



Leveraging the Chinese Remainder Theorem for outband D2D name-based content delivery

Marica Amadeo ^{*}, Filippo Battaglia , Giuseppe Campobello 

Department of Engineering, University of Messina, Italy
CNIT, Consorzio Nazionale Interuniversitario per le Telecomunicazioni, Italy

ARTICLE INFO

Keywords:

Device-to-Device Communications
Sidelink
Chinese Remainder Theorem
Named Data Networking
Information Centric Networking

ABSTRACT

The explosive growth of mobile data traffic demands scalable, resilient, and energy-efficient solutions for content dissemination, particularly at the network edge. In this context, 5G networks can rely on Outband Device-to-Device (O-D2D) communications, using technologies such as Wi-Fi Direct, to offload infrastructure, enhance coverage at the cell edge, and enable proximity-based services. However, O-D2D communications face critical challenges related to peer discovery, reliable data retrieval over lossy wireless links, and energy efficiency constraints. In this paper, we propose a novel framework that integrates Named Data Networking (NDN) and Chinese Remainder Theorem (CRT)-based fragmentation to support autonomous O-D2D content delivery. By leveraging name-based forwarding and in-network caching, NDN facilitates distributed content retrieval, making it inherently suited for dynamic and infrastructure-less D2D environments. CRT-based fragmentation, on the other hand, enhances reliability and energy efficiency by allowing data reconstruction from partial fragment sets, mitigating the effects of packet losses typical of wireless links. We implement the proposed architecture in ndnSIM v2.9, extending the NDN forwarding plane to support CRT-based multi-source data retrieval. Simulation results demonstrate that our solution significantly improves content delivery and energy efficiency compared to existing NDN retrieval schemes, particularly under lossy wireless conditions.

1. Introduction

The exponential growth in mobile data traffic driven by bandwidth-hungry services has prompted the need for alternative communication paradigms within cellular networks. In this context, Device-to-Device (D2D) communications enable direct communication between user devices, without traversing the Base Station (BS) [1].

Initially introduced to improve spectrum efficiency and reduce latency in 4G Long Term Evolution (LTE) systems [2], D2D communication is a pervasive technology supporting various applications, including content dissemination and proximity services, in 5G and beyond networks. In particular, in 5G networks the concept has been generalized with the sidelink interface (PC5) [3], which supports direct communication among devices, e.g., in scenarios like Industrial Internet of Things (IIoT) or vehicle-to-everything (V2X) [4].

D2D communications can be broadly classified, based on the spectrum used, into inband and outband categories. In inband D2D, devices communicate over the licensed cellular spectrum, either by reusing cellular resources (underlay) or using dedicated channels (overlay). While this approach benefits from tight integration with the cellular infrastructure, it introduces significant challenges related to interference management and resource allocation.

Conversely, outband D2D communications exploit unlicensed spectrum technologies (e.g., Wi-Fi Direct, Bluetooth), thereby decoupling D2D traffic from cellular spectrum usage. This mode can further be categorized into controlled (BS-assisted) and autonomous (device-managed) architectures [1]. Outband D2D (O-D2D) is particularly attractive for content distribution applications, as it enables efficient peer-to-peer data exchange without interfering with licensed cellular resources. However, despite its potential, O-D2D requires efficient and decentralized mechanisms for content discovery and retrieval. Although centralized infrastructure could support these functions, relying solely on it may limit scalability and increase signaling overhead, particularly in dense or dynamic environments.

In this context, the Named Data Networking (NDN) architecture [5] offers a compelling solution by shifting the communication model from host-centric to data-centric. Through name-based connection-less communication and in-network caching, NDN inherently supports distributed content discovery and facilitates content dissemination, making it a natural fit for autonomous O-D2D communication [6]. In NDN, content retrieval is driven by clients, who request contents by name by transmitting *Interest packets*. Any node owning the content,

^{*} Corresponding author at: Department of Engineering, University of Messina, Italy.
E-mail address: marica.amadeo@unime.it (M. Amadeo).

either the original producer or a cacher, can answer the requests with named *Data packets*.

Nevertheless, NDN alone cannot fully address the challenges posed by highly error-prone wireless environments such as autonomous O-D2D networks. In particular, NDN requires the successful delivery of complete Data packets, while short-range wireless links are often affected by packet losses, intermittent connectivity, and interference. Under such conditions, repeated retransmissions may be required, leading to increased latency and excessive energy consumption at battery-powered devices. Therefore, additional mechanisms are needed to enhance NDN's robustness against data loss, reduce retransmissions, and improve energy efficiency in disruption-prone O-D2D scenarios.

A promising approach to address these challenges is offered by the Chinese Remainder Theorem (CRT) [7], which enables robust and energy-efficient content delivery through packet fragmentation and reconstruction. CRT can be employed as a lightweight channel encoding technique in which a Data packet is divided into smaller redundant fragments that can be independently transmitted by different nodes [8]. The original Data packet can then be reconstructed at the receiver from any sufficiently large subset of these fragments. This property enables opportunistic data collection, allowing devices to recover the complete content even in the presence of losses, without requiring all fragments or retransmissions. The efficiency and algebraic simplicity of CRT make it especially appealing for energy-constrained wireless devices [9].

Moreover, CRT contributes to reducing the energy consumption at D2D devices. First, by dividing a packet into smaller fragments, each transmission becomes statistically less prone to losses, which reduces the number of retransmissions required in the presence of channel errors. Since the client can reconstruct the original content from any sufficient subset of fragments, this further lowers the retransmission overhead and the energy expenditure at both transmitting and receiving nodes. Second, CRT allows different nearby nodes to transmit distinct fragments, so that the delivery effort is shared across the network rather than concentrated on a single device, leading to more balanced energy consumption.

So far, CRT-based fragmentation schemes have been proposed with the aim of reducing energy consumption and improving the reliability of Wireless Sensor Networks [10] and, more recently, they have been used to improve the performance of content delivery over Bluetooth Low Energy (BLE) [11] and LoRa [12].

When integrated with NDN's data-centric forwarding and caching, CRT-based fragmentation can enhance the system's ability to handle lossy links, thus increasing the content dissemination performance and reducing the energy consumption in the network. In our previous work in [13], we pioneeredly introduced CRT fragmentation in Vehicular NDN (VNDN). Evaluation results, when considering vehicle-to-vehicle (V2V) and Vehicle-to-Infrastructure (V2I) communications, showed a consistent reduction of the content delivery latency.

In this work, instead, we explore the synergy between NDN and CRT in the context of O-D2D content distribution, aiming to provide an architecture that supports both autonomous content discovery and resilient *energy-efficient* data dissemination. More specifically, the contributions of our work are as follows:

- We design a CRT-based module for packet fragmentation and re-assembly within the Data Plane of NDN devices, enabling resilient D2D content delivery. The module is lightweight and seamlessly integrates with the NDN forwarding logic.
- We propose a content retrieval scheme tailored for O-D2D communication, where special Interest packets, called *Long Lived Interests* (LLI), are broadcast to trigger parallel responses from multiple neighboring devices. This mechanism enables multi-source fragment collection, without requiring prior knowledge of which nodes hold specific content fragments.

- To reduce the risk of wireless transmission collisions among CRT fragments, and consequently limit retransmissions and energy consumption, we define a randomized defer timer with over-hearing and a fine-tuned retransmission strategy associated with LLIs.
- We evaluate the feasibility and performance of the proposed approach through simulations conducted in ndnSIM [14], the reference simulator of the NDN research community. Results show that our solution significantly reduces content delivery time and enhances energy efficiency for both clients and content providers, compared to existing NDN delivery mechanisms.

The rest of the paper is organized as follows. Section 2 provides background on NDN and D2D communications. Section 3 presents the proposed approach, while Section 5 discusses the performance evaluation. Finally, Section 6 concludes the paper and outlines directions for future work.

2. Name-based D2D communications

2.1. D2D interactions

Content distribution through direct communication between devices has been extensively explored as a way to improve scalability and efficiency in mobile networks, particularly in dense environments or scenarios with intermittent infrastructure coverage [1].

In terms of system-level implementations, the 3rd Generation Partnership Project (3GPP) has progressively incorporated D2D communications under the term Proximity Services (ProSe), which has evolved into the New Radio (NR) sidelink interface in 5G systems [3]. NR sidelink introduces enhanced features such as flexible scheduling and unicast/groupcast communication, and is designed to operate in both coverage and out-of-coverage scenarios. Although current releases focus primarily on vehicular and public safety use cases, the underlying capabilities of NR sidelink could theoretically support advanced content dissemination strategies in different application domains aimed at boosting data rates and coverage. Moreover, the recent Release 19 introduces User Equipment-to-User Equipment (UE-to-UE) relaying, laying the groundwork for multi-hop communications in 6G networks [15].

From a resource allocation perspective, O-D2D communications, which are the focus of our work, are typically categorized into two main modes: controlled and autonomous [4]. In the controlled mode, the base station centrally schedules sidelink resources. Conversely, the autonomous mode allows user devices to select their own transmission resources without infrastructure assistance. This mode is essential for supporting out-of-coverage communications and fully decentralized systems, where each device monitors the channel over a sensing window to estimate resource availability.

Our framework aligns conceptually with autonomous D2D operation, where devices independently participate in content distribution using unlicensed spectrum (e.g., Wi-Fi Direct). As such, we leverage the principles of autonomous resource selection and distributed operation, and we consider the NDN architecture as an enabling communication paradigm.

2.2. Named data networking

NDN is the most prominent architecture within the broader Information-Centric Networking (ICN) paradigm [16]. It has been widely adopted and studied in both wired and wireless environments due to its content-centric communication model and native support for in-network caching and user mobility.

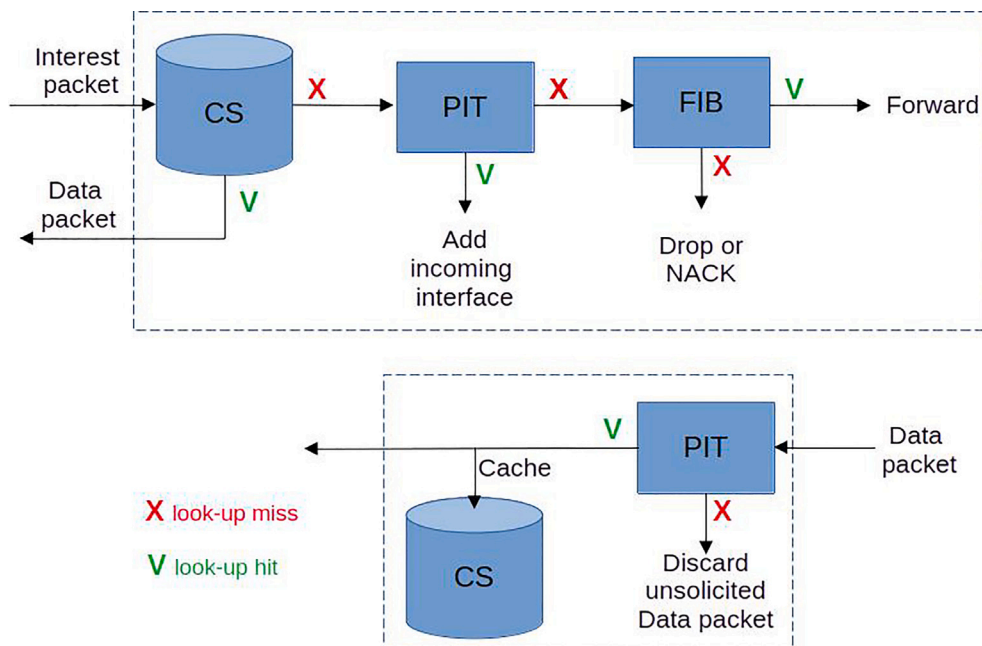


Fig. 1. NDN forwarding fabric.

2.2.1. Benefits

Implementing D2D communications via NDN can provide several benefits.

First, NDN shifts the communication paradigm from host-centric to content-centric, enabling devices to request data by name rather than by addressing specific hosts. This is particularly advantageous in D2D scenarios, where devices may dynamically join or leave the network, making it difficult or impractical for consumers to know or discover the corresponding IP addresses. By using content names, communication becomes more resilient to mobility and topology changes.

Second, NDN's in-network caching allows content that has already been retrieved by one device to be stored and reused by nearby devices thus limiting latency and network resource consumption. Moreover, by minimizing the need for repeated long-range communications or access to remote servers, caching can significantly lower energy consumption, which is an important advantage in D2D networks where devices often operate under strict energy constraints.

Third, NDN natively supports multicast delivery. This allows efficient data dissemination to multiple nearby devices with minimal overhead, which is useful for applications such as real-time content sharing, collaborative sensing, and proximity-based services.

These features make NDN a compelling candidate for enabling efficient and scalable D2D communication.

2.2.2. Node architecture

NDN communication is driven by consumers which issue Interest packets to request content by name [5]. When an Interest reaches a node that has the corresponding content, either as the original producer or from a cached copy, a Data packet is returned along the reverse path. There is a one-to-one correspondence between Interest and Data packets: for each Interest, at most one matching Data packet is sent back.

Each NDN node maintains the following three tables at the Data Plane:

- **Content Store (CS).** It is a local cache that temporarily stores incoming Data packets. The default caching strategy is to store all received Data packets and apply a Least Recently Used (LRU) replacement policy when the cache reaches its capacity.

- **Pending Interest Table (PIT).** It keeps track of forwarded Interests that have not yet been satisfied thus enabling stateful forwarding [17]. When an Interest is forwarded, an entry is created in the PIT with the Interest name and a record of the incoming interface(s). When the corresponding Data packet arrives, it is sent out on all interfaces listed in the matching PIT entry, and the entry is removed. As such, NDN natively supports multicast delivery.
- **Forwarding Information Base (FIB).** It contains name prefixes and corresponding outgoing forwarding interfaces. It is used to determine where to forward an Interest packet when no match is found in the CS or PIT.

As shown in Fig. 1, upon receiving an Interest, the node first checks the CS: if a matching Data packet is found, the request is satisfied locally, avoiding the need for further forwarding. Otherwise, the node looks in the PIT: if there is a matching entry, then the Interest is discarded. Otherwise, the node looks in the FIB to determine the appropriate outgoing interface(s).

Interest forwarding is driven by the so-called Strategy Module, which defines the actual policies, typically prefix-based, to transmit the packets. For example, the *multicast* policy may decide to transmit the same requests simultaneously over distinct access interfaces; while the *best-route* policy select the access interface with the lowest round-trip-time (RTT).

When a Data packet arrives, it is forwarded to the interfaces recorded in the PIT and optionally cached in the CS.

2.3. Related work

NDN for wireless intermittent networks. NDN has been widely investigated as a promising communication paradigm for wireless networks characterized by intermittent connectivity, node mobility, and the absence of fixed infrastructure. Its data-centric design, combined with in-network caching and receiver-driven communication, makes it particularly attractive for challenged environments where end-to-end paths may not always exist.

Several works have explored the use of NDN in intermittently connected IoT deployments. For instance, in [18], NDN is adopted in remote IoT scenarios following a store-carry-and-forward paradigm,

where data is opportunistically cached and delivered when connectivity becomes available. Similarly, [19] addresses both consumer and producer mobility in NDN-based IoT publish–subscribe systems, proposing mechanisms to repair broken paths and maintain data delivery despite frequent topology changes.

In Mobile Ad Hoc Networks (MANETs), NDN has also been studied as an alternative to traditional IP-based routing. In [20], the authors propose a location-aware multipath forwarding strategy coupled with distance-based caching, aiming to improve content delivery reliability under node mobility. Other approaches leverage additional control mechanisms or centralized components. For example, [21] integrates NDN forwarding with Software Defined Networking (SDN) to centrally manage mobility and energy-aware routing decisions in NDN-MANETs. In contrast to these works, our approach explicitly targets opportunistic D2D scenarios where connectivity is short-lived and packet losses are frequent.

NDN for D2D communications. One of the early attempts to explore the viability of NDN over D2D was presented by Shang et al. [6], who demonstrated the feasibility of deploying NDN stacks over various short-range wireless technologies, suggesting that NDN can effectively operate in infrastructure-less environments but requires careful system integration. In [22], NDN was advocated as a viable network-layer option within the 3GPP Cellular-V2X framework, targeting content delivery over both the traditional interface between user device and cellular BS (i.e., LTE-Uu) and the PC5 interface. A 3GPP V2X-NDN forwarding strategy is defined where decisions are context and interface-aware: requests for locally relevant data are sent over the PC5 interface, while Interests for remote content are routed over LTE-Uu. Furthermore, the BS integrates NDN's forwarding logic and utilizes Interest names to aggregate redundant grant requests for the same content. This enables Interest aggregation at the access layer, thereby reducing radio resource usage and dissemination latency.

The work in [23] investigates content-centric caching strategies in 5G D2D networks. The authors focus on how content can be efficiently distributed and cached among nearby devices, considering the dynamic and mobile nature of user nodes. They propose a time-based caching policy that adapts to user mobility and content popularity to optimize cache hit rate and reduce access latency.

The work in [24] proposes a cellular-assisted ICN framework for D2D content distribution, where BSs act as control-plane coordinators while content exchange occurs directly between nearby devices. The system relies on “helper” UEs, incentivized to cache popular content, which can be served via direct D2D links (e.g., Wi-Fi). When a user requests content, the BS, which is aware of content popularity and UE locations, resolves the request to a nearby helper if possible, thus offloading traffic from the infrastructure. Single-hop D2D transmissions are used to retrieve the content, while the centralized control function at the BS manages content assignment, caching/eviction, and D2D link setup.

Compared to these works, we propose a novel approach based on CRT to enhance the content delivery process while reducing the energy consumption of UEs in O-D2D communications. Our framework operates in a fully decentralized manner, enabling autonomous multi-source retrieval of content fragments among UEs with high resilience to wireless packet losses.

Packet fragmentation in NDN. Packet fragmentation in NDN has been extensively investigated, primarily to cope with Maximum Transmission Unit (MTU) limitations of the underlying links. Early work such as [25] focuses on the secure fragmentation of Interest and Data packets that exceed the MTU, while the approach proposed in [26] enables fragmentation and reassembly at intermediate routers to improve reliability under MTU constraints. However, these solutions are mainly conceived for wired multi-hop networks and rely on conventional sequence-based fragmentation mechanisms, where fragments are tightly coupled to the original Data packet and require full reassembly for caching and reuse.

More recently, NDN-based fragmentation has been revisited in multi-hop Low Earth Orbit (LEO) satellite networks [27], where large propagation delays and limited link capacities make hop-by-hop re-assembly particularly inefficient. In this context, fragmentation is still driven by MTU constraints, but novel mechanisms are introduced to forward fragments directly and enable fragment reuse at intermediate nodes in order to reduce end-to-end latency. Nevertheless, these approaches target stable multi-hop backbones with predictable connectivity and are not designed for highly dynamic, energy-constrained device-to-device environments.

In contrast to the above works, our approach addresses wireless D2D scenarios, where packet losses are mainly caused by channel impairments, interference, and intermittent connectivity rather than MTU limitations. Here, fragmentation is not a reactive mechanism imposed by the link layer MTU, but a deliberate network-layer strategy to improve content delivery robustness. By leveraging CRT-based fragmentation, our solution allows content reconstruction even in the presence of fragment losses, while avoiding repeated retransmissions of large Data packets.

Finally, in our previous work [13], we pioneered the integration of CRT-based fragmentation in a VNDN environment characterized by V2V and V2I interactions, with focus on a highway scenario under high mobility conditions, without explicitly considering energy constraints. Results demonstrated that the CRT-based delivery mechanism improves content accessibility and reduces retrieval delay. Compared to [13], the present contribution adopts a different system architecture tailored to energy-constrained O-D2D networks and introduces the following main design features:

- We consider a hybrid 5G + O-D2D architecture, where: (i) UEs may receive content via the cellular interface, (ii) CRT-based recovery is selectively activated over Wi-Fi Direct, and (iii) devices operate under autonomous sidelink conditions.
- We introduce an adaptive activation policy for the CRT mechanism, based on a short-term estimate of the packet error rate (PER), derived from standard MAC-layer retry/acknowledgment statistics available in IEEE 802.11 standards.
- Instead of focusing solely on reducing retrieval latency and improving content accessibility, one of the main objectives of this work is enhancing energy efficiency in battery-powered UEs operating in autonomous O-D2D environments. To this end, we consider an explicit energy consumption model and evaluate the energy cost of both NDN consumer and provider nodes in the O-D2D scenario.
- Although multi-source retrieval is conceptually present in both works, this paper refines the mechanism through a specific coordination strategy based on randomized defer timers explicitly tailored to IEEE 802.11 (and specifically to IEEE 802.11ax in our current implementation), thereby adapting fragment transmission control to contention-based O-D2D wireless access.

3. CRT for D2D communications

3.1. Motivations

Despite the benefits that NDN offers for D2D interactions, packet delivery in wireless environments remains highly vulnerable to errors due to interference, mobility, and link instability.

A key concern in such contexts is the *energy cost of retransmissions*. When packets are lost or corrupted, devices must resend them, incurring energy not only for the transmission itself but also for associated MAC-layer operations such as contention and channel sensing. This is particularly critical in O-D2D scenarios, where devices are often battery-powered and operate without infrastructure assistance.

The likelihood of transmission errors increases with packet size: larger packets contain more bits and are thus more prone to corruption, especially as the bit error rate (BER) worsens under dynamic

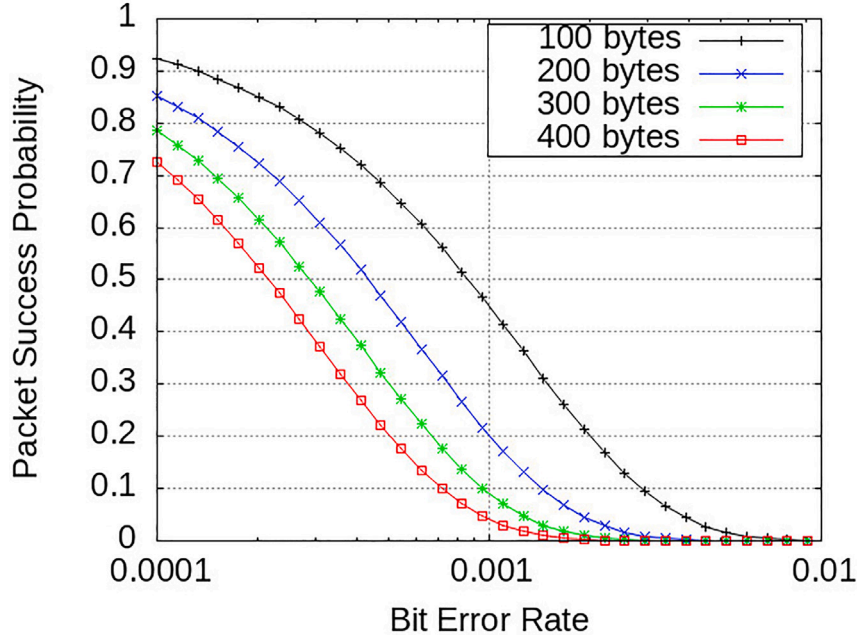


Fig. 2. Packet success probability when varying the bytes of the payload and the BER.

wireless conditions. Fig. 2 illustrates this by showing the *packet success probability*, expressed as:

$$P_{\text{success}} = (1 - \text{BER})^n, \quad (1)$$

for different payload sizes. As expected, larger packets experience a steeper drop in reliability as BER increases.

This drop in success probability has a direct impact on energy efficiency. Empirical studies (e.g., [28]) have shown that the energy consumed per bit increases non-linearly as retransmissions accumulate. Even high-throughput Wi-Fi technologies such as IEEE 802.11ac suffer from significant energy penalties when retransmissions are required to deliver large packets reliably.

To address this, *CRT-based fragmentation* offers a promising strategy [10]. By splitting a large Data packet into smaller, redundant, and independently transmittable fragments, CRT reduces the per-fragment error probability and thus the overall need for retransmissions. This improves delivery success rates and significantly lowers energy expenditure per successfully delivered bit.

Beyond transmission reliability, CRT's support for content reconstruction from a subset of fragments enhances resilience: fragments can be distributed across multiple devices, increasing the probability that enough of them are available locally for successful content retrieval. This spatial redundancy strengthens the robustness of content dissemination, particularly in networks characterized by intermittent connectivity and minimal infrastructure support.

Another of the advantages of using the CRT-based fragmentation scheme is the increased level of security it provides, because each transmitted fragment contains only partial, non-intelligible remainders rather than the full original data [8]. This implies that successful reconstruction requires intercepting a sufficient number of fragments. If these fragments are transmitted through different paths or originate from multiple sources, as in our scenario, the interception process becomes significantly more difficult. Moreover, compared to a simple fragmentation scheme, even if adversaries intercept enough fragments, they cannot reconstruct the complete message without also knowing the corresponding set of prime numbers, which effectively serves as a secret key. In this way, CRT-based fragmentation inherently adds a layer of confidentiality, in addition to its benefits in terms of energy efficiency, reliability, and latency reduction.

3.2. The Chinese remainder theorem

CRT is a popular theorem, used in both theoretical and applied fields of mathematics and computer science for its ability to simplify computations and solve modular arithmetic problems efficiently [7]. Basically, CRT enables the decomposition of a number into smaller components, named modular remainders or residues. In particular, according to the CRT, *given an unsigned integer number $m < M$ and a set of primes $PS = \{p_1, p_2, \dots, p_{N_{CRT}}\}$ whose product is $P = \prod_i p_i$, by assuming that the condition $P \geq M$ is met, then m can be represented through a linear combinations of the modular remainders $m^{(i)} = m \pmod{p_i}$, i.e., $m = \sum_i c_i \cdot m^{(i)} \pmod{P}$.¹*

Note that $m^{(i)} \in \{0, \dots, p_i - 1\}$ so that $\lceil \log_2(p_i) \rceil < \log_2(M)$ bits are needed to represent the generic residue $m^{(i)}$.

Therefore, CRT can be used to split a number m into N_{CRT} smaller components, i.e., the set of modular remainders $m^{(i)}$ with $i \in [1, N_{CRT}]$, which in turn can be used to re-obtain m using a simple linear combination.

It is worth mentioning that reconstruction is possible even if a subset S of the PS is considered such that $\prod_{i \in S} p_i \geq M$. In this case the CRT and all related equations remain valid by simply replacing P with $P_S = \prod_{i \in S} p_i$.

Henceforth, we assume that prime numbers are consecutive primes in ascending order, i.e., $p_1 < p_2 < \dots < p_{N_{CRT}}$, and that they are the smallest prime numbers such that $P \geq M$. In this case, it is straightforward to prove that, for fixed M and N_{CRT} , there exists a unique prime set, which we refer to as the Minimum Prime Set (*MPS*).

For instance, given $M = 256$ and $N_{CRT} = 3$, the corresponding *MPS* is $\{5, 7, 11\}$, since $5 \cdot 7 \cdot 11 = 385 \geq 256$ (in fact, the previous set of $N_{CRT} = 3$ consecutive primes, i.e., $\{3, 5, 7\}$, yields $3 \cdot 5 \cdot 7 = 105 < 256$, which does not satisfy the condition).

This property can be leveraged to determine the *MPS* in a distributed manner, relying solely on the knowledge of M and N_{CRT} . The same property also enables *MPS* determination via precomputed

¹ Coefficients c_i needed for reconstruction can be straightforwardly obtained from the PS as follows: $c_i = Q_i \cdot Q_i^{-1} \pmod{P}$ where $Q_i = \frac{P}{p_i}$ and Q_i^{-1} is its modular inverse, i.e. it solves $Q_i Q_i^{-1} \pmod{p_i} = 1$.

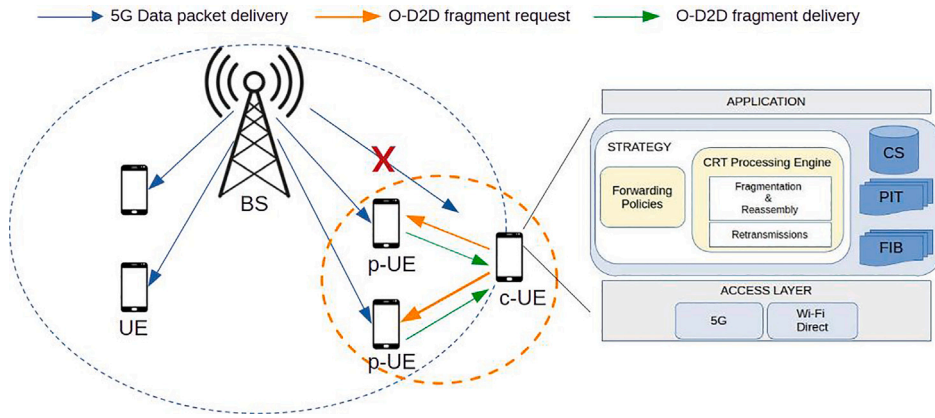


Fig. 3. Reference scenario: the BS transmits Data packets over 5G, while UEs cooperate via Outband D2D using Wi-Fi Direct.

lookup tables where M and N_{CRT} serve as lookup keys, as illustrated in [8].

Henceforward, we assume that $M = 2^w$ and that m can be represented as t 16-bit words, where $t = \lceil \frac{w}{16} \rceil$.

Under these conditions, if N_{CRT} is chosen such that $p_{N_{CRT}} < 2^b$ with $b \leq 16$, it can be proven that, the complexity of computing $m \bmod p_i$ on a 16-bit processor using the Euclidean algorithm [29] is $O(t)$. For reconstruction purposes, Garner's algorithm [7] can be employed with complexity $O(t \cdot b^2)$.

4. Name-based O-D2D communications with CRT

4.1. Reference scenario

The reference scenario, illustrated in Fig. 3, considers a 5G network enhanced with D2D cooperation for content dissemination. At the center of the architecture is a BS responsible for transmitting contents to a set of UEs, including traffic updates, map tiles, public safety notifications, and multimedia files for infotainment. Each UE is equipped with two wireless interfaces: a 5G interface for communicating with the BS, and a Wi-Fi Direct interface for establishing peer-to-peer, O-D2D links with nearby UEs. Without loss of generality, we assume that a recent release of the IEEE 802.11 standard is used, such as IEEE 802.11ax (Wi-Fi 6), which offers high throughput, enhanced spectral efficiency, and robust support for dense D2D environments.

UEs implement NDN as the network layer and, accordingly, the BS transmits contents in the form of NDN Data packets, which are then handled by UEs using the standard NDN processing fabric. The dual-interface setup allows UEs to act as both consumers and providers of contents. Devices with reliable 5G connectivity receive contents directly from the BS and cache them to serve future requests. A cache-everything strategy with LRU replacement is assumed as the default policy. More sophisticated cache management approaches could be adopted to further enhance performance, such as popularity-aware or cooperative strategies; however these aspects are beyond the scope of this paper.

UEs at the edge of BS coverage, or experiencing degraded signal quality, rely on O-D2D communications to retrieve missing data from neighboring UEs.

To improve transmission reliability and energy efficiency over D2D links, UEs can divide the received Data packets into CRT fragments, which can then be distributed, upon request, via Wi-Fi Direct to requesting devices, as explained in the following.

4.2. The CRT processing engine

To enable CRT packet fragmentation and reconstruction, we introduce a *CRT Processing Engine* (CPE) within the Strategy module of the NDN Data Plane of each UE, see Fig. 3. The CPE is responsible for both the fragmentation of Data packets (at the caching UEs) and the reassembly of fragments (at the consumer UE).

A consumer UE (c-UE) experiencing poor connectivity and failing to receive certain Data packets from the BS does not rely on re-transmissions over the infrastructure network but instead activates the O-D2D recovery mechanism over Wi-Fi Direct. The decision is based on the estimated quality of the O-D2D channel. Specifically, the c-UE maintains a short-term PER estimate of Wi-Fi Direct transmissions, which can be derived from standard MAC-layer retry/ACK statistics available in 802.11ax. A threshold-based policy can be adopted: if the PER estimation, computed over the last 20 packets [30], remains below a threshold (e.g., around 8% in our evaluated setting, with $N_{CRT} = 4$ and $f = 1$), the c-UE does not request CRT-based delivery, since the overhead would outweigh the benefit. If the estimated PER rises above the threshold, the c-UE activates the CRT mechanism to accelerate fragment retrieval and improve energy efficiency.

More in detail, the c-UE broadcasts a special request packet called a *Long-Lived Interest* (LLI) [31]. Unlike the standard NDN Interest, it consumes a single Data packet, enforcing a strict one-to-one correspondence, a LLI has a longer lifetime and it can be used to retrieve more Data packets. LLIs were originally conceived in [32] to support efficient multi-party audio conferencing in NDN, by allowing a single Interest to retrieve multiple sequential Data packets without repeated requests. Building on this idea, LLIs have subsequently been extended to more general push-based communication models and multi-source content dissemination in wireless scenarios, where multiple data producers or caching nodes can respond to a single outstanding request [31,33,34].

In our design, a LLI remains active for a longer period in the PIT of the c-UE, allowing the collection of multiple CRT fragments. Broadcasting allows multiple nearby UEs to opportunistically respond with distinct CRT fragments, enabling the consumer to retrieve them from different sources. This multi-source retrieval minimizes the number of transmitted Interests, thereby reducing energy consumption at the c-UE side, and parallelizes the content dissemination across several provider UEs (p-UEs), thus lowering their individual energy consumption. LLI transmission is always restricted to single-hop communications, i.e., LLIs are not forwarded beyond the local neighborhood. As a result, the PIT state introduced by LLIs is confined to the single hop D2D domain, and does not propagate through the network, limiting the retrieval process to nearby nodes.

The LLI carries the name of the requested Data packet, extended with a dedicated /CRT tag appended as the final component of the Interest name. The presence of the tag explicitly signals that the requester

expects to receive CRT-compliant fragments rather than standard NDN Data packets. CRT fragmentation of a Data packet is applied on demand, triggered by the reception of the first LLI. The p-UE fragments the packet, assigns each fragment the original content name extended with an *auxiliary header* for identification and reconstruction, caches the resulting fragments, and discards the original Data packet. This reactive strategy avoids storing both the original Data packet and its fragments in the CS, thus reducing memory usage and preventing unnecessary duplication of cached contents, which is particularly relevant in D2D scenarios characterized by limited storage and energy resources. Alternative designs are nevertheless possible. For example, a node could retain both the original Data packet and its fragments, or perform on-demand reassembly if a standard Interest is later received. A systematic evaluation of these alternative strategies, including their impact on memory usage and cache efficiency, is however beyond the scope of the current paper and will be addressed in future works.

4.3. Fragments retrieval

Upon receiving the LLI, each neighboring UE checks its CS to locate the requested Data packet or any pre-existing fragments:

- If a matching Data packet is found, the CPE is triggered to perform CRT-based fragmentation: the payload is split into multiple fragments. Each fragment is encapsulated in a Data packet augmented with an auxiliary CRT header, containing the necessary parameters to support correct reassembly at the consumer. The fragments are then stored in the CS, and one of them is randomly selected and scheduled for transmission over the O-D2D link. The original Data packet is discarded to free up memory.
- If the Data has already been fragmented and the p-UE has stored the fragments in its CS, one fragment is randomly selected and scheduled for transmission.
- If no matching content is found in the CS, the LLI is discarded and the UE remains silent.

To reduce the risk of fragment collisions, p-UEs apply a randomized defer timer with overheard, computed as $T_d = \text{rand}(0, T_{\max})$, where:

- $\text{rand}(0, x)$ represents a uniform random variable sampled from the interval $[0, x]$, ensuring variability in transmission timing;
- T_{\max} corresponds to the upper bound of the IEEE 802.11 Contention Window (CW) multiplied per a slot time, ensuring compatibility with the MAC layer and effective collision avoidance.

After selecting the fragment to send, the p-UE computes the defer time and waits. If it overhears the same fragment transmitted by another node, then it stops the waiting time, cancels its own transmission, schedules another fragment transmission and resumes the defer time.

At the c-UE side, the CPE collects incoming CRT fragments, temporarily stores them in a dedicated reconstruction buffer, and monitors the set of received moduli. Although the defer-based overheard mechanism helps reduce redundant transmissions, the hidden terminal problem may still lead to duplicate fragment receptions. In such cases, duplicates are detected and discarded.

Once the minimum number of distinct CRT fragments required for decoding has been collected, the CPE reconstructs the original Data packet. The recovered content is then passed to the Application layer through the standard NDN pipeline.

If decoding is not yet possible, the c-UE issues a new LLI including a Fragment Bitmap (FB) field. This field is a binary vector of length N_{CRT} , where each bit corresponds to one fragment in the original set. A bit is set to 1 if the corresponding fragment has already been received, and 0 otherwise. The FB allows responding UEs to avoid sending already obtained fragments, thereby reducing redundancy, saving energy, and improving the efficiency of the recovery process.

4.4. Fragmentation and redundancy

Fragmentation. Each NDN Data packet, D , transmitted by the BS, consists of four main components, as defined in the official NDN packet specification [35]: (i) the *Name field*, which uniquely identifies the Data packet and serves as the key for CS lookups; (ii) an optional *MetaInfo* field, specifying the content type and its freshness period; (iii) the *Payload*, containing the actual content; and (iv) a *Signature*, used to ensure data authentication and integrity verification.

CRT is applied to all fields of the Data packets, excluding the header information that ensures compatibility with legacy NDN caching and forwarding routines, e.g., content name and freshness period. For the sake of simplicity, henceforward we identify this information as the header (H).

We assume that p-UEs divide a Data packet D into N_{CRT} fragments, $\{D^{(1)}, D^{(2)}, \dots, D^{(N_{CRT})}\}$, each inheriting the original header H followed by an auxiliary header field (H_{CRT}) that is needed to enable reconstruction at the c-UE. We further assume that the payload of the Data packet is logically divided into N binary words of w bits each, i.e., $\{m_1, \dots, m_N\}$, and that CRT is applied to every word m_k so that the payload of the i th fragment contains the residues $m_k^{(i)} = m_k \pmod{p_i}$ with $k \in [1, N]$. A fragmentation example with $N_{CRT} = 3$ is reported in Fig. 4. After receiving the LLI, a fragment $D^{(i)}$ is transmitted as an independent Data packet.

Selection of prime numbers. We assume that prime numbers p_i with $i \in [1, N_{CRT}]$ are chosen as the smallest consecutive primes that satisfy the condition $\prod_{i=1}^{N_{CRT}-f} p_i > 2^w$ where $f < N_{CRT}$ is a design parameter.

In this case, according to the CRT, each word m_k can be recovered from every subset $\{m_k^{(i)}\}$ of at least $N_{CRT} - f$ residues, therefore the original Data packet can be reconstructed once at least $N_{CRT} - f$ fragments are collected. Here, f represents the number of allowed failures, reflecting the robustness of the CRT-based approach. Indeed, CRT enables successful packet reconstruction even in the presence of fragment losses (i.e., failures) by ensuring that the information required to reconstruct the original packet is distributed across the fragments. Finally, note that the overall payload length of the i th fragment is $L_i = N \cdot \lceil \log_2(p_i) \rceil \approx \frac{N \cdot w}{N_{CRT} - f} < N \cdot w$ bits, i.e., fragments are smaller than the original packet.

We assume that the BS distributes the CRT configuration parameters, namely w , N_{CRT} , f , and the ordered list of prime numbers, thus guiding the fragmentation process. This information can be disseminated using PC5 sidelink broadcast signaling, leveraging mechanisms similar to those adopted for group communications in 5G NR [36]. As a result, each fragment only needs to carry the index i of the prime number p_i in the small additional header field of the fragment, namely H_{CRT} , instead of including the full prime value or complex reconstruction metadata.

CRT Redundancy. CRT fragmentation introduces a redundancy that is related to the number of tolerable failures f , i.e., with the proposed fragmentation scheme, $r_{CRT} > 1$ bits must be transmitted for each bit of the original Data packet. It is straightforward to prove that, if the header length h of the original Data packets is negligible in comparison to the payload length $N \cdot w$, the CRT redundancy can be approximated by $r_{CRT} \approx \frac{N_{CRT}}{N_{CRT} - f}$. In particular, the above approximation is valid and tight if the condition $N \cdot w \gg N_{CRT} \cdot (h + h_{CRT})$ holds true. The proof is omitted for the sake of space. For instance, in the case of five fragments ($N_{CRT} = 5$) and one possible failure ($f = 1$), the expected CRT redundancy is $r_{CRT} \approx 1.25$, i.e., the number of bits that must be sent using the proposed CRT fragmentation scheme is only 25% more than the number of bits needed to send original Data packets without fragmentation.

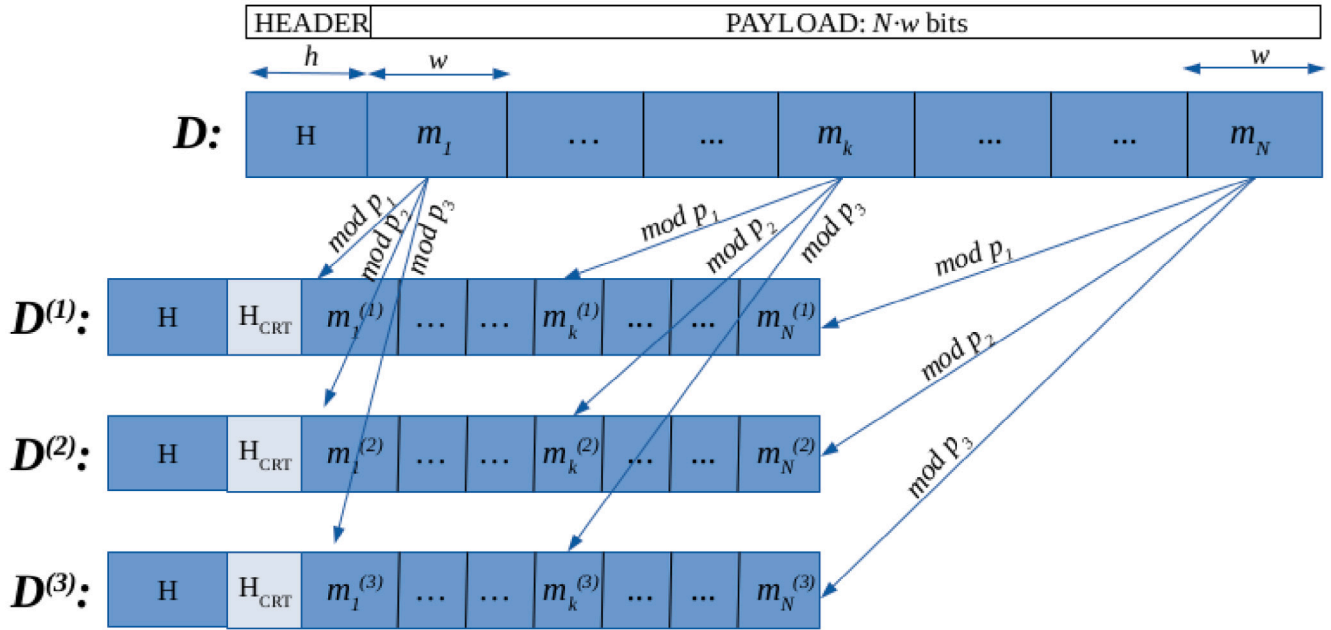


Fig. 4. Example of packet fragmentation generating $N_{CRT} = 3$ fragments.

5. Performance evaluation

To assess the performance of the proposed CRT-based O-D2D retrieval framework, we carried out an extensive simulation campaign using ndnSIM v2.9, the official simulation environment maintained by the NDN community [14]. We extended the NDN Forwarding Daemon (NFD) in ndnSIM to support LLIs and multi-source fragment retrieval mechanisms compliant with the envisioned CRT architecture. In particular, the CPE is implemented as an extension of the Forwarding Strategy, which can co-exist with the other one already available.

5.1. Simulation settings

To model D2D interactions, Wi-Fi Direct connectivity is assumed to comply with the IEEE 802.11ax standard, operating in the 2.4 GHz band with a channel width of 20 MHz [37]. We assume Best Effort access category in Enhanced Distributed Channel Access (EDCA), with $CW_{min} = 15$ and $CW_{max} = 1023$, and a default transmission queue length of 100 packets. Multi-hop forwarding is disabled, and D2D communication is restricted to direct, single-hop exchanges among neighboring UEs. Therefore, no NDN routing protocol is implemented.

Wireless channel and propagation are based on YansWifiPhy/YansWifiChannel with a Nakagami- m fading, which models small-scale fading effects. This configuration is a flat (per-packet) abstraction: it does not capture frequency-selective propagation and, in our study, does not include spatial correlation across links/antennas. However, these choices are common in network-layer evaluations to keep large scenarios tractable and to focus on packet-error-rate-driven request/response dynamics.

To reflect different conditions in outdoor urban scenarios, we consider two configurations of the Nakagami shape parameters m_0 , m_1 , and m_2 , which define the fading severity over different distance ranges:

- Rayleigh fading, representative of non-line-of-sight (NLOS) urban environments, with $m_0 = m_1 = m_2 = 1$;
- An extremely severe fading scenario, with $m_0 = m_1 = m_2 = 0.5$. In this setting, signal power exhibits greater variance, leading to

frequent deep fades and unreliable wireless links. Such a configuration is intentionally selected as a stress-test scenario to highlight the robustness and fault tolerance of the proposed CRT-based approach under adverse channel conditions.

We consider a population of 100 UEs located within the coverage area of a single 5G cell, assuming a typical urban radius of 200 m [38]. UEs move according to the Truncated Lévy Walk (TLW) mobility model, which has been shown to realistically replicate human mobility patterns in urban environments [39]. The pedestrian speed is set in the range between 1 m/s and 1.5 m/s. A variable subset of UEs, ranging from 4 to 24, is selected to act as D2D consumers requesting content. Interest packets are broadcast within the local neighborhood, and any node possessing the requested content may respond with the corresponding Data packet. The transmit power is set to 15 dBm, and the receiver sensitivity threshold is aligned with the IEEE 802.11ax specifications for MCS 7 operation (approximately -70 dBm) [40]. Following the parameters reported in [41], the instantaneous current draw is set to 550 mA for transmission and 107 mA for reception.

Each simulation run spans 1000 s of simulated time. Results are averaged over 50 independent seeds to ensure statistical significance and reported with 95% confidence intervals. Moreover, we compute the standard deviation across the simulation runs.

5.2. Benchmarks and metrics

We compare the proposed solution against two benchmark schemes. First, we consider the D2D-NDN approach described in [24] (simply referred to as NDN in the following), which assumes that requested content items are cached at D2D-capable devices, and the BS coordinates content delivery by instructing the consumer to retrieve data from a nearby provider. The standard one-Interest-per-one-Data-packet exchange is implemented in this case, without fragmentation. Second, we consider a baseline NDN fragmentation scheme, based on [25] (denoted as NDN+frag in the following). In this variant, the consumer issues a standard Interest, while the provider fragments the corresponding Data packet into smaller units and transmits them sequentially. Unlike

our proposal, this scheme does not include redundancy or support multi-source fragment retrieval.

We consider a content catalog of 1000 items. Each content is divided into 1000 NDN Data packets, each carrying a payload of 200 bytes, and represents data typically associated with proximity-based services, such as local point-of-interest information, traffic alerts, or public event notifications. In line with existing literature [42,43], we considered a freshness period for the Data packets in the range [200–500] seconds.

The request pattern generated by consumers follows a Zipf distribution, with the skewness parameter z set to either 2 or 0.8, respectively [44]. A value of $z = 2$ models a highly skewed demand, where a small subset of contents dominates the request pattern, while $z = 0.8$ represents a more uniform distribution, with a flatter popularity profile across the catalog.

In the CRT-based approach, each Data packet is further fragmented at the provider's Strategy Layer into smaller units. We assume a word size of $w = 50$ bits, a fault tolerance level of $f = 1$, and two distinct values for N_{CRT} , namely 4 and 5 fragments. These parameters were selected to maintain a balance between redundancy and efficiency, ensuring resilience to packet loss with minimal transmission overhead, in term of header bytes and more frames on the air, while maintaining a manageable word size to avoid excessive computational complexity. Fragments are treated as independent Data packets: once the expected number of fragments has been collected at the NDN Data Plane, the original Data packet is reconstructed and delivered to the consumer's application layer.

In the case of NDN-frag, we assume the original Data packet is divided into 4 distinct fragments, but no redundancy is applied to cope with losses.

Unless otherwise specified, in all schemes the c-UEs generate Interests using the window-based congestion control model described in [45], which dynamically adapts the sending rate to current network conditions. This choice provides a more realistic representation of content retrieval dynamics compared to fixed-rate models.

We selected an initial LLI lifetime of 500 ms, which acts as a conservative upper bound value, sufficiently large to accommodate MAC-layer contention, randomized defer times, and the collection of multiple fragments from neighboring providers, while remaining short enough to avoid excessive PIT occupancy. The lifetime also serves as a retransmission timeout, at the c-UE side, which can be dynamically adapted based on observed retrieval dynamics [45]. The c-UE monitors the round-trip time associated with the fragment collection process and updates the timeout using a smoothed estimator (i.e., an exponential moving average), in line with existing NDN approaches for Interest retransmission.

As a caching policy, we consider the legacy *cache everything* with LRU replacement. The capacity of the CS is set to 5% of the content catalog size [46]. At the beginning of the simulation, only three UEs are randomly selected to act as initial providers, reflecting a sparse but sufficient content availability within the D2D range. In the benchmark NDN approach, they serve the entire Data packets of 200 bytes, while in the CRT-NDN approach, they answer with fragments. Consumers cache incoming packets and may become themselves content providers over time.

To evaluate the performance of the proposed solution, we consider the following metrics:

- *Content retrieval delay.* This metric represents the average time required to retrieve the entire content. It is measured from the instant the first request (either an LLI or a legacy Interest) is transmitted, until the final Data packet is received at the consumer's application layer.
- *Request retransmissions.* This metric quantifies the total number of LLI or legacy Interest retransmissions issued by D2D consumers, expressed as a percentage of the total number of initial requests. It captures the robustness of CRT to handle packet delivery in the presence of harsh channel conditions.


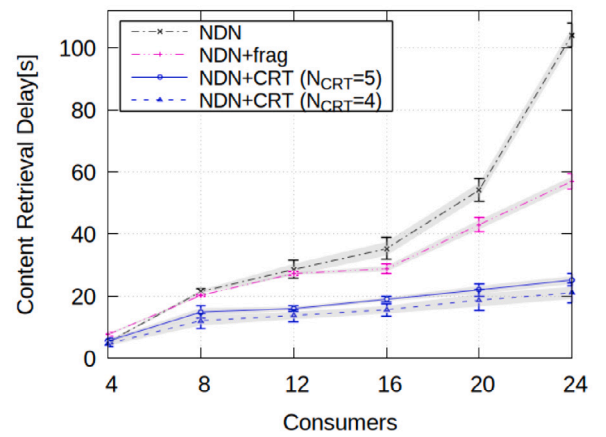
(a) $z = 2$.(b) $z = 0.8$.

Fig. 5. Content retrieval delay when varying the number of consumers and the Zipf skewness parameter, under Rayleigh fading conditions (vertical error bars indicate the confidence intervals, while shaded gray bands denote the standard deviation across the simulation runs).

- *Consumer energy cost.* This metric represents the average energy consumed per bit, in Joules, by nodes acting as content consumers throughout the retrieval session. It is computed by dividing the total energy consumed by each consumer by the number of content bits successfully received. This allows for a normalized assessment of energy efficiency during content acquisition.
- *Provider energy cost.* This metric captures the average energy consumed per bit, in Joules, by nodes acting as content providers. It is computed by dividing the total energy consumed during the simulation by the total number of bits corresponding to the original content demand across all consumers. This reflects the energy efficiency of the data dissemination process from the provider's perspective in the network.

5.3. Network performance results

5.3.1. Retrieval delay analysis

Fig. 5 reports the content retrieval delay metric as a function of the number of D2D consumers, when considering Rayleigh fading conditions and two distinct content popularity profiles. In both cases, the proposed NDN+CRT scheme achieves the lowest delay across the entire load range, while NDN+frag attains intermediate performance and the

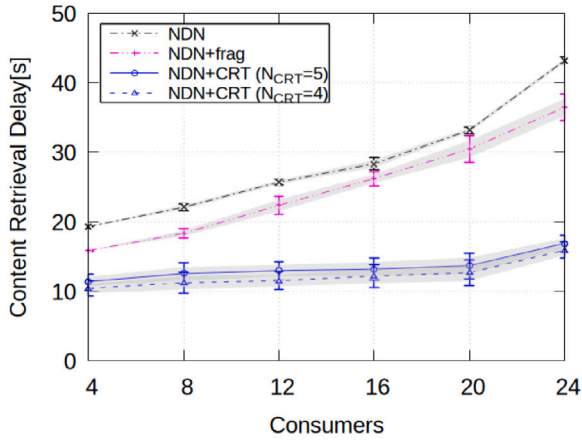
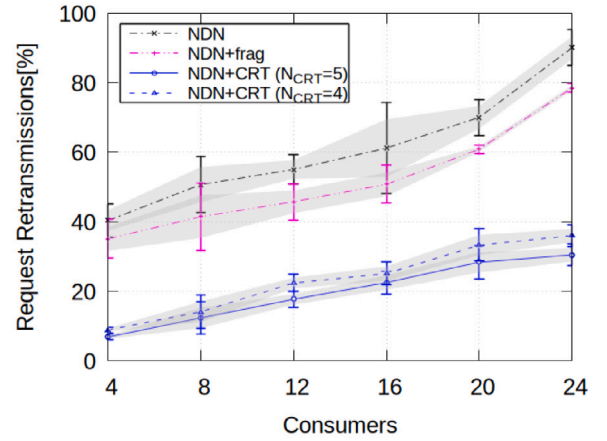

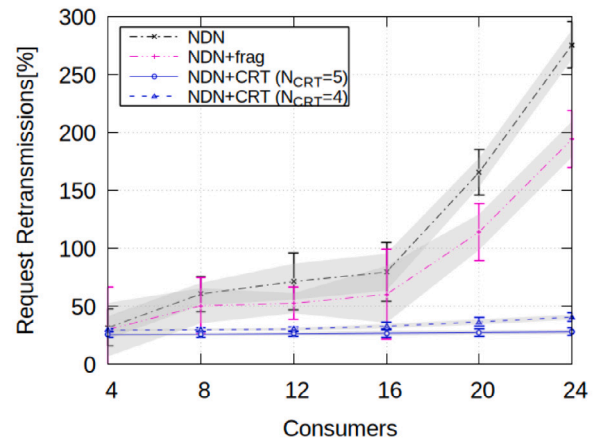
(a) $z = 2$.(a) $z = 2$.(b) $z = 0.8$.(b) $z = 0.8$.

Fig. 6. Content retrieval delay when varying the number of consumers and the Zipf skewness parameter, under Nakagami- m fading conditions (vertical error bars indicate the confidence intervals, while shaded gray bands denote the standard deviation across the simulation runs).

baseline NDN exhibits the largest delays. The benefit of CRT is already evident at moderate load and becomes more pronounced as the number of consumers grows, consistent with the reduced retransmissions and shorter on-air times provided by fragment redundancy and multi-source retrieval. Using a larger number of residues, i.e., $N_{CRT} = 5$, slightly increases the delay, because more fragments must be transmitted; nonetheless both configurations remain clearly superior to NDN and NDN+frag.

As expected, the popularity profile strongly affects latency. With a highly skewed catalog ($z = 2$, Fig. 5(a)), requests concentrate on a small subset of items, which amplifies cache reuse and keeps delays low for all schemes, with NDN+CRT nearly flat with load. When the distribution is flatter ($z = 0.8$, Fig. 5(b)), content diversity increases, reducing cache hits and raising contention on the wireless channel; all schemes see higher delays, yet NDN+CRT preserves a gentle growth and a substantial margin over the benchmarks.

Figs. 6(a)–(b) repeat the analysis for the more severe Nakagami- m channel. Delays are higher than under Rayleigh, reflecting the harsher propagation conditions. Nonetheless, the ordering of the schemes is unchanged: NDN+CRT provides the lowest delays, NDN+frag is intermediate, and NDN is the worst, particularly at high load. Again, increasing N_{CRT} from 4 to 5 yields a small delay penalty due to the

Fig. 7. Percentage of request retransmissions – LLI packets for NDN+CRT and Interest packets for NDN – when varying the number of consumers, under Rayleigh fading conditions (vertical error bars indicate the confidence intervals, while shaded gray bands denote the standard deviation across the simulation runs).

additional fragments, but the CRT curves remain well separated from the baselines. The gap between CRT and the other approaches widens with the number of consumers, indicating that CRT's tolerance to losses and ability to reconstruct Data from a subset of fragments effectively mitigates the collision-induced retransmissions that dominate at high offered load.

5.3.2. Retransmission analysis

To further assess the impact of losses and the system's ability to cope with them, we report the percentage of Interest retransmissions over the total number of original Interest packets for NDN and NDN+frag, and LLI packets for NDN+CRT. Results under Rayleigh fading, for $z = 2$ and $z = 0.8$, are shown in Fig. 7, while Fig. 8 reports the Nakagami- m case.

As expected, the NDN baseline exhibits a higher growth in retransmission rate as the number of consumers increases and channel conditions worsen. The increase is more pronounced with a flatter popularity profile ($z = 0.8$; Figs. 7(b), 8(b)) due to the reduced cache reuse compared to the highly skewed case ($z = 2$; Figs. 7(a), 8(a)). NDN+frag benefits from smaller packets and thus incurs fewer retransmissions than NDN; however, lacking redundancy, it must still




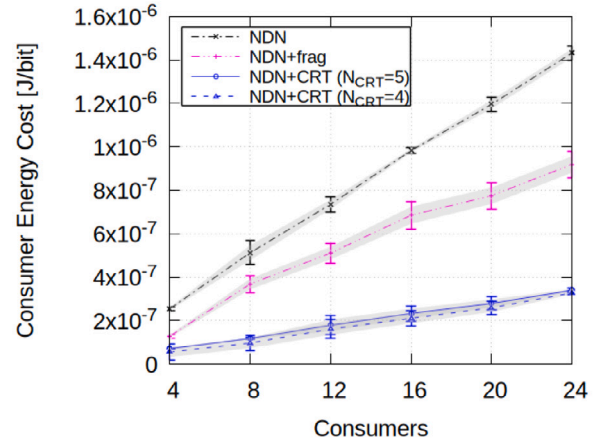
(a) $z = 2$.(a) $z = 2$.(b) $z = 0.8$.(b) $z = 0.8$.

Fig. 8. Percentage of request retransmissions – LLI packets for NDN+CRT and Interest packets for NDN – when varying the number of consumers, under Nakagami- m fading conditions (vertical error bars indicate the confidence intervals, while shaded gray bands denote the standard deviation across the simulation runs).

retransmit whenever any fragment is lost, which limits its robustness at higher load.

In contrast, NDN+CRT consistently attains the lowest retransmission percentages across all configurations. Two factors drive this advantage: (i) shorter fragments improve per-packet success probability, and (ii) redundant residues enable successful reconstruction once a sufficient subset is received, avoiding repeated requests for missing pieces. Moreover, within NDN+CRT, using a larger number of residues ($N_{CRT} = 5$) further reduces retransmissions compared to $N_{CRT} = 4$, with the gap widening at higher consumer counts where contention and collisions are more frequent. This trend holds under both Rayleigh and Nakagami- m fading, although retransmission levels are higher in the latter due to more severe propagation.

Overall, the results confirm that CRT-based retrieval not only tolerates wireless losses but also mitigates the need for request retries, particularly in dense, interference-prone D2D settings, by combining small packet sizes with loss-tolerant multi-source reconstruction.


5.4. Energy performance results

Fig. 9 and Fig. 10 report the average energy cost per bit for consumer and provider nodes, respectively, in the presence of Rayleigh

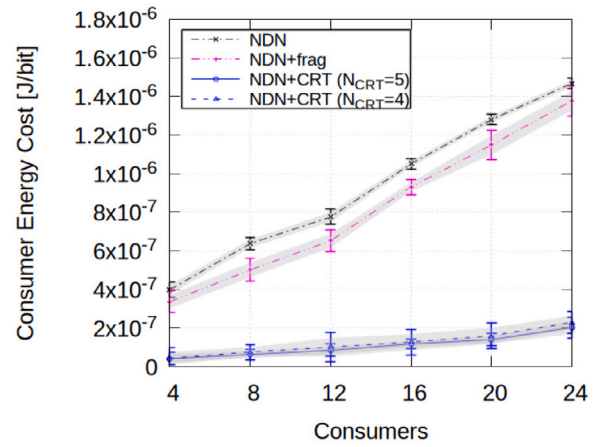
Fig. 9. Consumer energy cost metric when varying the number of consumers and the Zipf skewness parameter, under Rayleigh fading conditions (vertical error bars indicate the confidence intervals, while shaded gray bands denote the standard deviation across the simulation runs).

fading, under varying numbers of D2D consumers and Zipf parameters. Conversely, Fig. 11 and Fig. 12 report the same metrics under Nakagami- m fading.

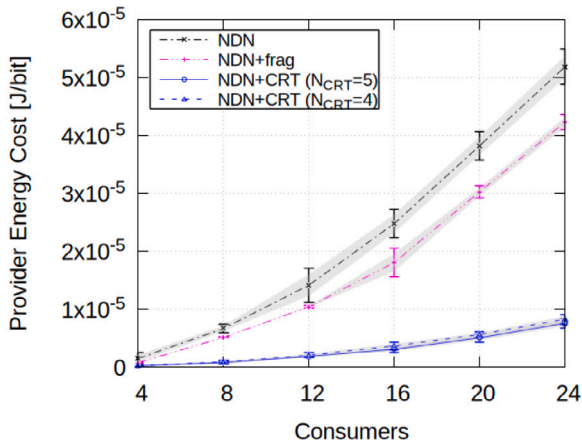
As expected, the NDN approach shows a higher increase in energy consumption at both the consumer and provider sides, especially under the severe Nakagami- m fading configuration. This behavior results from frequent packet losses and subsequent retransmissions, which increase the energy required to successfully retrieve or deliver each content bit. NDN+frag achieves intermediate results: the smaller fragment size reduces individual packet losses, improving energy efficiency compared to the unfragmented baseline. However, the lack of redundancy and the need to retrieve all fragments sequentially from a single provider prevent further gains, especially under poor channel conditions. In contrast, the proposed NDN+CRT approach significantly reduces energy expenditure. On the consumer side, the use of LLIs reduces the number of transmitted requests, while CRT-based reconstruction allows content recovery from partial fragment sets. This substantially lowers the energy consumption related to packet transmission and reception operations. On the provider side, the CRT mechanism reduces energy usage in two key ways. First, each fragment has a smaller payload, leading to shorter and more energy-efficient transmissions. Second, the retrieval load is distributed among multiple neighboring providers,



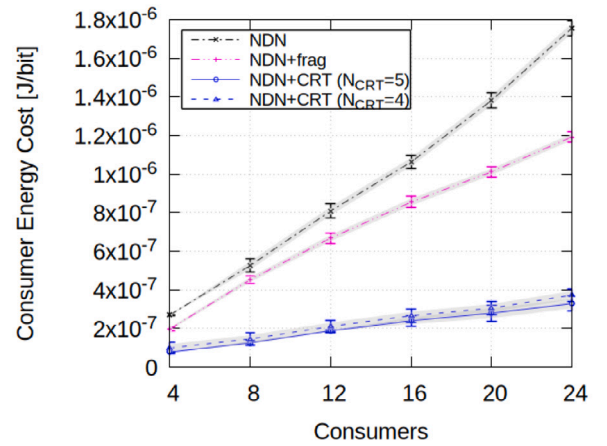
(a) $z = 2$.



(a) $z = 2$.



(b) $z = 0.8$.



(b) $z = 0.8$.

Fig. 10. Provider energy cost metric when varying the number of consumers and the Zipf skewness parameter, under Rayleigh fading conditions (vertical error bars indicate the confidence intervals, while shaded gray bands denote the standard deviation across the simulation runs).

Fig. 11. Consumer energy cost metric when varying the number of consumers and the Zipf skewness parameter, under Nakagami- m fading conditions (vertical error bars indicate the confidence intervals, while shaded gray bands denote the standard deviation across the simulation runs).

since a single LLI may elicit responses from different nodes. This cooperative fragment delivery avoids overburdening individual providers and improves energy balance across the network. Moving from $N_{\text{CRT}} = 4$ to $N_{\text{CRT}} = 5$ yields a slight energy saving on the c-UE side due to fewer retransmissions, despite one additional useful fragment needs to be received. Similarly, on the provider side, load sharing among sources and the reduction in retries largely offset the extra fragment, leading to per-provider energy consumption that is slightly lower.

Energy consumption is generally higher when $z = 0.8$. In this case, the content request pattern is more diverse and fewer opportunities for cache reuse exist, compared to the case with $z = 2$. This increases the transmission burden and leads to higher energy costs across all schemes, particularly under severe fading. Nonetheless, CRT-based retrieval remains the most energy-efficient strategy.

Overall, CRT-based fragmentation enhances energy efficiency at both ends of the communication, confirming its suitability for dense, energy-constrained O-D2D scenarios. These results also highlight the ability of NDN+CRT to maintain scalable and energy-aware performance in adverse network conditions.

5.5. Saturation analysis

To better understand the operating limits of the proposed scheme compared to the baselines, we investigate the network *saturation point*, i.e., the maximum request load that the system can sustain before delays increase dramatically. For this study, we adopt a Constant Bit Rate (CBR) request model, where each consumer issues Interests (or LLIs) at a fixed rate. This setup provides a controlled and deterministic offered load, in contrast to the adaptive window-based retrieval employed in the previous experiments.

To stress the system, we focus on the harshest propagation condition, namely Nakagami- m fading, and set the Zipf skewness parameter to $z = 0.8$. These settings reduce cache reuse and increase the probability of wireless errors, thereby accelerating the onset of saturation. Node positions are assigned deterministically to control the impact of network topology on the results. Specifically, we assume that c-UEs are placed uniformly on the circumference of a circle with radius 60 m, while the provider(s) are located at the center. All nodes employ a constant-position mobility model, so that the propagation conditions are determined exclusively by the channel model. This configuration provides a reproducible and worst-case layout, with relatively long wireless links subject to severe fading, and is therefore well suited for investigating the saturation point of the system.



(a) $z = 2$.(b) $z = 0.8$.

Fig. 12. Provider energy cost metric when varying the number of consumers and the Zipf skewness parameter, under Nakagami- m fading conditions (vertical error bars indicate the confidence intervals, while shaded gray bands denote the standard deviation across the simulation runs).

In this scenario, we set the number of c-UEs to 15 and vary the per-consumer Interest generation rate. We measure the resulting average Round-Trip Time (RTT), i.e., the time between the transmission of the Interest (or LLI) and the reception of the corresponding Data packet (or the last fragment, for the NDN+CRT approach). Therefore, unlike the content retrieval delay metric, which captures the time required to download the entire content composed of multiple Data packets, this metric focuses on a single Interest-Data exchange. Saturation is identified as the point where the mean RTT shows an abrupt growth and diverges, indicating that content retrieval becomes infeasible.

Fig. 13 reports the RTT when varying the request rate for NDN, NDN+frag, and NDN+CRT approaches. For NDN+CRT we considered $N_{CRT} = 4$, which proved to be more effective in terms of retrieval delay. It can be observed that, for the NDN benchmark, saturation occurs at around 250 Interests/s, beyond which the D2D connections fail to sustain timely delivery and the RTT grows unbounded. This behavior is explained by the larger packet size and the higher probability of collisions, which penalize the baseline scheme compared to the fragmentation-based approaches. In contrast, NDN+CRT sustains significantly higher request rates before saturation, also outperforming NDN+frag, confirming its robustness under harsh conditions.

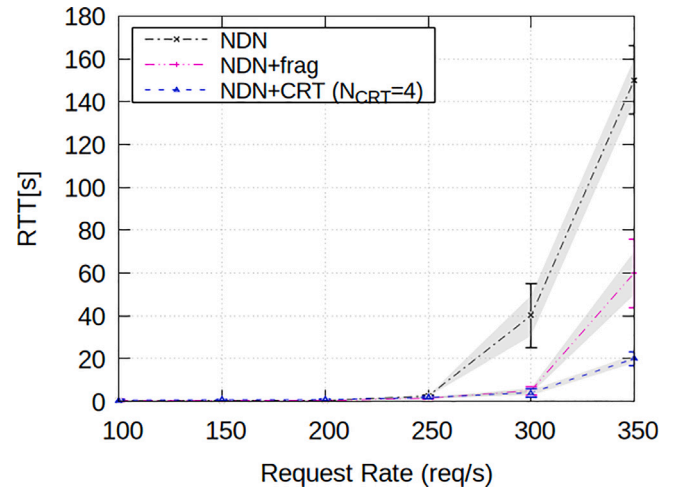


Fig. 13. RTT when varying the request (Interest or LLI) rate.

Table 1

Equivalent 64-bit additions needed for CRT reconstruction for different values of N_{CRT} and f when $w = 50$.

N_{CRT}	$f = 1$	$f = 2$	$f = 3$
5	83	88	–
6	94	99	106
7	106	111	117
8	112	121	130
9	125	128	134
10	140	140	142

5.6. Computational footprint analysis

To give further insight into the footprint introduced by the CRT-based operations, we focus on the CRT reconstruction cost and related memory footprint. We explicitly consider the reconstruction phase as it represents the dominant computational component of the CRT process, while fragmentation only involves modular reductions and data slicing, whose computational cost is negligible compared to reconstruction and packet transmission delays.

We consider the word size $w = 50$, and we test different values of N_{CRT} and f when executing the reconstruction of CRT fragments.

For each pair of N_{CRT} and f values, the average CRT reconstruction time is determined by performing 50,000 reconstructions using the Garner algorithm [7], which is a standard and widely adopted method for practical CRT reconstruction and is implicitly assumed in many CRT-based networking applications [12].

Fig. 14 visualizes achieved reconstruction times in microseconds, providing a clearer comparison across various configurations. As it is possible to observe, for the investigated configurations, reconstruction times range from 6 to 10 microseconds, which is several orders of magnitude smaller than the content retrieval time dynamics and therefore does not affect system performance. Table 1 reports the average execution times, expressed in terms of equivalent 64-bit additions. With this aim, the average execution time for 64-bit additions was measured by averaging the runtimes of 10 million 64-bit sums, and this value was used as a normalization factor to scale the execution time of the Garner algorithm.

Starting from the results reported in Table 1, a linear regression model was employed to estimate the number of equivalent 64-bit additions (N_A) required by Garner's CRT reconstruction algorithm as a function of the number of CRT components (N_{CRT}) and the parameter




Fig. 14. CRT reconstruction time for different values of N_{CRT} and f when $w = 50$.

f . The resulting model is expressed as:

$$N_A = 27.7 + 10.2 \cdot N_{CRT} + 5.3 \cdot f$$

where the coefficients were obtained using ordinary least squares (OLS) regression, i.e., minimizing the sum of squared errors across all available measurements. The model achieves a maximum absolute error of 5 equivalent additions and a mean squared error of about 5.18, confirming its adequacy for prediction.

This formulation captures the approximately linear growth of the reconstruction cost with respect to both parameters, providing an interpretable relationship between algorithmic complexity and design choices.

Regarding memory usage, each CRT fragment stores a few residues of w bits plus minimal metadata, resulting in a memory overhead that grows linearly with N_{CRT} and remains limited to a few hundred bytes in the evaluated scenarios. Overall, the computational and memory overhead introduced by CRT is negligible compared to the gains in delivery robustness and reduced retransmissions.

6. Conclusions and future works

In this paper, we proposed a novel framework that integrates NDN with CRT-based fragmentation to support resilient content retrieval in O-D2D networks. By introducing LLs and enabling multi-source fragment collection, our approach enhances both reliability and efficiency in lossy wireless environments. Extensive simulations demonstrate that NDN+CRT significantly reduces content retrieval delay, request retransmissions, and energy consumption compared to existing benchmarks, particularly under severe fading conditions, varying data traffic and increasing consumer density.

The performance gains observed are attributed to three key features: the robustness of CRT-based reconstruction from partial data, the reduced packet size which improves transmission success rates, and the opportunistic, parallel retrieval enabled by LLs. These properties make NDN+CRT particularly suited to dense, infrastructure-less scenarios, such as urban crowdsensing or edge-assisted offloading.

Several interesting directions emerge for future investigation. First, we plan to design a CRT-aware caching policy that strategically maximizes the diversity of cached fragments across D2D nodes. By spreading different CRT fragments throughout the network, such a policy could enhance multi-source retrieval opportunities and further reduce retrieval latency and redundancy.

Another promising direction for future work is the extension of the proposed architecture to multi-hop D2D scenarios. While this paper focuses on single-hop O-D2D communications in order to isolate and evaluate the benefits of CRT-based fragmentation, extending the approach to multi-hop NDN networks would allow exploiting path diversity and in-network caching across multiple hops. In such a setting, CRT-based fragmentation could be naturally combined with name-based forwarding strategies to improve robustness against losses and intermittent connectivity, at the cost of additional design challenges related to routing, forwarding state management, and fragment coordination.

CRedit authorship contribution statement

Marica Amadeo: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Conceptualization. **Filippo Battaglia:** Writing – review & editing, Investigation, Conceptualization. **Giuseppe Campobello:** Writing – review & editing, Writing – original draft, Validation, Methodology, Investigation, Formal analysis, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

No data was used for the research described in the article.

References

- [1] A. Asadi, Q. Wang, V. Mancuso, A survey on device-to-device communication in cellular networks, *IEEE Commun. Surv. Tutorials* 16 (4) (2014) 1801–1819.
- [2] J. Liu, N. Kato, J. Ma, N. Kadowaki, Device-to-device communication in LTE-advanced networks: A survey, *IEEE Commun. Surv. Tutorials* 17 (4) (2014) 1923–1940.
- [3] K. Ganesan, 5G advanced: Sidelink evolution, *IEEE Commun. Stand. Mag.* 7 (1) (2023) 58–63.
- [4] A. Bazzi, A.O. Berthet, C. Campolo, B.M. Masini, A. Molinaro, A. Zanella, On the design of sidelink for cellular V2X: A literature review and outlook for future, *IEEE Access* 9 (2021) 97953–97980.
- [5] L. Zhang, A. Afanasyev, J. Burke, V. Jacobson, K. Claffy, P. Crowley, C. Papadopoulos, L. Wang, B. Zhang, Named data networking, *ACM SIGCOMM Comput. Commun. Rev.* 44 (3) (2014) 66–73.

- [6] W. Shang, A. Afanasyev, Y. Li, J. Burke, L. Zhang, Device-to-device communication with named data networking, in: Proceedings of the 4th ACM Conference on Information-Centric Networking, 2017, pp. 184–185.
- [7] D. Pei, et al., Chinese Remainder Theorem: Applications In Computing, Coding, Cryptography, World Scientific Publishing Company, 1996.
- [8] G. Campobello, A. Leonardi, S. Palazzo, Improving energy saving and reliability in wireless sensor networks using a simple CRT-based packet-forwarding solution, *IEEE/ACM Trans. Netw.* 20 (1) (2012) 191–205, <http://dx.doi.org/10.1109/TNET.2011.2158442>.
- [9] G. Campobello, A. Segreto, N. Donato, An Experimental Evaluation of CRT-based Forwarding Technique, in: 2019 IEEE International Symposium on Measurements Networking (M&N), 2019, pp. 1–6, <http://dx.doi.org/10.1109/IWMN.2019.8805035>.
- [10] G. Campobello, A. Leonardi, S. Palazzo, On the use of Chinese Remainder Theorem for energy saving in wireless sensor networks, in: 2008 IEEE Int. Conf. on Communications, IEEE, 2008, pp. 2723–2727.
- [11] F. Battaglia, G. Gugliandolo, G. Campobello, N. Donato, EEG-Over-BLE: A Low-Latency, Reliable, and Low-Power Architecture for Multichannel EEG Monitoring Systems, *IEEE Trans. Instrum. Meas.* 72 (2023) <http://dx.doi.org/10.1109/TIM.2023.3268471>.
- [12] F. Battaglia, et al., CRT-LoRa: An efficient and reliable MAC scheme for real-time industrial applications, in: IEEE International Workshop on Metrology for Industry 4.0 & IoT, 2023, pp. 30–35, <http://dx.doi.org/10.1109/MetroInd4.0IoT57462.2023.10180184>.
- [13] M. Amadeo, F. Battaglia, G. Campobello, Improving content retrieval in vehicular NDN using the Chinese remainder theorem, *IEEE Netw. Lett.* (2025).
- [14] S. Mastorakis, A. Afanasyev, L. Zhang, On the evolution of ndnSIM: An open-source simulator for NDN experimentation, *ACM SIGCOMM Comput. Commun. Rev.* 47 (3) (2017) 19–33.
- [15] X. Lin, The Bridge Toward 6G: 5G-Advanced Evolution in 3GPP Release 19, *IEEE Commun. Stand. Mag.* 9 (1) (2025) 28–35.
- [16] S. Fayyaz, M.A.U. Rehman, M.S.U. Din, M.I. Biswas, A.K. Bashir, B.-S. Kim, Information-centric mobile networks: A survey, discussion, and future research directions, *IEEE Access* 11 (2023) 40328–40372.
- [17] C. Yi, A. Afanasyev, I. Moiseenko, L. Wang, B. Zhang, L. Zhang, A case for stateful forwarding plane, *Comput. Commun.* 36 (7) (2013) 779–791.
- [18] V. Demiroglou, C.-A. Sarros, V. Tsaoussidis, NoD: A content retrieval scheme for intermittently-connected IoT networks, *Ad Hoc Networks* 130 (2022) 102825.
- [19] D. Hernandez, L. Gameiro, C. Senna, M. Luís, S. Sargento, Handling producer and consumer mobility in IoT publish–subscribe named data networks, *IEEE Internet Things J.* 9 (2) (2021) 868–884.
- [20] R.A. Rehman, B.-S. Kim, LOMCF: Forwarding and caching in named data networking based MANETs, *IEEE Trans. Veh. Technol.* 66 (10) (2017) 9350–9364.
- [21] A. Khalid, R.A. Rehman, M. Burhan, CBILEM: A novel energy aware mobility handling protocol for SDN based NDN-MANETs, *Ad Hoc Networks* 140 (2023) 103049.
- [22] M. Amadeo, C. Campolo, A. Molinaro, J. Harri, C.E. Rothenberg, A. Vinel, Enhancing the 3GPP V2X architecture with Information-Centric Networking, *Futur. Internet* 11 (9) (2019) 199.
- [23] C. Xu, M. Wang, X. Chen, L. Zhong, L.A. Grieco, Optimal information centric caching in 5G device-to-device communications, *IEEE Trans. Mob. Comput.* 17 (9) (2018) 2114–2126.
- [24] G. Chandrasekaran, N. Wang, R. Tafazolli, Caching on the move: Towards D2D-based information centric networking for mobile content distribution, in: IEEE 40th Conference on Local Computer Networks, LCN, 2015, pp. 312–320.
- [25] C. Ghali, A. Narayanan, D. Oran, G. Tsudik, C.A. Wood, Secure fragmentation for content-centric networks, in: IEEE 14th International Symposium on Network Computing and Applications, 2015, pp. 47–56.
- [26] C. Cao, K. Xue, H. Yue, J. Xu, FFRD: Fragment forwarding and reassembly decoupling based chunk transmission in NDN, in: IEEE WCSP, 2017, pp. 1–6.
- [27] W. Diao, J. An, T. Li, C. Zhu, Y. Zhang, X. Wang, Z. Liu, Low delay fragment forwarding in LEO satellite networks based on named data networking, *Comput. Commun.* 211 (2023) 216–228.
- [28] Y. Zeng, P.H. Pathak, P. Mohapatra, A first look at 802.11 ac in action: Energy efficiency and interference characterization, in: 2014 IFIP Networking Conference, 2014, pp. 1–9.
- [29] D.E. Knuth, The Art of Computer Programming, Volume 2: Seminumerical Algorithms, third ed., Addison-Wesley, Boston, MA, 1997.
- [30] Y. Tao, W.L. Tan, A reinforcement learning approach to wi-fi rate adaptation using the reinforce algorithm, in: 2024 IEEE Wireless Communications and Networking Conference, WCNC, IEEE, 2024, pp. 1–6.
- [31] M. Amadeo, C. Campolo, A. Molinaro, Multi-source data retrieval in IoT via Named Data Networking, in: Proceedings of the 1st ACM Conference on Information-Centric Networking, 2014, pp. 67–76.
- [32] Z. Zhu, S. Wang, X. Yang, V. Jacobson, L. Zhang, ACT: audio conference tool over named data networking, in: Proceedings of the ACM SIGCOMM Workshop on Information-Centric Networking, 2011, pp. 68–73.
- [33] H. Hui, J. Grant, K. McLaughlin, D. Laverty, S. Sezer, Secure real-time industrial IoT communications in smart grids using named data networking, in: IEEE 21st International Conference on Industrial Informatics, INDIN, 2023, pp. 1–6.
- [34] S. Dyerowicz, O. Aponte, P. Mendes, NDN Operation in Opportunistic Wireless Networks, *NDNcomm*, Memphis, USA, 2017.
- [35] Named Data Networking Project, NDN Packet Format Specification – Data Packet, 2020, (Accessed 10 April 2025), <https://docs.named-data.net/NDN-packet-spec/0.2.1/data.html>.
- [36] R. Shrivastava, S. Hegde, O. Blume, Sidelink evolution toward 5G-a/6G future considerations for standardization of group communications, *IEEE Commun. Stand. Mag.* 7 (1) (2023) 24–30.
- [37] M.S. Afaqui, E. Garcia-Villegas, E. Lopez-Aguilera, IEEE 802.11 ax: Challenges and requirements for future high efficiency WiFi, *IEEE Wirel. Commun.* 24 (3) (2016) 130–137.
- [38] D. Xu, A. Zhou, X. Zhang, G. Wang, X. Liu, C. An, Y. Shi, L. Liu, H. Ma, Understanding operational 5G: A first measurement study on its coverage, performance and energy consumption, in: Proceedings of the Annual Conference of the ACM Special Interest Group on Data Communication on the Applications, Technologies, Architectures, and Protocols for Computer Communication, 2020, pp. 479–494.
- [39] I. Rhee, M. Shin, S. Hong, K. Lee, S.-J. Kim, S. Chong, On the levy-walk nature of human mobility, *IEEE/ACM Trans. Netw.* 19 (3) (2011) 630–643.
- [40] D.X. Yang, Y. Guo, O. Aboul-Magd, 802.11 ax: The coming new WLAN system with more than 4x MAC throughput enhancement, in: IEEE 86th Vehicular Technology Conference (VTC-Fall), 2017, pp. 1–5.
- [41] F. Lemerrier, A.-C. Orgerie, Towards an energy-efficient wi-fi: An experimental study on recent standards power consumption, in: IEEE Wireless Communications and Networking Conference, WCNC, 2025, pp. 01–06.
- [42] M. Amadeo, C. Campolo, G. Ruggeri, A. Molinaro, Beyond edge caching: Freshness and popularity aware iot data caching via ndn at internet-scale, *IEEE Trans. Green Commun. Netw.* 6 (1) (2021) 352–364.
- [43] Y. Atif, S. Kharrazi, D. Jianguo, S.F. Andler, Internet of things data analytics for parking availability prediction and guidance, *Trans. Emerg. Telecommun. Technol.* 31 (5) (2020) e3862.
- [44] Y. Meng, Z. Zhang, Y. Huang, Cache-and energy harvesting-enabled D2D cellular network: Modeling, analysis and optimization, *IEEE Trans. Green Commun. Netw.* 5 (2) (2021) 703–713.
- [45] K. Schneider, C. Yi, B. Zhang, L. Zhang, A practical congestion control scheme for named data networking, in: Proceedings of the 3rd ACM Conference on Information-Centric Networking, 2016, pp. 21–30.
- [46] M. Amadeo, G. Ruggeri, C. Campolo, A. Molinaro, Diversity-improved caching of popular transient contents in vehicular named data networking, *Comput. Netw.* 184 (2021) 107625.