In CCN, this means that each node of the network should know where each Interest $i_{p,n}$ should be forwarded such that all the Interests for the segment $c_{p,n}$ from different clients are aggregated in the optimal point in the network that reduces the number of copies of $c_{p,n}$ transmitted before reaching all the clients interested in it. However, finding the optimal set of multicast distribution trees in a distributed manner is a very complicated task [2], which requires the knowledge of the network topology. In the simple example shown in Fig. 1b, if all the clients send the Interest $i_{p,n}$ over the LTE face, the segment $c_{p,n}$ will be transmitted from the source to the LTE network and then to the clients. However, if some of the clients decide to send the Interest $i_{p,n}$ over the WiFi face, the segment $c_{p,n}$ will also be transmitted from the source to the WiFi network, wasting resources that could have been used to transmit another segment.
- *Multi-source multicast*: In this scenario, both of the above mentioned problems manifest, as it is a combination of both scenarios. Hence, it is more challenging for CCN to efficiently deliver the content to the clients.

B. Enabling network coding in CCN

The shortcomings of the CCN architecture discussed in Section III-A can be dealt with by enabling network coding [3], a technique in which the segments delivered to the clients are coded at sources and intermediate nodes. The key idea behind introducing network coding in CCN is that clients no longer need to request specific data segments, but rather encoded data segments as they all have the same amount of information. This removes the need to coordinate the forwarding of Interests and leads to a more efficient use of the available network bandwidth.

Differently from the original CCN where an Interest message $i_{p,n}$ requests a specific segment $c_{p,n}$, in a network coding enabled CCN variant an Interest message \hat{i}_p requests a coded segment \hat{c}_p , without specifying a particular segment id. The sources or intermediate nodes can reply to these Interests with network coded segments, generated by combining the segments in the CS of the node that match the name prefix p . In matrix form, this can be expressed as $\hat{c}_p = \mathbf{A} \cdot \mathbf{CS}_p$, where \mathbf{A} is a vector of coding coefficients drawn from a Galois field, and \mathbf{CS}_p is a vector of the segments in the CS of the node that match the name prefix p . When the coding coefficients in \mathbf{A} used to combine the packets are chosen uniformly at random from a large enough Galois field, the generated segments have high probability of being linearly independent, and thus innovative. Whenever a client interested in a content object C_p collects N innovative coded segments \hat{c}_p , it can decode the original segments that compose C_p .

To illustrate the benefits that network coding brings to CCN, let us revisit the scenarios described in Section III-A.

- *Multi-source unicast*: Differently to the original CCN, when network coding is allowed, the client and the intermediate nodes do not need to know over which face they can reach a particular source, since they send Interests for coded segments rather than for specific segments. This implies that the FIB of the clients and intermediate nodes do not need an entry for each segment, but only a single entry for the name prefix leading Interests towards the sources is enough. Each source then reply to these Interests with coded segments \hat{c}_p , generated by combining the segments that match the name prefix p .
- *Single-source multicast*: When network coding is allowed the clients do not need to coordinate what Interests they send over each face, since all of them are for coded segments. Thus, when all the clients send an Interest \hat{i}_p requesting a coded segment over a face (e.g., LTE) it will be aggregated by the intermediate nodes, and only one Interest requesting a coded segment will reach the source.
- *Multi-source multicast*: In this scenario, when network coding is allowed, neither clients nor intermediate nodes need to coordinate the forwarding of the Interests, since they are for coded data and can be satisfied by any coded segment.

Network coding has also been shown to improve the throughput when bottlenecks are present in the network and the resiliency to packet erasures. In order to illustrate these benefits, let us consider the scenario in Fig. 1c, commonly known as the butterfly network. We consider that two clients are interested in a content object $C_p = \{c_{p,1}, c_{p,1}\}$ that consists of two segments. Each source holds only one of the segments. In this case, if network coding is not enabled, one of the clients will always receive the content with higher delay. This is because the link between nodes r_3 and r_4 becomes a bottleneck since the only way in which the clients u_1 and u_2 can get $c_{p,2}$ and $c_{p,1}$ respectively, is through node r_4 . Thus, the

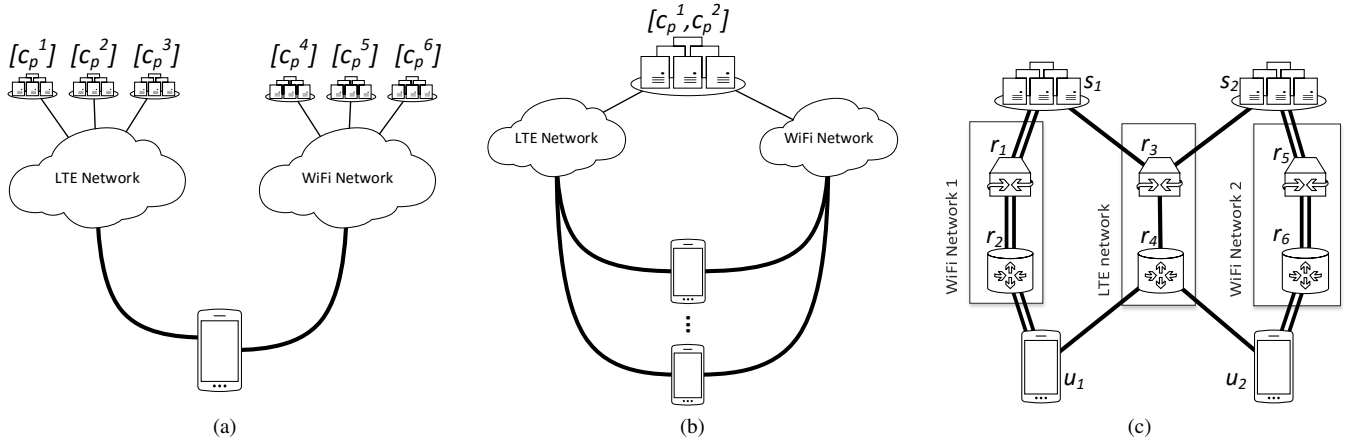


Fig. 1. Mobile devices retrieving segments in parallel over LTE and WiFi: (a) multi-source unicast; (b) single-source multicast; (c) multi-source multicast (butterfly network).

Interests sent by the clients u_1 and u_2 cannot be aggregated in the node r_4 , since they are for different segments. In contrast, when network coding is enabled, the Interests sent by the clients u_1 and u_2 can be aggregated in the node r_4 , as they are both for coded data. If the node v_3 applies network coding to the segments received from the sources, the resulting coded segment will be useful for both clients u_1 and u_2 .

C. Challenges

As discussed in Section III-B, enabling network coding in CCN nodes brings benefits that can potentially improve the performance of content object retrieval under certain scenarios. However, some issues arise when the Interest messages do not specify the segment id.

One of the issues that arises is that any node that has a single coded segment \hat{c}_p cached in its CS will reply with it to all the Interests \hat{i}_p , as the name prefix in the Interest matches that of the cached segment \hat{c}_p . This is undesirable, since the intermediate nodes will always reply with the same cached segment \hat{c}_p , while clients need to receive N innovative coded segments in order to decode the original segments. Therefore, the intermediate nodes need a way to determine when they cannot provide a coded segment that is innovative to client, and thus a new coded segment need to be retrieved. In [4] the authors propose to solve this problem by allowing the clients to include information about the coded segments they have collected so far. Intermediate nodes will reply to an Interest only if they can provide innovative information. However, it is not clear how intermediate nodes can aggregate Interests with different information from the clients.

Another challenge that emerges when network coding is enabled in CCN is related to the *pipelining* procedure, *i.e.*, a client sending multiple concurrent Interests for different segments of the same content object. In the original CCN when a node receives an Interest that it cannot satisfy with content stored in its CS, the node checks its PIT. If the node finds an entry in the PIT indicating that an Interest for the same name has been received previously over the same face,

it will consider this new Interest as a duplicate and will not forward it. Since Interests for different segments will have different names, as the segment id is appended to the name prefix, pipelining is supported. However, when network coding is enabled in CCN, concurrent Interests for the same content object will be considered duplicated by the intermediate nodes, as Interests for different (linearly independent) coded segments of the same object will have the same name.

IV. THE NETCODCCN PROTOCOL

In this section, we present NetCodCCN, a practical implementation of a network coding enabled content-centric networking protocol. We build our proposal on the CCN [1] architecture. We start by defining the content segmentation and naming scheme in NetCodCCN. Then, we describe how Interests and Data messages are processed in our proposed protocol.

A. Content Segmentation

As in CCN, in NetCodCCN content objects are split into smaller segments, $C_p = \{c_{p,1}, \dots, c_{p,N}\}$, that fit into Data messages. Network coded segments \hat{c}_p are generated by the sources by randomly combining the set of L segments, $L \leq N$, with name prefix p that are stored in their CS. This set of segments is denoted as \mathbf{CS}_p . Thus, $\hat{c}_p = \sum_{l=1}^L a_l \cdot c_{p,l}$, where a_l is a randomly selected coding coefficient and $c_{p,l}$ in the l th segment in \mathbf{CS}_p . As the segments \mathbf{CS}_p used to create \hat{c}_p and the coding coefficients a_l are only known by the node performing network coding, information about both the segments and the coding coefficients should be included in each network coded segment, in order to allow decoding by the clients. In NetCodCCN, this is accomplished by following the approach presented in [9], where an encoding vector g_n is associated to the segment $c_{p,n} \in C_p$. The initial value of this encoding vector is an n th unit vector, meaning it is a vector that has value 1 in the n th position and 0 otherwise. Network coding operations are performed on both the segments and their corresponding encoding vectors, generating

$\hat{c}_{p,g} = \sum_{l=0}^{L-1} a_l \cdot c_{p,l}$, where $\hat{c}_{p,g}$ is a coded vector associated with the encoding vector g and $g = \sum_{l=0}^{L-1} a_l \cdot g_l$. Intermediate nodes generate network coded segments as the sources, but they combine segments that have been already coded, thus $\hat{c}_{p,g} = \sum_{l=0}^{L-1} a_l \cdot \hat{c}_{p,g_l}$, where $g = \sum_{l=0}^{L-1} a_l \cdot \hat{g}_l$.

The clients and intermediate nodes keep track of the received innovative encoding vectors g_l , $\mathbf{G}_p = [g_1; \dots; g_L]$, so that the original set of segments can be retrieved by performing Gaussian elimination when the matrix \mathbf{G}_p is full rank, *i.e.*, it contains $L = N$ linearly independent coding vectors.

The use of the encoding vectors introduces a communication overhead that leads to waste of network resources, especially when the content object is segmented in a large number of content segments. To limit this overhead, we adopt the concept of *generations* [9], where the original set of segments that compose C_p is partitioned into smaller groups of segments, *i.e.*, generations, and the coding operations are restricted only between segments that belong to the same generation. For example, let us consider that C_p is partitioned into K generations of H_k segments each one. Hence, the H_k segments of the k th generation are denoted as $C_{p,k} = \{c_{p,k,1}, \dots, c_{p,k,H_k}\}$, where k is the *generation id*. The size of the generation, H_k , controls the tradeoff between the decoding delay, the packet diversity and the overhead required to communicate the encoding vector. Overall, the encoding vectors do not pose any limitations to our system as there are approaches to compress them efficiently [10], [11]. In order to avoid mixing packets from different generations, the segments are tagged with the generation id.

B. Content Naming

From the discussion above, it is obvious that the naming in NetCodCCN should have two additional components, *i.e.*, the encoding vector g and the generation id k . Let us consider a data object C_p with content name $p = /provider/videos/largevideo.h264$. This C_p is partitioned into K generations of H_k segments each one. Thus, in NetCodCCN the first segment of the k th generation, $c_{p,k,1}$, associated with the unit vector $[1, 0, 0, 0]$ when $H_k = 4$, is named $\{p, k, 1\} = /provider/videos/largevideo.h264/k/1000$.

For the sake of clarity, and without loss of generality, hereafter we consider that the name prefix p in $\hat{c}_{p,g}$, considers both the name prefix and the generation id. Note that, the proposed naming scheme is compatible with the original CCN and can support the delivery of non coded segments.

C. Interest Message Processing

Similarly to CCN, in our protocol the data communication is triggered by the clients who send Interest messages \hat{i}_p for data with name prefix p . In the proposed NetCodCCN protocol, the Interests have a *NetworkCodingAllowed* field that takes the value “1” when network coded segments are expected, otherwise, the field is not present or its value is set to “0”. Nodes receiving an Interest with the *NetworkCodingAllowed* field activated process the Interest messages following the

NetCodCCN procedure explained below and summarized in Algorithm 1. Otherwise, the Interests are treated following the original CCN procedures.

Algorithm 1 Interest Processing in NetCodCCN

Require: $\hat{i}_p, f, \mathbf{CS}_p \leftarrow$ segments that match p in the CS

- 1: **if** Decoded(p) **then**
- 2: $\hat{c}_{p,g} \leftarrow \sum_{l=0}^{L-1} a_l \cdot \hat{c}_{p,g_l}$
- 3: Send segment $\hat{c}_{p,g}$ over face f
- 4: **else**
- 5: $\nu^{p,f} = \text{Rank}(\mathbf{G}_p) - n_{sent}^{p,f}$
- 6: **if** $\nu^{p,f} > 0$ **then**
- 7: $\hat{c}_{p,g} \leftarrow \sum_{l=0}^{L-1} a_l \cdot \hat{c}_{p,g_l}$
- 8: Send segment $\hat{c}_{p,g}$ over face f
- 9: **else**
- 10: InsertPIT (p, f)
- 11: **if** $n_p^{pending} < n_{p,f}^{pending}$ **then**
- 12: PropagateInterest(\hat{i}_p)
- 13: **end if**
- 14: **end if**
- 15: **end if**

When a node $v \in \mathcal{V}$ receives an Interest \hat{i}_p for a network coded segment over a face f , it either (i) replies to the Interest with a coded segment generated with the set of segments \mathbf{CS}_p ; or (ii) forwards the Interest to other nodes in order to receive a new linearly independent segment. This is further explained in the following.

Replying an Interest: The node v replies to an Interest \hat{i}_p when (i) it has collected enough network coded segments to decode the original set of segments C_p ; or when (ii) a segment generated by node v has high probability to be innovative for the node connected through face f from where the Interest arrived. The number of coded segments $\nu^{p,f}$ that can be generated by the node v and have high probability to be innovative for the node connected through face f , is given by $\nu^{p,f} = \text{rank}(\mathbf{G}_p) - \delta_{sent}^{p,f}$. The parameter $\delta_{sent}^{p,f}$ denotes the number of segments belonging to \mathbf{CS}_p that have been previously sent over face f . When $\nu^{p,f}$ is higher than 0, v generates a new coded segment and sends it over face f .

Forwarding an Interest: The node v forwards an Interest \hat{i}_p to its neighbors when $\nu^{p,f}$ is equal to 0, since it needs a new segment that increases the rank of \mathbf{G}_p before it can reply to Interest \hat{i}_p . As in CCN, prior to forwarding an Interest, the node v checks its PIT. However, in order to support pipelining, in NetCodCCN the PIT verification procedure is modified. Specifically, if the node finds a matching PIT entry indicating that an Interest for the same name prefix p has been previously received over the same face f , the Interest is not considered duplicated, but it is treated as a request for additional network coded segments from the same face, as shown in Fig. 2. This means that the face f can appear multiple times in the PIT entry for the name prefix p . To decide whether the Interest \hat{i}_p should be forwarded, the node v computes the number ν_p of innovative coded segments matching the name prefix

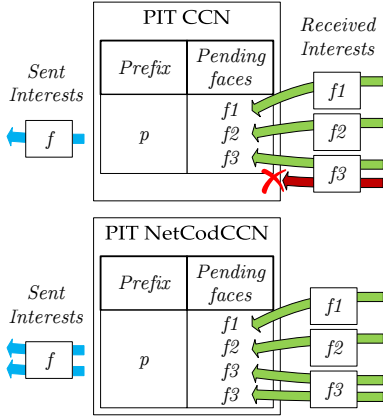


Fig. 2. Comparison of the PIT in CCN and in NetCodCCN.

p that it is expecting to receive before the Interest \hat{i}_p expires. If $\nu_p > n_{p,f}^{pending}$, where $n_{p,f}^{pending}$ is the number of Interest received over face f that are pending for a reply, the node v does not forward the Interest \hat{i}_p , as it will receive enough coded segments to satisfy all the pending Interests, including the received Interest \hat{i}_p . Otherwise, if $\nu_p \leq n_{p,f}^{pending}$, the node v forwards the Interest.

To compute the expected value of ν_p , the node v needs a probabilistic model that can take into consideration the loss rate of the system, the delays that segments may suffer, etc. For the sake of simplicity, we make the assumption that nodes follow a simple model in which any forwarded Interest brings an innovative segment before its expiration. In this case, $\nu_p = n_p^{pending}$, where $n_p^{pending}$ is the total number of pending interests for the name prefix p . Thus, the node will forward the Interest if $n_p^{pending} \leq n_{p,f}^{pending}$.

When a segment with name prefix p is removed from the CS of node v (e.g., when the CS eviction policy decides that segments of name prefix p needs to be removed from the cache), the corresponding vector should also be removed from \mathbf{G}_p , and $\delta_{sent}^{p,f}$ needs to be decreased by 1 for all the faces.

D. Data Message Processing

The use of network coding in CCN imposes also modifications in the Data message forwarding procedure. Specifically, when a node $v \in \mathcal{V}$ receives a coded segment $\hat{c}_{p,g}$ on face f , it should determine whether $\hat{c}_{p,g}$ is innovative. The segment $\hat{c}_{p,g}$ is innovative for the node v if the encoding vector g is linearly independent in \mathbf{G}_p , i.e., if it increases the rank of \mathbf{G}_p . A non innovative segment is considered as a duplicated segment and therefore it is discarded by the node v . When the segment $\hat{c}_{p,g}$ is innovative, the node v inserts it into its CS, and updates the encoding matrix \mathbf{G}_p that should now contain the received encoded vector g . Then, the node v checks its PIT. If the node v finds a matching PIT entry, meaning that an Interest with name prefix p is pending, it generates a network coded segment $\hat{c}_{p,g'} = \sum_{l=0}^{l=L} a_l \cdot \hat{c}_{p,g_l}$ and sends it once over each of the faces specified in the PIT entry. It is important to note that since the face f can appear multiple times in the

PIT entry for the name prefix p , as a consequence of allowing pipelining, a coded segment sent over face f consumes only one of the appearances of f in the corresponding PIT entry. If no matching PIT entry for the name prefix p is found, the received coded segment is considered unsolicited and is not further transmitted. However, it can be kept in the CS, according to the CS insertion policy, as it can be useful to serve future Interests. This procedure is outlined in Algorithm 2.

Algorithm 2 Content Processing in NetCodCCN

Require: $\hat{c}_{p,g}$

- 1: **if** g increases the rank of \mathbf{G}_p **then**
- 2: Insert $\hat{c}_{p,g}$ in the CS
- 3: $PIT_p \leftarrow$ PIT entry for prefix p
- 4: **if** $PIT_p = \emptyset$ **then**
- 5: Discard $\hat{c}_{p,g}$
- 6: **else**
- 7: **for all** f in PIT_p **do**
- 8: **if** f has not been served **then**
- 9: $\hat{c}_{p,g'} = \sum_{l=0}^{l=L} a_l \cdot \hat{c}_{p,g_l}$
- 10: Send segment $\hat{c}_{p,g'}$ over face f
- 11: $n_{sent}^{p,f} \leftarrow n_{sent}^{p,f} + 1$
- 12: Remove one appearance of f from PIT_p
- 13: **end if**
- 14: **end for**
- 15: **end if**
- 16: **else**
- 17: Discard $\hat{c}_{p,g}$
- 18: **end if**

It is important to note that the application of network coding adds complexity to the Interest and Data message processing in NetCodCCN. Performing algebraic operations on the segments before forwarding them adds, effectively, some complexity to the CCN node. In particular, as we have seen in Section IV-A, a node generates a new coded segment $\hat{c}_{p,g'}$ as $\hat{c}_{p,g'} = \sum_{l=0}^{l=L} a_l \cdot \hat{c}_{p,g_l}$. If we consider that the operations are performed in a Galois field of size 2^8 and that segments are of size X symbols, each time a node needs to generate a new coded segment, it should perform $X \cdot L$ multiplications and $X \cdot (L - 1)$ additions. This complexity does not pose limitations to our scheme as there are efficient implementations of network coding [12]. Further, as it has been shown in [13], a network coding coder and decoder can operate at wire-speed with rates of up to 1000Mbps.

V. NETCODCCN EVALUATION

In this section, we evaluate the performance of NetCodCCN in various scenarios, and compare the results to the performance of the standard CCN. First, we describe the simulation setup. Then, we evaluate the performance of NetCodCCN in the butterfly network. This toy network provides a controllable environment which permits to verify the expected behavior of NetCodCCN, and facilitates illustration of its benefits. Finally, we present the simulation results in a more realistic

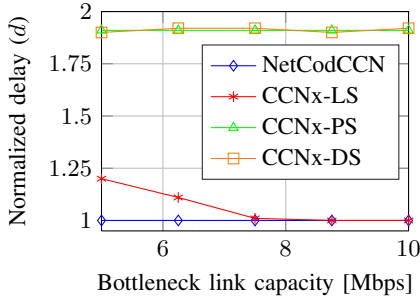


Fig. 3. Normalized delivery delay versus the capacity of the bottleneck link in the butterfly network.

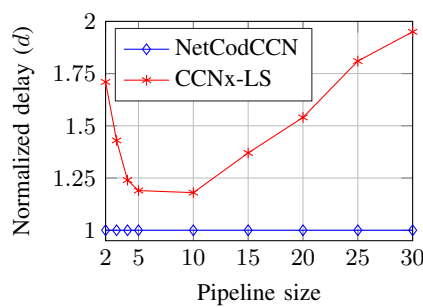


Fig. 4. Normalized delivery delay versus the pipeline size in the butterfly network.

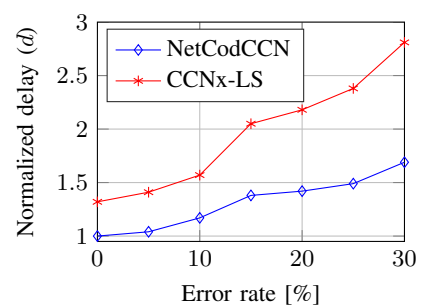


Fig. 5. Normalized delivery delay versus the error rate in the butterfly network.

network topology, which is generated based on real network measurements taken from the Planetlab project [14].

A. Simulation Setup

We implemented NetCodCCN by integrating the changes to the CCN architecture described in Section IV into the CCNx 0.8.2 [8] code, and we compare its performance to that of the unmodified CCNx. The network topology is simulated using the NS-3 network simulator [15]. The software forwarders/routers for CCNx and NetCodCCN are installed on NS-3 nodes using the Direct Code Execution framework (DCE) [16].

We consider that the clients are interested in a content object composed of $N = 100$ segments. The size of each segment is 5KB. The data segments are stored in a set of sources that are connected to the clients through a network of intermediate nodes. We consider that the intermediate nodes have sufficient CS space to store all the incoming data. We assume that the $N = 100$ source data segments comprise a single generation, *i.e.*, $H = N$ and $K = 1$. The finite field in which the network coding operations are performed is of size 2^8 . In order to evaluate our protocol in a challenging scenario, we consider that all the clients send Interests for the content segments during the same interval of time. In this way, we demonstrate that by using our protocol, nodes are able to aggregate Interests adequately.

For the evaluation of CCNx, we consider the three main Interest forwarding strategies implemented in CCNx 0.8.2, and described in [17]:

- The *default* (DS) strategy selects the fastest responding face based on the face statistics.
- The *loadsharing* (LS) strategy distributes the Interest forwarding load over all the available faces, sending each Interest over the face with the smallest pending Interest queue.
- The *parallel* (PS) strategy sends the Interests in parallel over all the faces indicated on the FIB.

For the evaluation of NetCodCCN, we always consider the *parallel* strategy, since by sending a single Interest over all its faces the client can receive multiple useful segments, *i.e.*, linearly independent.

To evaluate the performance of NetCodCCN, we measure the time $\Delta t_{measured}$ that a client needs in order to get the N segments available at the sources. In the standard CCN, $\Delta t_{measured}$ is defined as the elapsed time between the transmission of the first Interest and the reception of the N th missing segment. In NetCodCCN, $\Delta t_{measured}$ is defined as the elapsed time between the transmission of the first Interest and the reception of the N th linearly independent network coded segment, which permits to decode the whole generation of segments. We consider that clients can have heterogeneous network resources, thus, in order to make a fair comparison of the delivery delay, we define the *normalized delivery delay* as $d = \Delta t_{measured} / \Delta t_{min}$, where Δt_{min} is the theoretical lower bound on the time that a client would need in order to receive all the segments if it was alone in the network and was able to receive at max-flow rate. Thus, a normalized delivery delay equal to 1 means that the client was able to receive the complete set of packets at the maximum rate. Note that $\Delta t_{measured} \geq \Delta t_{min}$, or equivalently, $d \geq 1$ always holds.

B. Butterfly Topology

We begin by evaluating NetCodCCN in the butterfly topology presented in Fig. 1c. We consider that every data segment is stored randomly in at least one of the two sources, and a copy of the same segment is also placed in the Content Store of the other source with a duplication probability $\phi \in [0, 1]$. We set the capacity of every link in the network to $5Mbps$.

In the first set of experiments, we consider that $\phi = 1$, *i.e.*, both sources hold a copy of each data segment in their CS. This corresponds to the single source multicast case presented in Section III-A. In this case, clients u_1 and u_2 can reach a copy of any segment over any of their faces. However, as explained in Section III-B, with the original CCN protocol the maximum performance can be achieved only if both clients coordinate and send Interest messages for the same segments over the faces that connect them to the node r_4 . In contrast, when network coding is employed, the need for coordination is eliminated, since clients do not send Interests for a specific segment but rather for any network coded segment.

Fig. 3 depicts the normalized delivery delay as a function of the capacity of the bottleneck link between nodes r_3 and r_4 . We can see that NetCodCCN achieves the optimal

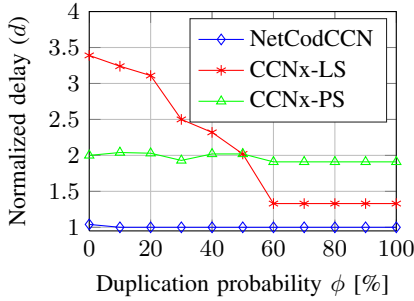


Fig. 6. Normalized delivery delay versus the source content duplication probability in the butterfly network.

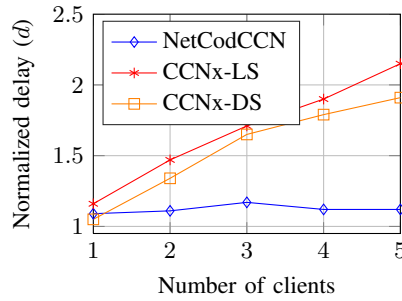


Fig. 7. Normalized delivery delay versus the number of clients in the network in the PlanetLab topology.

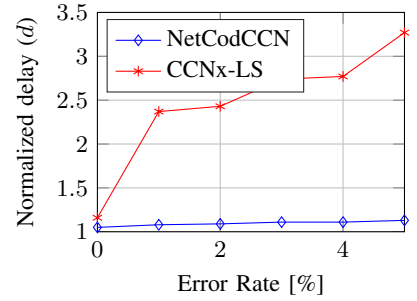


Fig. 8. Normalized delivery delay versus the segment transmission error rate in the PlanetLab topology.

performance in the whole range of link capacity values. This is due to the fact that network coding removes the need for coordinating the forwarding of Interest messages. In contrast, the CCN forwarding strategies perform poorly and only the LS strategy can achieve the performance of NetCodCCN but requires significantly higher link capacity. When the bottleneck link has the same capacity as all the other links, the average delivery time d of CCNx-LS is around 1.2 times the minimum delivery delay, Δt_{min} . This is caused by the randomness introduced by the LS strategy when choosing the faces over which Interests are transmitted when all the faces have the same load. This creates two extreme cases. In one case, all the Interests sent by both clients to node r_4 are the same, thus d tends to one. In the other case, all the Interests are different, thus d tends to 1.33. This happens because each node will receive $2/3$ of the segments through the link connecting them to the sources, and $1/3$ over the face connecting them to the node r_4 , which means that $2/3$ of the total segments will travel on the bottleneck link. With the DS and the PS strategies, the average delivery time is close to 2, as expected. With the DS strategy, the face connecting the clients to the sources will be chosen as the best, and thus most of the segments will be received over that face. With the PS strategy, each client will forward every Interest over both faces, thus bringing one copy of each segment over each face.

We now investigate how the number of concurrent Interests that a client can send, also known as the pipeline size, affects the performance in terms of the average normalized delivery delay. As shown in Fig. 4, the performance of CCN is optimized for a pipeline size value between 5 and 10, where the normalized delivery delay seen by the clients is 1.2. This is due to the fact that clients need to send at least 4 Interest messages over the faces connecting them to the node r_4 in order to create a continuous flow of segments. Since the LS strategy distributes the Interests over all available faces, a client has to send 4 Interests over each face while it also has sent 3 or 4 Interests over the other face, which amounts to 7 or 8 Interests in total. For smaller pipeline sizes, the continuous flow is not set, while for larger pipeline sizes the number of Interests sent over the bottleneck link increases, thus worsening the client coordination problem. In contrast,

the performance of NetCodCCN is not affected by the pipeline size, as can be verified in Fig. 4. This can be explained by the fact that NetCodCCN eliminates the necessity that the clients request the same segments over the bottleneck link. For the rest of the experiments, we choose a pipeline size of 10.

In Fig. 5, we depict the influence of the data segment loss rate on the performance of NetCodCCN and of the original CCN. We consider losses that are caused both by the transmission losses and the errors during the processing of the segments. We can see that the performance of CCN with the LS strategy degrades faster than the performance of NetCodCCN as the segment error rate increases. This is caused by the fact that in CCN, the client will be able to react to a segment loss only when the corresponding Interest expires, since any earlier re-transmission of an Interest with the same prefix will be prevented by the PIT. Instead, with NetCodCCN, the clients can send Interests for new coded segments until they have a sufficient number of coded segments in order to recover the original ones. It is important to note that the maximum amount of concurrent Interests that a client can send is controlled by the pipeline size.

Finally, we evaluate the performance of NetCodCCN for different values of the duplication probability ϕ . This corresponds to the multi-source multicast case presented in Section III-A. In CCNx, when $\phi < 1$, the clients should not only coordinate the requests sent over the bottleneck link as in the previous scenario, but they also should have the knowledge of the segments that each source stores, in order to avoid sending Interests over the face connecting them directly to the source that does not hold a copy of the requested segment. In Fig. 6 we can see that CCN with the LS strategy takes 3.4 times longer to deliver all the segments to the clients, when each segment is stored only in one of the sources. When the PS strategy is employed, the clients do not need to know how the content is distributed since each Interest message is sent over both of its faces. However, since a copy of every segment will cross the bottleneck link, the traffic over the bottleneck link will be doubled compared to the network coding case. When the probability that the segments are stored in both sources increases, the performance of CCNx with the LS strategy improves, but eventually saturates at 1.2 times the minimum

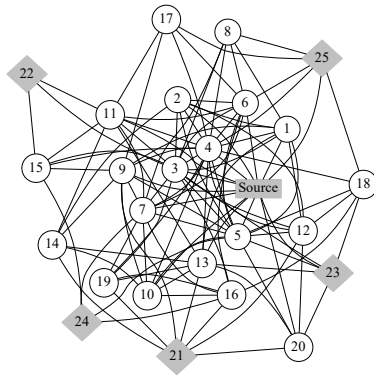


Fig. 9. Planetlab topology used.

delay, which is consistent with the results depicted in Fig. 4.

C. Planetlab Topologies

We now evaluate our protocol in more realistic network topologies captured by the PlanetLab project [14]. We use the network topology shown in Fig. 9 that consists of one source node, 5 client nodes and 20 intermediate nodes. The links connecting the nodes have a capacity of 12Mbps. The topology was generated using the procedure described in [18]. We measure the normalized delivery delay d for each client and then compare its average.

First, we investigate how the performance is affected by the number of clients in the network. In Fig. 7 we can see that with a single client (in this case node 24), NetCodCCN and CCNx perform similarly. In this case, network coding does not introduce any gains since there is only one client in the network and no losses are considered. However, the performance of CCNx starts to degrade with the introduction of more clients, as they start to compete for the network resources. In contrast, we can see that the performance of NetCodCCN does not deteriorate with the addition of new clients to the network topology. These results show that the NetCodCCN protocol uses more efficiently the available network resources.

We also evaluate how the error in segment transmission affects the performance of the NetCodCCN for larger topologies. For this evaluation, we choose to keep only one client, in order to compare the results with the performance of the CCN. As with the butterfly topology, we consider losses that are caused both by the transmission losses and the errors during the processing of the segments. In Fig. 8 we can see that NetCodCCN maintains the delivery delay close to the expected one, while the performance of CCN degrades very fast with the introduction of errors. As in the butterfly topology, this fast degradation is due to the fact that when a segment is lost, the client needs to wait until the corresponding Interest expires before it can re-send a new one.

VI. CONCLUSIONS

In this paper, we have presented NetCodCCN, a protocol that integrates network coding in CCN. Specifically, we have defined novel algorithms to process Interest and Data messages so that consider network coded content. In NetCodCCN, the clients express Interest messages for coded segments of a given prefix instead of asking a specific segment as in CCN. The network nodes combine the Data messages by means of RLNC before forwarding them in order to take advantage of the network diversity. Our protocol is able to (i) simplify the aggregation of Interest for coded content; (ii) reduce the number of duplicate segments; and (iii) allow clients to send multiple Interests for the same content in parallel. The overall system has been tested in networks with multiple clients and sources, where we have observed large performance gains in terms of the time needed to retrieve the demanded content.

Our future research includes the investigation of optimal Interest forwarding strategies that enable flow control in the case of multiple different content objects. We will also consider the transmission of video content characterized by strict delivery deadlines. Furthermore, we will focus on enabling content security on NetCodCCN.

ACKNOWLEDGMENTS

This work has been partially funded by the Swiss National Science Foundation under grant number 149225.

REFERENCES

- [1] 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*, New York, NY, USA, 2009, pp. 1–12. [Online]. Available: <http://doi.acm.org/10.1145/1658939.1658941>
- [2] Y. Wu, P. Chou, and K. Jain, "A Comparison of Network Coding and Tree Packing," in *Proc. of IEEE ISIT'04*, June 2004, pp. 143–.
- [3] R. Ahlswede, N. Cai, S.-Y. Li, and R. Yeung, "Network Information Flow," *IEEE Trans. Information Theory*, vol. 46, no. 4, pp. 1204–1216, Jul 2000.
- [4] M.-J. Montpetit, C. Westphal, and D. Trossen, "Network Coding Meets Information-Centric Networking: An Architectural Case for Information Dispersion Through Native Network Coding," in *Proc. of the 1st ACM NoM Workshop*. New York, NY, USA: ACM, 2012, pp. 31–36. [Online]. Available: <http://doi.acm.org/10.1145/2248361.2248370>
- [5] J. Sundararajan, D. Shah, M. Medard, S. Jakubczak, M. Mitzenmacher, and J. Barros, "Network coding meets tcp: Theory and implementation," *Proceedings of the IEEE*, vol. 99, no. 3, pp. 490–512, March 2011.
- [6] Q. Wu, Z. Li, and G. Xie, "CodingCache: Multipath-Aware CCN Cache with Network Coding," in *Proc. of the 3rd ACM ICN Workshop*. New York, NY, USA: ACM, 2013, pp. 41–42. [Online]. Available: <http://doi.acm.org/10.1145/2491224.2491240>
- [7] J. Llorca, A. Tulino, K. Guan, and D. Kilper, "Network-Coded Caching-Aided Multicast for Efficient Content Delivery," in *Proc. of IEEE ICC'13*, June 2013, pp. 3557–3562.
- [8] "Project CCNx ®, version 0.8.2," <http://www.ccnx.org/releases/ccnx-0.8.2/doc/>.
- [9] P. Chou and Y. Wu, "Network Coding for the Internet and Wireless Networks," *IEEE Signal Processing Magazine*, vol. 24, no. 5, pp. 77–85, Sept 2007.
- [10] N. Thomas and P. Frossard, "Toward one Symbol Network Coding Vectors," *IEEE Communications letters*, vol. 16, no. 11, pp. 1860–1863, Nov. 2012.
- [11] D. E. Lucani, M. V. Pedersen, J. Heide, and F. H. P. Fitzek, "Fulcrum Network Codes: A Code for Fluid Allocation of Complexity," available at <http://arxiv.org/abs/1404.6620>, 2014.

- [12] M. Pedersen, J. Heide, P. Vingelmann, and F. Fitzek, "Network coding over the $2^{32} - 5$ prime field," in *Communications (ICC), 2013 IEEE International Conference on*, June 2013, pp. 2922–2927.
- [13] M. Zhang, H. Li, F. Chen, H. Hou, H. An, W. Wang, and J. Huang, "A general co/decoder of network coding in hdl," in *Network Coding (NetCod), 2011 International Symposium on*, July 2011, pp. 1–5.
- [14] "PlanetLab," <https://www.planet-lab.org/>.
- [15] "The network simulator - ns3," <http://www.nsnam.org/>.
- [16] "Direct Code Execution (DCE)," <https://www.nsnam.org/overview/projects/direct-code-execution/>.
- [17] "CCND(1) Manual Page, Project CCNx ®, version 0.8.2," <https://www.ccnx.org/releases/ccnx-0.8.2/doc/manpages/ccndc.1.html>.
- [18] N. Cleju, N. Thomos, and P. Frossard, "Selection of Network Coding Nodes for Minimal Playback Delay in Streaming Overlays," *IEEE Trans. Multimedia*, vol. 13, no. 5, pp. 1103–1115, Oct 2011.