

Gain More for Less: The Surprising Benefits of QoS Management in Constrained NDN Networks

Cenk Gündoğan
HAW Hamburg
cenk.guendogan@haw-hamburg.de

Jakob Pfender
Victoria University of Wellington
jpfender@ecs.vuw.ac.nz

Michael Frey
Safety IO
michael.frey@safetyio.com

Thomas C. Schmidt
HAW Hamburg
t.schmidt@haw-hamburg.de

Felix Shzu-Juraschek
Safety IO
felix.juraschek@safetyio.com

Matthias Wählisch
Freie Universität Berlin
m.waehlich@fu-berlin.de

ABSTRACT

Quality of Service (QoS) in the IP world mainly manages forwarding resources, i.e., link capacities and buffer spaces. In addition, Information Centric Networking (ICN) offers resource dimensions such as in-network caches and forwarding state. In constrained wireless networks, these resources are scarce with a potentially high impact due to lossy radio transmission. In this paper, we explore the two basic service qualities (i) *prompt* and (ii) *reliable* traffic forwarding for the case of NDN. The resources we take into account are forwarding and queuing priorities, as well as the utilization of caches and of forwarding state space. We treat QoS resources not only in isolation, but correlate their use on local nodes and between network members. Network-wide coordination is based on simple, predefined QoS code points. Our findings indicate that coordinated QoS management in ICN is more than the sum of its parts and exceeds the impact QoS can have in the IP world.

CCS CONCEPTS

• **Networks** → **Network protocol design**; **Network experimentation**; **Network performance analysis**;

KEYWORDS

Internet of Things; ICN; Quality of Service; measurement

ACM Reference Format:

Cenk Gündoğan, Jakob Pfender, Michael Frey, Thomas C. Schmidt, Felix Shzu-Juraschek, and Matthias Wählisch. 2019. Gain More for Less: The Surprising Benefits of, QoS Management in Constrained NDN Networks. In *ICN '19: Conference on Information-Centric Networking, September 24–26, 2019, Macao, China*. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3357150.3357404>

1 INTRODUCTION

Quality of Service (QoS) in IP networks has been around for two decades, but so far has experienced remarkably little deployment. Its hesitant adoption is commonly understood to have two reasons:

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ICN '19, September 24–26, 2019, Macao, China

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6970-1/19/09...\$15.00

<https://doi.org/10.1145/3357150.3357404>

limited scalability (IntServ [14, 15]) and plain resource trading (Diff-Serv [12, 40])—the latter is often referred to as managed unfairness. While QoS in the IP world is mainly restricted to managing forwarding resources (link capacities and buffer spaces) [55] and the IEEE follows this approach on Layer 2 with Time Sensitive Networking (TSN) [26, 37], Information-Centric Networking (ICN) [3, 56] offers additional resource dimensions such as in-network caches and forwarding states that can shape network performance significantly.

Constrained wireless networks used to be the playground for experimental research and tinkering, but have recently turned into a fast-growing market. The advent of the Internet of Things (IoT) introduced the vision of omnipresent and always connected sensors and actuators that generate business models from new products, innovative processes, and data—a total of 1.6 Zettabytes is soon expected [36] in this rapidly growing business segment. LoRa [34] and NB-IoT [1] have already started deployment for a wide-area outreach to the embedded edge. As of today, 5G technologies have appeared on the horizon with the promise of tailored technologies that can deliver dedicated service qualities to their users in vertically sliced networks.

Key use cases for assured QoS are raised by industry. The Industrial Internet-of-Things (IIoT) is on track to interconnect time- and safety-critical infrastructure in industrial environments consisting of embedded microcontrollers with nearby edge intelligent and remote cloud services. Its use cases continually require high reliability and very low latency. Traditionally, these sensors and actuators are wired via field buses, which have been the state-of-the-art for decades, but often lack interconnectivity, interoperability, and security. Communication infrastructure components in industrial environments are also often physically separated from one another.

5G facilitates the creation of networks that tackle these challenges and allows companies to create their own private 5G-based networks on site. Companies can provide their own 5G infrastructure by making use of a key 5G concept called network slicing. Network slicing enables the creation of sub-networks for specific services and users that can preselect 5G network parameters such as end-to-end latency, maximum throughput, and traffic density. It allows companies to deploy ultra-reliable low-latency networks for critical infrastructure.

ICN has been a promising candidate for networking the IoT edge for a while [5, 7, 11, 45, 48, 49]. Recent experimental studies [17, 21] confirmed that ICN can sustain high reliability at moderate latency penalty even on lossy wireless links, and showed how it could be integrated into the 5G edge architecture. Nevertheless,




Figure 1: Manageable resources in IP vs. NDN.

current analyses and solutions are built purely on equal resource sharing. Studies of ICN behavior under prioritization and active QoS management are missing.

In this paper, we perform a first exploration of QoS management impacts on NDN [28, 58] in the resource-constrained IoT using RIOT [9]. We start from the two basic service attributes (i) *prompt* and (ii) *reliable* traffic forwarding and define a simple yet efficient management scheme. We carefully implement these QoS semantics by employing not only the NDN resources *forwarding capacity*, *PIT state*, and *cache* alone, but correlating resources internally on a node and also externally between nodes. These correlations can be performed without additional signaling overhead.

Our findings from extensive experiments in a large testbed confirm the efficacy of our approach. To our surprise, however, the QoS measures do not sacrifice the performance of unprioritized traffic. Even though resources shift to the prime packets, best effort flows still uphold their performance or even improve. A thorough analysis reveals the positive effects of resource correlation, which raises the overall network performance to a higher level than in the state of uncoordinated resource allocation.

The remainder of this paper is structured as follows. The problem space of managing distributed resources is discussed in Section 2 along with related work. Section 3 introduces the QoS building blocks and how we manage the resources. Our implementations and evaluations are presented in Section 4, followed by the conclusions and an outlook on future work in Section 5.

2 THE PROBLEM OF DISTRIBUTED RESOURCE MANAGEMENT IN ICN

2.1 Problem Statement

Implementing service differentiation and assurance in a network raises the challenge of managing distributed resources without sacrificing them. Common Internet approaches follow a flow-based (e.g., IntServ) or a class-based (e.g., DiffServ) concept. Flow-based resource reservation requires dedicated signaling and state, which quickly reaches scalability limits with resource exhaustion. In the presence of ubiquitous in-network caching and request aggregation, content endpoints are unspecified and data paths are inherently multi-source, multi-destination, and possibly widely disjoint. This makes ICN flows difficult to identify and to maintain.

Resource allocation according to packet classes requires ingress shaping and filtering, since unforeseen traffic bursts quickly exhaust the per-class reservations and counteract service assurances. In several ICN flavors including NDN and CCNx, link occupancy and




Figure 2: PIT decorrelation terminates data paths.

forwarding demands are steered hop-by-hop in a request-response fashion. Small requests trigger data replies of unknown size, provenance, and timing. This complicates reliable resource predictions for responses in NDN. Shaping and dropping Interests can prevent resource exhaustion, but may leave the network underutilized. Restricting ingress only to data may lead to bursts of unsatisfied Interests, which waste network resources.

In-network caches in ICN enrich the field of manageable resources. Caches reduce latency and forwarding load and often take the role of a (large, delay-tolerant) retransmission buffer. With NDN/CCNx, additional resources come into play in the form of Pending Interest Tables (PITs) that govern stateful forwarding. The overall resource ensemble is visualized in Figure 1 and raises concerns. Capacities in forwarding, caching, and pending Interest state may be largely heterogeneous. A wirespeed forwarder may supply negligible cache memory compared to its transmission capacity, for example. In the IoT the opposite is often true in that flash memory is normally shipped in ‘infinite’ sizes when compared to the main memory (PIT) or the wireless data rate. A beneficial resource management faces the problem of how to carefully balance these resources and arrive at an overall optimized network performance [30].

Resource complexity, however, extends beyond a single system. The impact of distributed resources is easily flawed if management cannot jointly coordinate contributions. Neighboring caches, for instance, are less effective if filled with identical data. A more delicate problem arises from PIT state management. If neighboring PITs diverge and no longer represent common forwarding paths (see Figure 2) all data flows terminate and forwarding resources are wasted. This problem of state decorrelation was first reported in [54].

2.2 QoS Qualifiers

QoS extensions for ICN have recently attracted attention and generated various efforts within the IRTF ICN research group [22, 29, 38, 41, 42]. A proposed flow classification mechanism [38] differentiates traffic flows based on content name prefixes using two different methods. The first method, *EC3*, introduces a new message header entry in Data packets that indicates the prefix length of the content name characteristic for the classification process. Once a consumer learns about such an equivalence class, it can also include equivalence class indicators into subsequent Interests. The second method, *ECNCT*, encodes classification indicators directly into names at content creation using a new type of name component. This solution does not inflate messages with additional headers, which is advantageous for constrained IoT deployments. Nevertheless, encoding flow classification indicators in typed name components leads to an inflation of names in the routing system. Identical content published with different classification values will lead to duplicate names that cannot be aggregated in Forwarding

Information Bases (FIBs). Analogously, the default matching functions used for the Pending Interest Table (PIT) and Content Store (CS) will consider names that differ only in their flow classification as dissimilar, which conflicts with Interest aggregation and cache utility.

In the Internet Research Task Force (IRTF), two approaches to handle QoS treatments are under discussion. An end-to-end QoS framework [29], in which non-routable QoS markers are appended as suffixes to content names. In analogy to DiffServ, these markers are then used to apply different resource allocation mechanisms to the request and response messages. The FIB is extended to ignore QoS markers, and the PIT is modified to disaggregate pending requests that have differing QoS markers. This disaggregation may lead to PIT inflation in particular for setups with a high diversity of QoS markers. The work at hand has also been presented to the IRTF [22].

Tsilopoulos *et al.* [53] identify three different types of traffic based on two characteristics in ICN traffic. The authors introduce two extensions to CCN in order to handle these traffic types, *Persistent Interests* and *Reliable Notifications*. Persistent Interests are valid for Data packets which are produced during a pre-defined period of time. Reliable Notifications inform receivers that real-time data is available and are propagated reliably on a hop-by-hop basis. Notifications that are not acknowledged in time are retransmitted. If a receiver successfully receives a notification but does not receive data in a given time, new Interests are created to renew the request.

2.3 Distributed Forwarding Resources

MIRCC [35] introduces a rate-based, multipath-aware congestion control scheme for ICN. Each Data message in MIRCC contains a rate value which in turn is used to calculate per-link rates. It is inspired by the Rate Control Protocol (RCP) [20] for IP networks.

Another approach [4] manages forwarding resources and is experimentally evaluated for the PURSUIT architecture. This work implements a QoS differentiation scheme, where each forwarder manages virtual links that include packet queues with varying traffic rates and a designated traffic shaper. QoS information is encoded into the names and determines the mapping of traffic flows to low or high priority virtual links.

2.4 Distributed Cache Management

Caching policies that employ heuristics to inform a caching decision instead of caching all incoming content can be broadly organized into a number of different families, depending on what information they use to reach their caching decision.

The easiest way to achieve higher cache diversity without increasing the complexity of the caching policy is to cache probabilistically. The static version of this approach, commonly known as *Prob(p)*, uses a static probability p that governs how likely a given node will cache a given content chunk. It has been shown [24, 60] that *Prob(p)* outperforms the default strategy of caching everything in terms of cache diversity, and that lower values for p correlate with higher diversity [6, 24, 25, 46, 51, 60].

Instead of using the same static caching probability for all incoming content, a caching strategy may also dynamically compute a probability for each individual node or even for each content

chunk in order to adapt the caching behavior to the state of the network. These strategies may be based purely on node-local information, such as CS saturation or battery levels; on properties of the incoming content, such as its name, freshness, type, or producer; or on information from the wider network, such as the position of the caching node in the network topology or the CS contents of neighboring nodes. Examples include *ProbCache* [46], which computes the caching probability of a given content chunk based on the total number of hops between its producer and the consumer that requested it, and *pCASTING* [24], which considers the freshness of the content as well as the node's battery level and CS saturation when calculating p . Both strategies have been found to increase the cache hit ratio, reduce the average number of hops required to hit requested content, and reduce the number of cache evictions [43]. Various other dynamic probabilistic caching strategies have been proposed, with decisions based on content freshness [19, 24], content popularity [18], or whether the content is already in a neighboring CS [59].

Not all caching strategies use the probabilistic approach. Instead, some exploit knowledge about the network topology [16, 27, 39, 50, 61], which has the advantage of taking global knowledge about the network into account but often comes at significant costs such as lengthy setup times, communications and memory overhead, and vulnerability to changes in the topology [44]. Another class of caching policies has nodes cooperate with their neighbors, either explicitly by exchanging information [33] or implicitly by using pre-defined rules [23, 32, 57].

3 QOS BUILDING BLOCKS FOR NDN

3.1 Traffic Flow Classification

General purpose networks simultaneously host competing traffic flows that exhibit varying resource requirements and time constraints. This also holds for typical IoT deployments, in which flows originate from sensors and actuators, or from remote cloud services that connect via gateways to the IoT domain. A flow classification is necessary for differentiating packets that belong to separate message flows, whenever the network should allocate distinct resources and treat them in a differentiated manner.

In the IP world, traffic flows are defined by the application endpoints and identified by address/port tuples (IPv4) or addresses plus flow label (IPv6). Since ICN abandons the host-centric paradigm, this definition of traffic flows no longer holds. In the presence of delocalized content and in-network caching, the concept of application endpoints becomes meaningless.

The characterizing property of these packets are content names (or prefixes). In the example of CCNx and NDN, content names appear in related Interest and Data packets. This ubiquity of names and their potential to impose hierarchies on content make them a distinguished component for identifying flows. Accordingly, we propose a traffic flow classification mechanism for NDN that builds on hierarchical names, prefixes, and longest common prefix match. It is explicitly designed to not put a strain on typically resource constrained IoT devices. In this regard, our scheme is computationally simple and does not require an additional overhead in message headers.




Figure 3: QoS Service Levels.

In this work, we consider service differentiation with respect to two quality dimensions: *latency* and *reliability*. For simplicity we only use a plain distinction in each quality, which results in a matrix breakdown of service levels as shown in Figure 3. More sophisticated differentiations apply analogously.

Service classes are assigned to flows according to a list of prefixes that are marked with a traffic class and maintained by each node. Incoming Interest and Data messages are then mapped onto traffic classes by applying a longest prefix match against this list of traffic identifiers. Traffic class names may look as follows:

Traffic Class	Priority
/HK/ACM/ICN	<Reliable, Regular>
/HK/ACM/ICN/site/A/alarm	<Reliable, Prompt>
/HK/ACM/ICN/site/B/temp	<Regular, Prompt>

Given this example, Interest and Data messages for the name prefix */HK/ACM/ICN/site/C* would map to the class of */HK/ACM/ICN*.

In this work, we assume such lists deployed at all nodes within a network domain. The distribution and maintenance of QoS configurations may be use case dependent and introduce a complex dynamic. Its generic treatment is beyond the scope of this work, which focuses on a lightweight flow-to-priority mapping for the purpose of quantifying the benefits of QoS resource management.

It may be noted that fine-grained service differentiation within a complex name hierarchy can result in large QoS prefix tables that consume significant memory. For the constrained IoT use case, though, we argue that sensor readings and actuator settings are machine type communications (MTC) with (short) names configurable according to processing needs. Hence QoS overheads should easily comply to resource constraints.

3.2 Manageable Resources

3.2.1 Link Layer. The link layer manages access to the media and provides space to buffer packets. In low-power wireless networks, media access times are highly susceptible to media saturation and buffer spaces are small. While time slotted technologies such as the IEEE 802.15.4e TSCH mode and Bluetooth Low Energy access media in a deterministically scheduled manner, the unslotted CSMA/CA version of IEEE 802.15.4 or long-range radios such as LoRa are more susceptible to packet collisions between neighboring nodes.

In the IoT, content producers that generate sensor readings may produce egress traffic at a rate that is much higher than the average media access time. In addition, nodes may further need to forward ingress traffic in multi-hop scenarios. For this purpose, buffering egress traffic is necessary to cope with traffic spikes.

Queuing and buffering take the same role in ICN as in the IP world. Class-based forwarding queues will process packets of the *prompt* flow class before packets with a regular priority while buffer space will prevent packet drops. It should be noted, though, that rapidly forwarded packets in NDN will also quickly satisfy PIT entries and thereby free forwarding resources for unprioritized traffic on the same path.

3.2.2 Pending Interest Table. The Pending Interest Table (PIT) enables the stateful forwarding plane of CCNx and NDN and thereby governs the flows in the network. The size of the PIT resource effectively dictates the maximum number of simultaneous open requests and the coherence of PIT entries along a path determines whether flows can propagate without barriers. In normal NDN operation, PIT state is allocated when an Interest message is processed, and it is removed in two scenarios. Either a returning Data message consumes the PIT state, or a timeout after a succession of retransmissions clears the state.

PIT saturation is common even in overprovisioned networks, but is far more likely to occur in IoT deployments. Limited RAM resources, slow processing power, delayed media access, and packet loss in low-power networks with intermittent connectivity can all cause the PIT to reach its maximum capacity.

The typical way of handling incoming Interest messages at a saturated PIT is to drop them in order to avoid cancelling active but incomplete request operations. The penalty for dropping such Interests is an increase in latency due to retransmissions, which usually happen on the scale of seconds. To avoid high latencies for time-sensitive traffic flows, an elaborate PIT eviction strategy is necessary, which accounts for (i) unhindered forwarding of prioritized Data, and harmonizes with (ii) the retransmission mechanisms of the regular NDN.

3.2.3 Content Store. Due to memory constraints, the Content Store (CS) is typically small, necessitating heuristics for deciding what content to cache at which node instead of indiscriminately caching all incoming content. There is a wealth of existing research on how to make the most efficient use of limited CS space, with a number of different strategies employing various heuristics to decide whether or not to cache incoming content (see Section 2.4). This aspect of caching is called the *caching decision strategy*.

The introduction of traffic flow priorities adds an additional dimension to the caching decision. Regardless of which specific caching decision strategy is employed, content marked as *reliable* should always be cached as it is imperative that this content is available throughout the network. Thus, reception of *reliable* content should not trigger the caching decision strategy; instead, control should be handed directly to the cache replacement strategy (see below). The question whether *prompt* content should be cached with a higher priority than content with *regular* latency requirements is not as clear-cut. Caching *prompt* content with higher priority would have a positive effect on future transmissions of that content object (either by retransmission of the original request or by new requests) and thus have a positive effect on latency, although the potential gain in this aspect is dependent on path length. Any content that is marked as *regular* in both QoS dimensions should be treated as normal; in other words, the caching decision strategy is consulted.

After a node has decided to cache a new content object, an additional step may have to be taken in case the CS is at capacity. This aspect of caching is the *cache replacement strategy*. In most cases, CS contents will be replaced using a simple heuristic such as Least Recently Used (LRU). However, once again the introduction of traffic flow priorities adds an additional dimension to this decision.

In general, incoming content should not replace content of a higher priority. Therefore, content with *regular* latency requirements should not replace *prompt* content and content with *regular* reliability should not replace *reliable* content. When it comes to the correlation between latency and reliability, the primary goal of the CS should be to ensure content availability, which places a stronger emphasis on the reliability aspect. Thus, *reliable* content with *regular* latency should be able to replace *prompt* content if no other content is eligible to be replaced. If all content is of the same priority class, regular replacement rules (e.g., LRU) should apply.

In probabilistic caching, as introduced in Section 2.4, each node caches incoming content according to a certain probability p . Regardless of how exactly p is determined (whether statically or by one of the dynamic methods discussed in Section 2.4), the probabilistic approach may be refined by differentiating between two separate probabilities p_{reg} for regular content and p_{rel} for reliable content, with $p_{rel} > p_{reg}$. This has the effect that a CS at each node will have different contents, thus contributing to CS diversity across the network by making a larger range of content available as cached copies, while giving consideration to service classes ensures that higher-priority content is still treated preferentially.

3.3 Distributed QoS Management

We are now ready to present our approaches to distributed resource management for supporting QoS in ICN. The corresponding mechanisms fall into three classes, depending on the level of interdependence between resources on the same device or between devices:

3.3.1 Locally Isolated Decisions. The straightforward allocation of independent resources to packet forwarding follows three simple rules:

Prioritized forwarding applies to flows marked as *prompt*.

Cache (re-)placement decisions obey the priority order *reliable* (highest) to *regular* (lowest). In the presence of probabilistic caching strategies, the weights are set accordingly.

PIT management operates in favor of rapid packet forwarding, so PIs enter the PIT in the order *prompt* (highest) to *regular* (lowest). Newly arriving Interests that meet a PIT saturated with entries of equal or higher priority will be dropped.

3.3.2 Local Resource Correlations. These are decisions that entail interaction between mechanisms on the same device (intra-device correlations). This includes the correlation between the caching decision and cache replacement strategies, where e.g. the caching decision may pre-empt the cache replacement decision in the case of *reliable* flows, while the cache replacement decision may drop content even in the case of a positive caching decision if no content of the same or lower service class can be replaced. In detail we take the following steps:

w/ PIT entry If arriving Data meets a valid PIT entry, Data is forwarded according to priorities *and cached* with priority, if marked as *reliable*. In the case of exhausted prioritized forwarding queue, *prompt* traffic will be cached with the highest priority.

w/o PIT entry If arriving Data meets no valid PIT entry, cache placement will still be initiated for *prompt* and *reliable* data in subsequent order. For probabilistic caching, weights are adjusted accordingly.

In balanced, unconstrained NDN networks, returning regular Data meets open PIT states. For saturated PITs, however, PIT entries may time out quickly, or resource management may enforce eviction of PIT entries in favor of other requests. Allowing Data without corresponding PIT entries to be cached may introduce the threat of cache poisoning attacks. However, a simple rate limiting on incoming Data packets and a reduced cache time for these CS entries may reduce the attack surface in our constrained environment. Further analysis of related effects is left to future work.

3.3.3 Distributed Resource Coordination. Such mechanisms affect resources across multiple or all devices in the network (inter-device correlations). These include maintaining PIT coherence by ensuring that all nodes apply uniform QoS mechanisms when replacing content of different service classes, as well as achieving CS diversity by introducing probabilistic caching based on priority classes. In our system, distributed coordination is achieved as follows.

PIT coherence is increased by applying the same PIT eviction strategy at all nodes, i.e., evict *regular* before *reliable* before *prompt*.

Cache efficiency increases with probabilistic caching using the coordination of equal cache weights. It is noteworthy that probabilistic caching reduces the risk of starvation for low priority content due to higher cache diversity, even if high priority flows dominate the network.

4 EXPERIMENTAL EVALUATIONS

4.1 Implementation and Experiment Setup

4.1.1 Software and Hardware Platforms. We conduct our experiments on the FIT IoT-Lab testbed using typical class 2 [13] IoT devices that feature an ARM Cortex-M3 MCU with 64 kB RAM and 512 kB ROM. Each device further contains an Atmel AT86RF231 [8] 2.4 GHz transceiver to operate on the IEEE 802.15.4 radio.

All devices run on RIOT OS [10] version 2019.04 with the integrated NDN network stack CCN-lite [52], which we extended with our QoS management scheme. In addition to the PIT and CS management strategies, a very lightweight prioritized forwarding was implemented using a single packet double-buffer that allows for pairwise packet re-ordering. We also note that network stack performance usually exceeds link speed. While IEEE 802.15.4 provides a theoretical maximum of 250 kbit/s, I/O and data processing in the network stack is at least one order of magnitude faster [31].

Following the large-scale deployment in [21], we configure a maximum number of four retransmissions and a retransmission interval of two seconds for CCN-lite. To analyze our approach under different levels of network saturation, we configure the maximum capacities of PIT and CS to range between 5 and 30 elements.

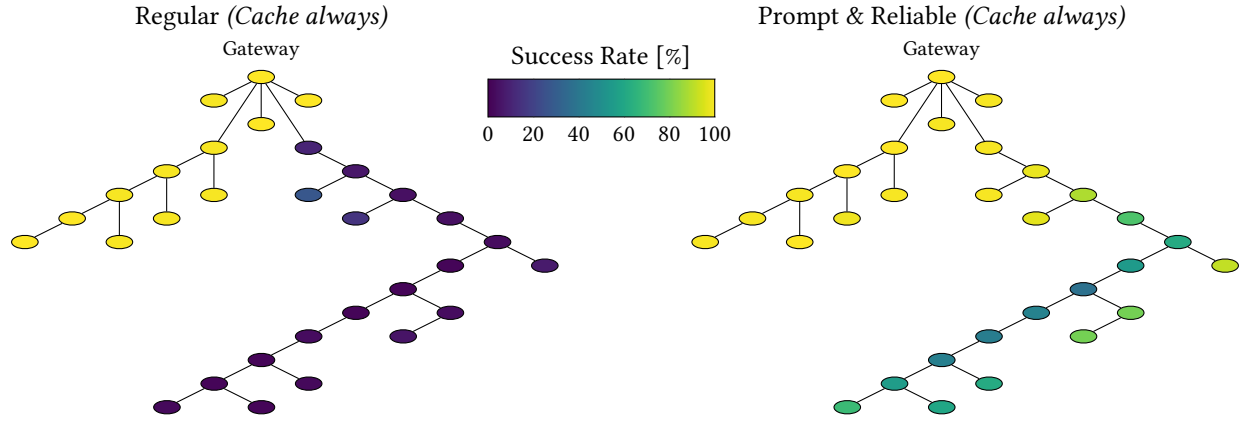


Figure 4: Nodal success rates for *Scenario 1* using regular traffic (left) and *reliable actuator traffic* (right).

Notably, the estimated RAM usage for 30 PIT entries and 30 CS elements with name lengths of ≈ 32 bytes is already approximately 11 KiB, which is around 17% of the total available RAM for our hardware platform.

4.1.2 Topology Setup. The testbed provides access to multiple sites with varying numbers of M3 IoT devices and network characteristics. We deploy our applications on the *Grenoble* site, as this supports significantly complex multi-hop paths. We choose 31 M3 devices arbitrarily, where one device acts as a gateway node and the other 30 devices act as sensors and actuators. Since convergecast is the most predominant traffic pattern in common IoT scenarios, we arrange our devices to form a *Destination Oriented Directed Acyclic Graph (DODAG)* that is rooted at the gateway node. Approximately 60% of the nodes are reachable from the root within 4–5 hops, while the remaining devices have path lengths up to 12 hops. The topology is visualized in Figure 4. Extended left and right wings in the routing topology result from long hallways in the *Grenoble* site.

4.1.3 Scenarios. We want to quantify the efficiency of our QoS management scheme in saturated multi-hop deployments that display typical traffic patterns. With this in view, we analyze our approach under two different scenarios.

Scenario 1: Mixed Sensors and Actuators. The gateway node requests temperature readings from the 30 sensor nodes every 10 s with ± 2 s jitter interval. Thus, on average, the request rate at the gateway approximates to 3 *packets/s* and including the reception rate of responses, the gateway handles 6 *packets/s*. The naming scheme for each request consists of a prefix, a device-specific *node id*, and an increasing sequence number. We refer to this traffic equally as *sensor readings* or *gateway initiated traffic*.

In addition, all 30 devices further act as actuators that periodically request a device-specific state from the gateway node every 5 s with ± 1 s jitter. This yields a request reception rate of 6 *packets/s* on the root node. The naming scheme for these requests similarly consists of a prefix, a device-specific *node id*, and an increasing sequence number. We name this traffic *actuator initiated traffic*.

Scenario 2: Sensing and Lighting Control. In the second scenario, we change the role of actuators but leave the sensor readings unchanged. Instead of independent actuators that receive disjoint instructions, we envision a scenario of lighting in which groups of fixtures switch lights in a coordinated way. Hence, these groups receive identical commands and caching becomes applicable.

To explore the event space by mixing group memberships, our experiments proceed as follows. For each request, an actuator randomly joins one out of five possible groups. The naming scheme for such requests is changed to include the selected *group id*, instead of a device-specific *node id*. Besides naming, we use the same request parameters for the *actuator initiated traffic* as in the first scenario. We repeat this process 240 times for each configuration in order to explore the state space of unevenly distributed groups and converge statistics.

4.1.4 Caching Parameters. We consider the caching procedure to consist of two fundamental steps: (i) caching decision, and (ii) cache replacement. In our experiments, we use two different caching decision strategies and one cache replacement strategy.

The first decision strategy is to *cache always* incoming Data packets, with the restriction for regular traffic that Data packets without a corresponding PIT entry are dropped and not cached. For *reliable* traffic, we adjust this strategy, such that Data packets without PIT entries are still considered for caching. The second decision strategy is to *cache probabilistically*. Every node caches incoming Data packets with a probability p_{reg} of 30% for regular traffic, and with a probability p_{rel} of 70% for *reliable* traffic. For a saturated CS, our cache replacement strategy evicts content store elements using the *least recently used (LRU)* policy.

4.2 Results

4.2.1 Analyses and metrics. We analyze the network utilization for gateway and actuator initiated traffic separately in the presented scenarios with and without the proposed QoS features. Furthermore, we measure the *success rates*, *goodputs* and *time to completion* for each node in the topology.

4.2.2 Scenario 1: Mixed Sensors and Actuators. In this scenario, sensor and actuator data are device-specific and thus only single

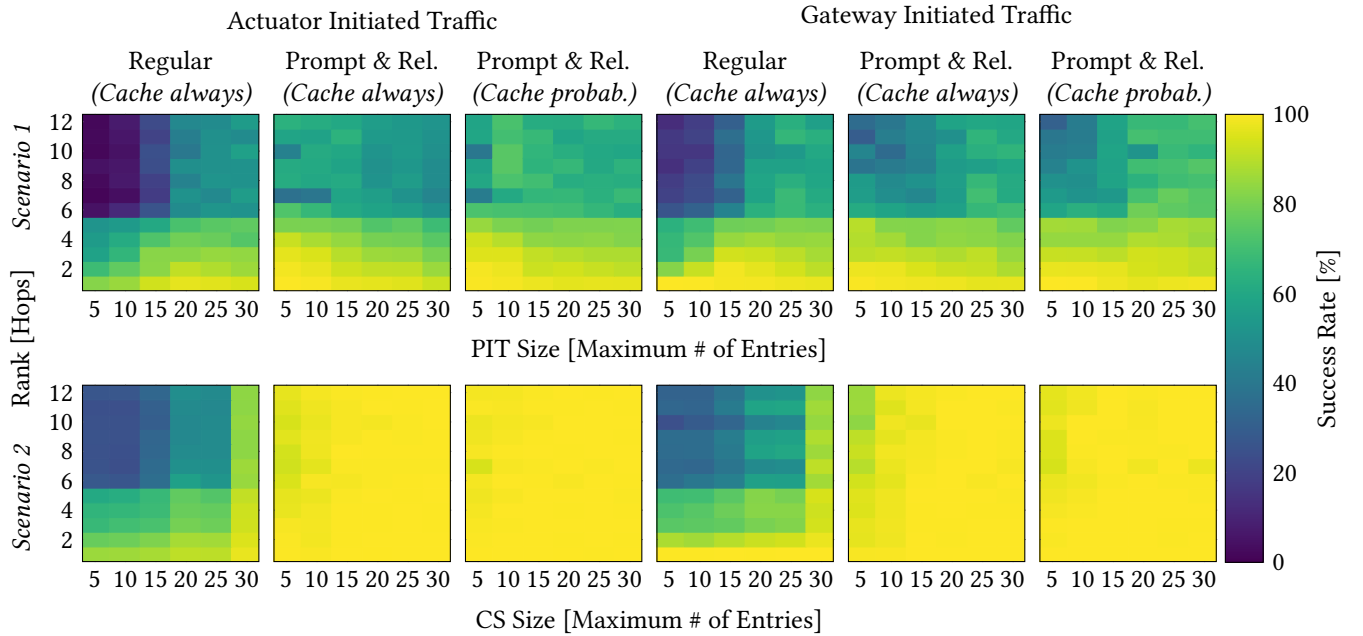


Figure 5: Success rates per rank for *Scenario 1* and *Scenario 2* using varying PIT and CS sizes.

destinations benefit from on-path caching during the narrow window of Interest retransmissions.

Success rates. In our first experiment we focus on the nodal success rates using the first scenario with a PIT and CS limitation of 5 entries. The gateway is configured with a PIT limitation of 50. Figure 4 shows the resulting success rates for (i) the regular operation of NDN on the left hand side, and (ii) a setup with prioritized actuator initiated traffic using the *prompt* and *reliable* QoS service levels on the right hand side. The success rates per node are color coded and range from 0% (purple) to 100% (yellow).

Figure 4 clearly depicts huge differences in success rates for both configurations. In the normal NDN operation, nodes close to the gateway, as well as the left wing of the topology, achieve 100% success rates. Strikingly, the right wing exhibits major network stress, with nearly all actuators having a success rate below 10% due to PIT overflows. Conversely, with QoS service levels enabled, the right wing shows much enhanced network performance, which results in overall higher success. With this configuration, 70–100% of the packets arrive at actuators close to the gateway, and 40–70% at more distant nodes. Leaf nodes farther away show a greater improvement in success rates than forwarding ancestor nodes.

This striking example nicely illustrates the positive effect of PIT correlation under QoS. While in the regular case Interests at nodes with exhausted PIT are discarded as they randomly arrive, the QoS marking preselects those requests that are prioritized throughout the network, leading to fewer retransmissions and a more efficient use of the overall PIT space available in the network.

We now dig deeper into network reliability and examine the success rates for a range of PIT sizes plotted against the node ranks (first row of Figure 5). As previously observed, the success rate

for *regular* actuator traffic collapses at higher ranks. The gateway traffic performs similarly poor. Increasing the maximum PIT size gradually improves the nodal success rates. In addition to the *regular* traffic, we further show actuator traffic with *prompt & reliable* QoS service levels. Overall, the setups with prioritized traffic show great improvements for the smaller PIT sizes, while the *probabilistic* caching decision strategy performs slightly better than the *always* strategy. A surprising effect is observed with a PIT size of 30: the success rates for all configurations decline slightly for lower ranks. This is caused by an increased retransmission overhead per node, resulting in link saturation.

Network utilization. In typical convergecast settings, a majority of the traffic traverses the gateway to reach remote endpoints. Thus, a robust operation of the gateway node is crucial to ensure adequate performance of the entire IoT network. Due to the increasing sequence numbers used in the naming scheme for *Scenario 1*, virtually no response contributes to future in-network cache hits. To gauge the network load on the gateway, we analyze the number of outgoing requests and incoming responses during a setup in which actuator initiated traffic is added to the gateway initiated traffic after approximately eight minutes. We first perform this experiment without any QoS features enabled, and then repeat it with the adjustment that actuator initiated traffic is prioritized using the *reliable* and *prompt* service levels. The PIT of the gateway is configured to a maximum of 50 entries, while the remaining nodes have a PIT maximum of 5 entries. The CS is limited to 5 entries for all nodes.

We observe in Figure 6 that the gateway initiated traffic exhibits a steady request-response flow of about 180 *packets/min* for both

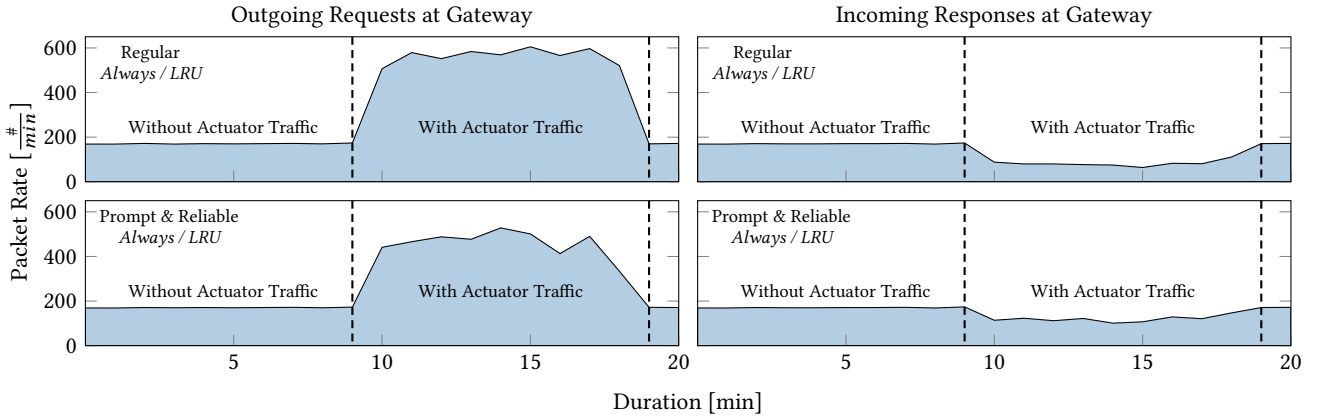


Figure 6: Packet transmission rate per minute for requests and responses measured at the gateway.

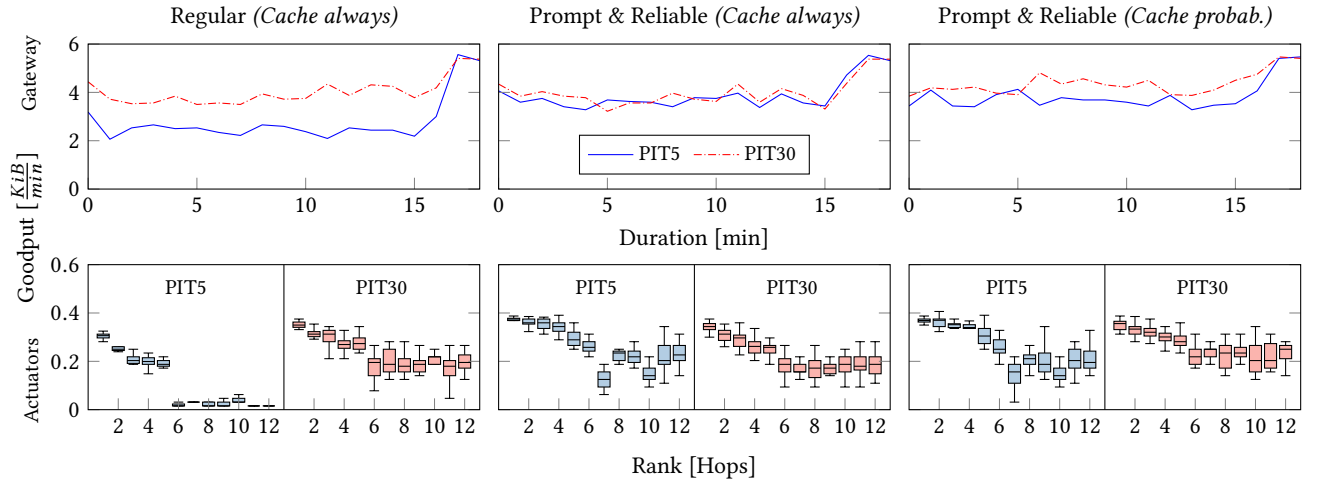


Figure 7: Goodput evolution for Scenario 1 with actuator and gateway traffic using a CS size of 5.

requests and responses. As soon as the actuators initiate their periodic requests (at minute eight), the network load at the gateway increases. The number of requests spikes threefold due to network layer retransmissions, while the number of returning responses drops to half. In contrast, the setup with prioritized actuator traffic clearly admits a reduced number of requests at the gateway while achieving a higher response rate.

These results reveal that prioritizing the actuator initiated traffic has a positive effect on the overall network load due to reduced retransmissions.

Goodputs. Figure 7 shows that while PIT sizes have a significant impact on the goodput at the gateway and at each actuator during normal operation, this effect is reduced when QoS service classes are introduced. When using service classes and the *always* caching decision strategy, network members with small PIT sizes can reach a level of goodput that is comparable to those of the largest PIT size during normal operation. With *probabilistic* caching in operation, the network performance increases even further for larger

PITs, which support a higher number of concurrent flows that can leverage cache diversity better. This overall enhanced transport performance is caused by flows that complete with delay based on segment-wise available network capacities and retransmissions. The corresponding temporal effects can be observed from Figure 8.

At actuator nodes far from the gateway, our QoS mechanisms mitigate the effects of PIT exhaustion, which in normal operation leads to an abrupt collapse of the throughput at around a rank of 6. Our QoS mechanisms cause a smooth, gradual decline in performance instead.

Time to completion. We can see from Figure 8 that content arrival times are significantly reduced for smaller PIT sizes when QoS service classes are introduced—for traffic at gateways and even more at the prioritized actuators. Simultaneously, we see again a signature of enhanced traffic delivery for QoS-coordinated flows that shows doubled success rates from 40% to 80%. It is worthwhile to (re-)observe that the sensor traffic also experiences improvements

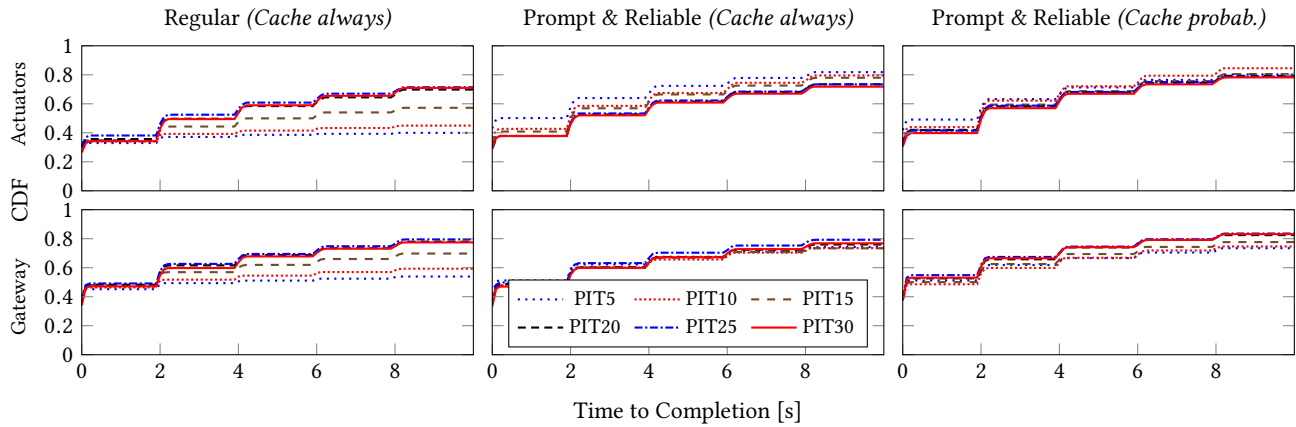


Figure 8: Time to completion for *Scenario 1* with actuator and gateway traffic using a CS size of 5.

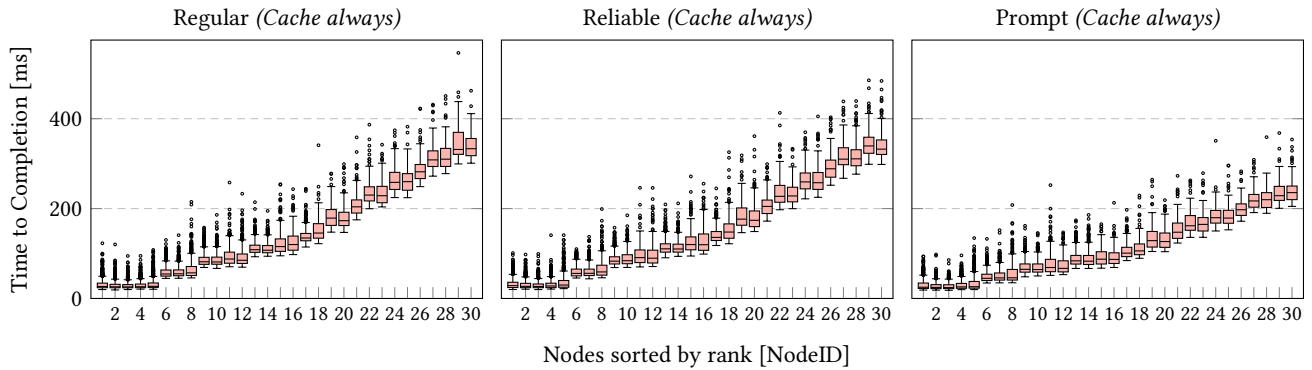


Figure 9: Time to completion per actuator and quality dimension in *Scenario 1* using PIT and CS sizes of 5.

due to a more efficient balancing of resources—small PIT sizes in particular.

We now quantify the content arrival times separately for each quality dimension. Figure 9 illustrates the nodal content arrival times with PIT and CS sizes both set to 5. In all three displayed configurations, completion time increases with the distance to the gateway. While arrival times range from 5 ms to 350 ms for the majority of requests with the *regular* and *reliable* configurations, the *prompt* configuration yields noticeably faster delivery times for all nodes. Especially nodes far away from the gateway experience a reduction of about 100 ms ($\approx 30\%$).

4.2.3 Scenario 2: Sensing and Lighting Control. In this scenario, prioritized actuator traffic flows from the gateway to multiple destinations and thus benefits from on-path caching. Accordingly, we expect the network performance to improve over that for Scenario 1.

Success rates. The overall success values presented in the second row of Figure 5 confirm these expectations. Success in content delivery nicely approaches 100% in most QoS settings. Even for the constrained number of 5 PIT entries, increasing the CS sizes from 5 to 30 suffices to turn traffic from QoS classes into a reliable service. While we do see some failures at higher ranks for the small

CS size of 5 similar to *Scenario 1*, increasing the CS capacity to 10 is already sufficient to attain success rates above 80%. The most striking contrast is found in comparison to the results of regular NDN operation, where even with a maximum CS size of 30 the failure rate stays at 30–40% at higher ranks. Regular NDN traffic apparently profits less from cacheable content. This is mainly due to PIT decorrelation, which breaks content flows.

Time to completion. QoS service classes have a significant impact on the content arrival times for both the actuators and the gateway, as was already observed for *Scenario 1*. Figure 10 reflects the same qualitative picture for *Scenario 2*, but at significantly reduced probabilities of retransmission. The latter is due to actuator traffic that is coordinated in QoS classes and coherently serviced from caches—a large reduction in overall network load. Results slightly improve for enhanced cache diversity in probabilistic caching, with 90% of the packets arriving promptly (< 100 ms) at cache sizes of at least 10 packets. Similar to *Scenario 1*, CS sizes become less relevant in the presence of QoS marking, since prioritized traffic arrives quickly at its destination and remains unaffected by regular cache replacement.

While we observe that increased CS sizes contribute to reduced completion times and improved success rates, we also notice that

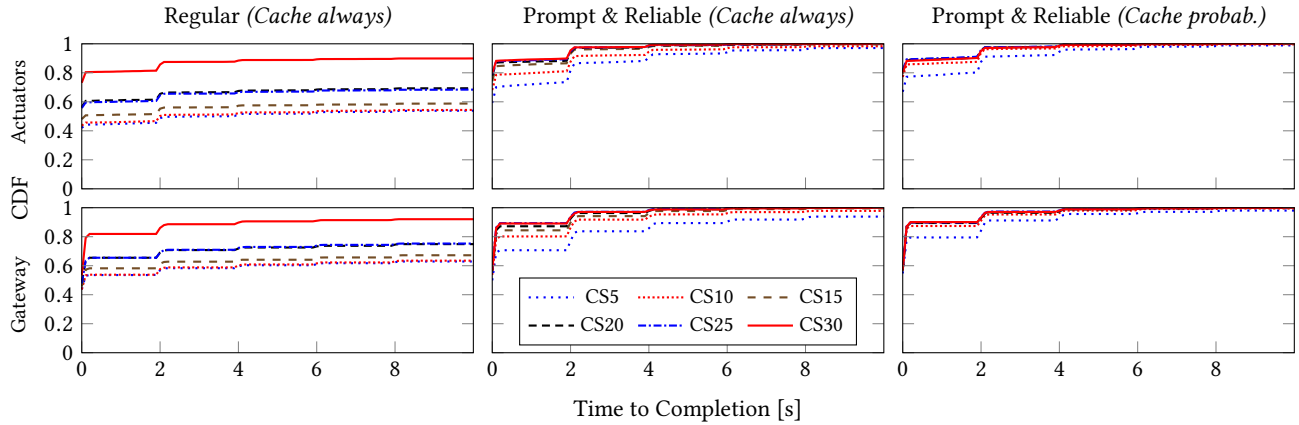


Figure 10: Time to completion for Scenario 2 with actuator and gateway traffic using a PIT size of 5.




Figure 11: Cache hit for Scenario 2 and a PIT size of 5.

even a CS size of 30 does not suffice for the *regular* configuration. On the other hand, the configurations that use QoS service levels display close to 100% success rate, even with severely limited CS sizes, and about 70% of all requests for each traffic type complete in less than 100 ms.

Cache hits. Analysing the cache efficiencies supports these observations. Figure 11 displays the relative in-network cache hits for actuator traffic in setups with varying CS sizes. While the regular NDN operation yields a marginally improving cache hit ratio for increasing CS sizes, both QoS enabled setups exhibit a noticeable enhancement. This improved cache efficacy is caused by the privileged cache resource utilization for data of the *reliable* actuator traffic, whereas data of the gateway traffic is more likely to be evicted. Another expected observation is that *probabilistic* caching further improves the cache hit ratio thanks to its increased CS diversity.

5 CONCLUSIONS AND OUTLOOK

We presented and analyzed QoS extensions to NDN that are suitable for constrained devices. Starting from a name-oriented flow classification scheme, we introduced the two service dimensions *prompt* and *reliable* network forwarding. Strategies were defined that not only foster local, isolated resource allocations, but take into account coordinative effects between different internal resources of a node, as well as correlations between nodes. Here we exploited the rich set of forwarding and caching options that NDN includes.

We were able to validate our approach in real-world experiments on a large testbed using a realistic multi-hop wireless setup. Moreover, we learned that QoS management in NDN is not confined to simple resource trading, but can lead to a global enhancement of network performance by optimizing the interplay between various resource consumptions. In particular, we found evidence that (i) coordination of PIT and CS has a prevailing effect on the overall performance of the networked system, and (ii) incorporation of Interests in QoS treatment is vital to cater for resource coordination.

In future work, we will enhance and optimize our service elements and implementations. Due to the constrained nature of our experimental platform, several resource decisions had to remain less elaborate than desired. In particular, the forwarding queues on RIOT had to stay limited due to memory constraints in this challenged regime. We further want to explore mechanisms to securely administer and distribute QoS settings at runtime. Since in our experiment setups the high priority flows did not dominate the network, we plan to investigate the effects of our proposed QoS mechanisms in overbooked network settings that are threatened by service starvation. An experimental deployment in our industrial environment is on schedule as well.

A Note on Reproducibility

We fully support reproducible research [2, 47] and perform all our experiments using open source software and an open access testbed. Code and documentation will be available on Github at <https://github.com/5G-I3/ACM-ICN-2019-QOS>.

Acknowledgements

This work was inspired by many fruitful discussions in the IRTF ICN research group, as well as by industrial deployment demands. Notably, we want to thank Dave Oran, who not only stimulated major design thinking, but also helped as a thoughtful shepherd to improve this paper. It was supported in part by the German Federal Ministry for Education and Research (BMBF) within the projects *I3 – Information Centric Networking for the Industrial Internet*, *RAPstore – RIOT App Store*, and the Hamburg *ahoi.digital* initiative with *SANE*.

REFERENCES

- [1] 3GPP. 2016. *Evolved Universal Terrestrial Radio Access (E-UTRA) and Evolved Universal Terrestrial Radio Access Network (E-UTRAN): Overall description; Stage 2*. Technical Report TS 36.300 13.4.0. 3GPP.
- [2] ACM. Jan., 2017. Result and Artifact Review and Badging. <http://acm.org/publications/policies/artifact-review-badging>.
- [3] Bengt Ahlgren, Christian Dannewitz, Claudio Imbrenda, Dirk Kutscher, and Börje Ohlman. 2012. A Survey of Information-Centric Networking. *IEEE Communications Magazine* 50, 7 (July 2012), 26–36.
- [4] Mays F. Al-Naday, Andreas Bontozoglou, Vassilios G. Vassilakis, and Martin J. Reed. 2014. Quality of Service in an Information-Centric Network. In *Proceedings of Globecom 2014*. 1861–1866.
- [5] M. Amadeo, C. Campolo, A. Iera, and A. Molinaro. 2015. Information Centric Networking in IoT scenarios: The case of a smart home. In *Proc. of IEEE International Conference on Communications (ICC)*. IEEE, Piscataway, NJ, USA, 648–653.
- [6] Sobia Arshad, Muhammad Awais Azam, Mubashir Husain Rehmani, and Jonathan Loo. 2017. Information-Centric Networking Based Caching and Naming Schemes for Internet of Things: A Survey and Future Research Directions. *arXiv preprint arXiv:1710.03473* (2017).
- [7] Onur Ascigil, Sergi René, George Xylomenos, Ioannis Psaras, and George Pavlou. 2017. A Keyword-based ICN-IoT Platform. In *Proc. of 4th ACM Conference on Information-Centric Networking (ICN)*. ACM, New York, NY, USA, 22–28.
- [8] Atmel. 2009. *Low Power 2.4 GHz Transceiver for ZigBee, IEEE 802.15.4, 6LoWPAN, RF4CE, SP100, WirelessHART, and ISM Applications*. Atmel Corporation. <http://www.atmel.com/images/doc8111.pdf>
- [9] Emmanuel Baccelli, Cenk Gündogan, Oliver Hahm, Peter Kietzmann, Martine Lenders, Hauke Petersen, Kaspar Schleiser, Thomas C. Schmidt, and Matthias Wählisch. 2018. RIOT: an Open Source Operating System for Low-end Embedded Devices in the IoT. *IEEE Internet of Things Journal* 5, 6 (December 2018), 4428–4440. <http://dx.doi.org/10.1109/JIOT.2018.2815038>
- [10] Emmanuel Baccelli, Oliver Hahm, Mesut Günes, Matthias Wählisch, and Thomas C. Schmidt. 2013. RIOT OS: Towards an OS for the Internet of Things. In *Proc. of the 32nd IEEE INFOCOM. Poster*. IEEE Press, Piscataway, NJ, USA, 79–80.
- [11] Emmanuel Baccelli, Christian Mehlig, Oliver Hahm, Thomas C. Schmidt, and Matthias Wählisch. 2014. Information Centric Networking in the IoT: Experiments with NDN in the Wild. In *Proc. of 1st ACM Conf. on Information-Centric Networking (ICN-2014)*. ACM, New York, 77–86. <http://dx.doi.org/10.1145/2660129.2660144>
- [12] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss. 1998. *An Architecture for Differentiated Services*. RFC 2475. IETF.
- [13] C. Bormann, M. Ersue, and A. Keranen. 2014. *Terminology for Constrained-Node Networks*. RFC 7228. IETF.
- [14] R. Braden, D. Clark, and S. Shenker. 1994. *Integrated Services in the Internet Architecture: an Overview*. RFC 1633. IETF.
- [15] R. Braden, L. Zhang, S. Berson, S. Herzog, and S. Jamin. 1997. *Resource ReSerVation Protocol (RSVP) – Version 1 Functional Specification*. RFC 2205. IETF.
- [16] Wei Koong Chai, Diliang He, Ioannis Psaras, and George Pavlou. 2013. Cache “Less for More” in Information-Centric Networks (Extended Version). *Computer Communications* 36, 7 (2013), 758–770.
- [17] Asit Chakraborti, Syed Obaid Amin, Aytac Azgin, Satyajayant Misra, and Ravishankar Ravindran. 2018. Using ICN Slicing Framework to Build an IoT Edge Network. In *Proceedings of the 5th ACM Conference on Information-Centric Networking (ICN '18)*. ACM, New York, NY, USA, 214–215.
- [18] Bo Chen, Liang Liu, Zhao Zhang, Wenbo Yang, and Huadong Ma. 2016. BRR-CVR: A Collaborative Caching Strategy for Information-Centric Wireless Sensor Networks. In *12th International Conference on Mobile Ad-Hoc and Sensor Networks (MSN)*. IEEE, 31–38.
- [19] Dong Doan Van and Qingsong Ai. 2018. An Efficient in-Network Caching Decision Algorithm for Internet of Things. *International Journal of Communication Systems* 31, 8 (2018), e3521.
- [20] N. Dukkkipati. 2008. *Rate control protocol (rcp): Congestion control to make flows complete quickly*. Ph.D. Dissertation. Stanford, CA, USA.
- [21] Cenk Gündogan, Peter Kietzmann, Martine Lenders, Hauke Petersen, Thomas C. Schmidt, and Matthias Wählisch. 2018. NDN, CoAP, and MQTT: A Comparative Measurement Study in the IoT. In *Proc. of 5th ACM Conference on Information-Centric Networking (ICN)*. ACM, New York, NY, USA, 159–171. <https://conferences.sigcomm.org/acm-icn/2018/proceedings/icn18-final46.pdf>
- [22] Cenk Gündogan, Thomas C. Schmidt, Matthias Wählisch, Michael Frey, Felix Shzu-Juraschek, and Jakob Pfender. 2019. *Quality of Service for ICN in the IoT*. IRTF Internet Draft – work in progress 01. IRTF. <http://tools.ietf.org/html/draft-gundogan-icng-iotqos>
- [23] Oliver Hahm, Emmanuel Baccelli, Thomas C. Schmidt, Matthias Wählisch, Cedric Adjih, and Laurent Massoulié. 2017. Low-power Internet of Things with NDN and Cooperative Caching. In *Proc. of 4th ACM Conference on Information-Centric Networking (ICN)*. ACM, New York, NY, USA, 98–108.
- [24] Mohamed Ahmed Hail, Marica Amadeo, Antonella Molinaro, and Stefan Fischer. 2015. Caching in Named Data Networking for the Wireless Internet of Things. In *International Conference on Recent Advances in Internet of Things (RIoT)*. IEEE, 1–6.
- [25] Mohamed Ahmed M. Hail, Marica Amadeo, Antonella Molinaro, and Stefan Fischer. 2015. On the Performance of Caching and Forwarding in Information-Centric Networking for the IoT. In *Proceedings of the International Conference on Wired/Wireless Internet Communication (WWIC)*. Springer, Malaga, Spain, 313–326.
- [26] IEEE. 2018. *IEEE Standard for Local and Metropolitan Area Network–Bridges and Bridged Networks*. Technical Report. 1–1993 pages. <https://doi.org/10.1109/IEEESTD.2018.8403927>
- [27] Stratis Ioannidis and Edmund Yeh. 2017. Jointly Optimal Routing and Caching for Arbitrary Network Topologies. In *4th ACM Conference on Information-Centric Networking (ACM-ICN '17)*. ACM, New York, NY, USA, 77–87.
- [28] Van Jacobson, Diana K. Smetters, James D. Thornton, and Michael F. Plass. 2009. Networking Named Content. In *5th Int. Conf. on emerging Networking Experiments and Technologies (ACM CoNEXT'09)*. ACM, New York, NY, USA, 1–12.
- [29] Anil Jangam, Prakash suthar, and Milan Stolic. 2019. *QoS Treatments in ICN using Disaggregated Name Components*. Internet-Draft – work in progress 00. IETF.
- [30] D. Kutscher, S. Eum, K. Pentikousis, I. Psaras, D. Corujo, D. Sauced, T. Schmidt, and M. Waehlich. 2016. *Information-Centric Networking (ICN) Research Challenges*. RFC 7927. IETF.
- [31] Martine Lenders, Peter Kietzmann, Oliver Hahm, Hauke Petersen, Cenk Gündogan, Emmanuel Baccelli, Kaspar Schleiser, Thomas C. Schmidt, and Matthias Wählisch. 2018. *Connecting the World of Embedded Mobiles: The RIOT Approach to Ubiquitous Networking for the Internet of Things*. Technical Report arXiv:1801.02833. Open Archive: arXiv.org. <https://arxiv.org/abs/1801.02833>
- [32] Zhe Li and Gwendal Simon. 2011. Time-Shifted TV in Content Centric Networks: The Case for Cooperative In-Network Caching. In *Proceedings of the IEEE International Conference on Communications (ICC)*. Kyoto, Japan, 1–6.
- [33] Yinlong Liu, Dali Zhu, and Wei Ma. 2016. A Novel Cooperative Caching Scheme for Content Centric Mobile Ad Hoc Networks. In *IEEE Symposium on Computers and Communication (ISCC)*. IEEE, 824–829.
- [34] LoRa Alliance. 2015. *A Technical Overview of LoRa and LoRaWan*. Technical Report. <https://lora-alliance.org/sites/default/files/2018-04/what-is-lorawan.pdf>
- [35] Milad Mahdian, Somaya Arianfar, Jim Gibson, and Dave Oran. 2016. MIRCC: Multipath-aware ICN Rate-based Congestion Control. In *Proceedings of the 3rd ACM Conference on Information-Centric Networking (ACM-ICN '16)*. ACM, New York, NY, USA, 1–10.
- [36] A. Markkanen and D. Shey. 2015. *Edge Analytics in IoT*. Technical Report. ABI Research.
- [37] Philipp Meyer, Till Steinbach, Franz Korf, and Thomas C. Schmidt. 2013. Extending IEEE 802.1 AVB with Time-triggered Scheduling: A Simulation Study of the Coexistence of Synchronous and Asynchronous Traffic. In *2013 IEEE Vehicular Networking Conference (VNC)*. IEEE Press, Piscataway, New Jersey, 47–54.
- [38] Ilya Moiseenko and David Oran. 2019. *Flow Classification in Information Centric Networking*. Internet-Draft – work in progress 04. IETF.
- [39] Sheneela Naz, Rao Naveed Bin Rais, Peer Azmat Shah, Sadaf Yasmin, Amir Qayyum, Seungmin Rho, and Yunyoung Nam. 2018. A Dynamic Caching Strategy for CCN-Based MANETs. *Computer Networks* (2018).
- [40] K. Nichols, S. Blake, F. Baker, and D. Black. 1998. *Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers*. RFC 2474. IETF.
- [41] David Oran. 2019. *Considerations in the development of a QoS Architecture for CCNx-like ICN protocols*. Internet-Draft – work in progress 00. IETF.
- [42] David Oran. 2019. *Maintaining CCNx or NDN flow balance with highly variable data object sizes*. Internet-Draft – work in progress 01. IETF.
- [43] Jakob Pfender, Alvin Valera, and Winston K. G. Seah. 2018. Performance Comparison of Caching Strategies for Information-Centric IoT. In *Proceedings of the 5th ACM Conference on Information-Centric Networking (ICN '18)*. ACM, New York, NY, USA, 43–53.
- [44] Jakob Pfender, Alvin Valera, and Winston K. G. Seah. 2019. Content Delivery Latency of Caching Strategies for Information-Centric IoT. *arXiv:1905.01011 [cs]* (May 2019). <http://arxiv.org/abs/1905.01011> arXiv: 1905.01011.
- [45] George C. Polyzos and Nikos Fotiou. 2015. Building a reliable Internet of Things using Information-Centric Networking. *Journal of Reliable Intelligent Environments* 1, 1 (2015), 47–58.
- [46] Ioannis Psaras, Wei Koong Chai, and George Pavlou. 2012. Probabilistic In-Network Caching for Information-Centric Networks. In *Proceedings of the Second Edition of the ICN Workshop on Information-Centric Networking*. Helsinki, Finland, 55–60.
- [47] Quirin Scheitle, Matthias Wählisch, Oliver Gasser, Thomas C. Schmidt, and Georg Carle. 2017. Towards an Ecosystem for Reproducible Research in Computer Networking. In *Proc. of ACM SIGCOMM Reproducibility Workshop*. ACM, New York, NY, USA, 5–8.
- [48] E. M. Schooler, D. Zage, J. Sedayao, H. Moustafa, A. Brown, and M. Ambrosin. 2017. An Architectural Vision for a Data-Centric IoT: Rethinking Things, Trust and Clouds. In *IEEE 37th Intern. Conference on Distributed Computing Systems (ICDCS)*. IEEE, Piscataway, NJ, USA, 1717–1728.
- [49] Wentao Shang, Adeola Bannis, Teng Liang, Zhehao Wang, Yingdi Yu, Alexander Afanasyev, Jeff Thompson, Jeff Burke, Beichuan Zhang, and Lixia Zhang. 2016.

- Named Data Networking of Things (Invited Paper). In *Proc. of IEEE International Conf. on Internet-of-Things Design and Implementation (IoTDI)*. IEEE Computer Society, Los Alamitos, CA, USA, 117–128.
- [50] Yong Sun, Tiankui Zhang, Ruoqi Wang, Weidong Feng, and Pu Chen. 2016. Topology Potential Based Probability Caching Strategy for Content Centric Ad Hoc Networks. *Journal of Residuals Science & Technology* 13, 6 (2016).
- [51] Saran Tarnoi, Kalika Suksomboon, Wuttipong Kumwilaisak, and Yusheng Ji. 2014. Performance of Probabilistic Caching and Cache Replacement Policies for Content-Centric Networks. In *Proceedings of the IEEE 39th Conference on Local Computer Networks (LCN)*. Edmonton, AB, Canada, 99–106.
- [52] Christian Tschudin, Christopher Scherb, et al. 2018. CCN Lite: Lightweight implementation of the Content Centric Networking protocol. <http://ccn-lite.net>
- [53] Christos Tsilopoulos and George Xylomenos. 2011. Supporting Diverse Traffic Types in Information Centric Networks. In *Proceedings of the ACM SIGCOMM Workshop on Information-Centric Networking (ICN '11)*. ACM, New York, NY, USA, 13–18.
- [54] Matthias Wählisch, Thomas C. Schmidt, and Markus Vahlenkamp. 2013. Backscatter from the Data Plane – Threats to Stability and Security in Information-Centric Network Infrastructure. *Computer Networks* 57, 16 (Nov. 2013), 3192–3206. <http://dx.doi.org/10.1016/j.comnet.2013.07.009>
- [55] Michael Welzl. 2005. *Network Congestion Control - Managing Internet Traffic*. Wiley, Hoboken, NJ, USA.
- [56] George Xylomenos, Christopher N. Ververidis, Vasilios A. Siris, Nikos Fotiou, Christos Tsilopoulos, Xenofon Vasilakos, Konstantinos V. Katsaros, and George C. Polyzos. 2014. A Survey of Information-Centric Networking Research. *IEEE Communications Surveys and Tutorials* 16, 2 (2014), 1024–1049.
- [57] Yuguang Zeng and Xiaoyan Hong. 2011. A Caching Strategy in Mobile Ad Hoc Named Data Network. In *6th International ICST Conference on Communications and Networking in China (CHINACOM)*. IEEE, 805–809.
- [58] Lixia Zhang, Alexander Afanasyev, Jeffrey Burke, Van Jacobson, kc claffy, Patrick Crowley, Christos Papadopoulos, Lan Wang, and Beichuan Zhang. 2014. Named Data Networking. *SIGCOMM Comput. Commun. Rev.* 44, 3 (2014), 66–73.
- [59] Li Zhang, Jiayan Zhao, and Zhenlian Shi. 2015. LF: A Caching Strategy for Named Data Mobile Ad Hoc Networks. In *Proceedings of the 4th International Conference on Computer Engineering and Networks*. Springer, 279–290.
- [60] Meng Zhang, Hongbin Luo, and Hongke Zhang. 2015. A Survey of Caching Mechanisms in Information-Centric Networking. *IEEE Communications Surveys & Tutorials* 17, 3 (2015), 1473–1499.
- [61] Le Zhou, Tiankui Zhang, Xiaogeng Xu, Zhimin Zeng, and Yinlong Liu. 2015. Broadcasting Based Neighborhood Cooperative Caching for Content Centric Ad Hoc Networks. In *IEEE/CIC International Conference on Communications in China (ICCC)*. IEEE, 1–5.