Lessons from the Past: Why Data-driven States Harm Future Information-Centric Networking

Matthias Wählisch
Freie Universität Berlin
waehlisch@ieee.org

Thomas C. Schmidt
HAW Hamburg
t.schmidt@ieee.org

Markus Vahlenkamp
HAW Hamburg
markus@vahlenkamp.net

Abstract—Information-centric networking (ICN) raises data objects to first class routable entities in the network and changes the Internet paradigm from host-centric connectivity to data-oriented publish/subscribe. We revisit the data-centric paradigm from the perspective of security and resilience and question its applicability in an open, widely distributed routing and forwarding service. Current concepts of content routing are built on data-driven protocol events and thereby introduce a strong coupling of the control to the data plane in the underlying routing infrastructure. In this paper, we explore the vulnerability of the distribution backbone. Based on a straight-forward analytical model we show that local systems cannot be protected from the threats of data-driven state management on an Internet scale. By practical evaluations using the example of the CCNx implementation, we further analyze threats to stability and performance of a data-driven infrastructure that refrains from separating the control from the data plane. We identify intrinsic attack vectors, as well as possibilities and limitations to mitigate them. Our overall findings suggest that major architectural refinements are required prior to global ICN deployment in the real world.

Index Terms—Security, vulnerability, performance evaluation, content-centric routing, denial of service (DoS)

I. Introduction

A basic design principle of the traditional Internet is to restrict control information to topological guides on where to direct packets. Control plane states are generated from provider configurations and dynamic updates among neighbors, all carried in explicit router-to-router communication. States are used but not altered by the forwarding plane. Consequently, suppliers of Internet infrastructure can architect routers that separate the control from the data plane.

Inspired by the use case of widely deployed Content Delivery Networks (CDNs), current trends of *Information-Centric Networking (ICN)* shift the Internet towards data awareness. In ICN, consumers shall retrieve content by name directly from a network that provides storage, caching, content-based rendezvous, and searching at times. Thereby data sets become first class routable objects and content names require exposure to the control plane. Several ICN proposals have been presented in recent years [1], among them TRIAD [2], DONA [3], NDN [4], [5], PSIRP [6], and NetInf [7], which differ in several design choices. As we are interested in the stability and security of ICN infrastructures, we will concentrate on the aspects of routing and forwarding.

Essentially two approaches to routing exist in current ICN proposals, an evolutionary path that resolves names to locators and routes on IP (or a related location scheme), and 'clean slate' concepts that route directly on content names. NetInf extends the current Internet by a resolution service that maps content names to topological IDs like IP addresses, but alternatively supports name-based routing. TRIAD, DONA, and NDN perform content retrieval by routing on names. Route responses and the data itself are then forwarded along reverse paths (RPF), either by using IP as a lower layer, or without IP but by dedicated RPF states. PSIRP publishes content objects to a resolution system that incloses full knowledge of the network topology. Requesters trigger the mapping system to generate source routing identifiers in the form of Bloom filters that aggregate IDs of forwarding links.

All solutions operate on the content itself, and force the network infrastructure into a content awareness. A mapping service is not only required to resolve *file* names to source locations, but must answer a request by advising a nearby replica, the existence of which it learned from the data distribution system. Content routers need to rely on (often aggregated) names in its interface tables and—for RPF-based forwarding schemes—a reverse state for every data unit. This control information is highly dynamic and requires regular updates from the data plane. The ICN paradigm thereby opens up the control plane to continuous modifications from the data plane. This is in contrast to the current Internet, where DNS and routing states remain unaltered when a Web page is published, a file is transfered, or data is cached.

In this paper, we study the impact of traffic conditions on the control plane. We are in particular interested in threats to the stability and security of the ICN infrastructure, whose impacts we evaluate in a theoretical analysis and experimental trials. Experiments are performed in test networks running PARC's CCNx software. We want to stress, though, that our tests only attribute for the core concepts of content routing and do not evaluate implementation properties of the CCNx prototype. Following basic insights gained from theoretical and practical analysis, we contribute a sample set of attacks that ground on this correlation of data with control states. We argue that the novelty of these exposures derives from an intrinsic binding to ICN concepts so that attacks—even if reminiscent from today's Internet—cannot be mitigated by simple protocol provisions.

The remainder of this paper is organized as follows. The specific problems in protecting the ICN infrastructure are stated in § II along with related work on ICN security. We theoretically analyze basic threats to stability in § III and discuss related implications. Based on practical experiments, threatening scenarios and their effects on the routing system are demonstrated in § IV. These general insights lead to concrete attack scenarios in § V. The paper concludes with a discussion in § VI.

II. WHY ICN IS CHALLENGED BY DESIGN

A. ICN System Model

Information-centric networking involves two functional blocks within the network infrastructure, (1) content publications or announcements, and (2) content subscriptions or (asynchronous) access. Throughout this paper, we assume a generic ICN system model that is composed of these two subsystems, both of which introduce routing or forwarding states at the network layer. Even though not all ICN proposals are constructed equally pronounced in both parts, they all update corresponding table entries in response to data operations of the network infrastructure. In addition, we assume that universal caching is implemented in the content-centric routing system. Universal caching is common to all ICN solutions.

Content requests and delivery do not follow an end-to-end design, but require a dynamic set-up of paths between the requester and a (nearby) copy of the data. Commonly, this is done by Reverse Path Forwarding (RPF), at which each content request triggers a trail of 'bread crumb' states on routers along the path (NDN, DONA, NetInf). Alternative approaches that route on an underlying routing substrate like IP (NetInf), or construct source routing identifiers based on complete knowledge of the topology (PSIRP), are not considered further in this work.

B. Problem Statement

Publishing and subscribing in current ICN solutions introduce network control states that generate the following management problems.

(1) Addressable content items need advertisement in the route resolution system. Consequently, any end user who can publish requires admission to modify the control plane.

(2) Content is conceptually delocalized by universal caching. Data replication thus imposes updates of the routing systems—a change of control state initiated by the data plane.

(3) Reverse Path Forwarding requires state initiation and consumption at routers along the path. Corresponding control state updates are not only driven by the data plane, but require processing in wire-speed.

These state operations raise the following threat classes in ways that are unique to ICN.

Resource Exhaustion Infrastructural entities need to offer accumulating resources like memory and processing power for provisioning, maintaining, and exchanging content states. They are therefore threatened by resource exhaustion due to misuse or uncontrolled load. In addition, the asymmetry in size between data requests and delivery leads to traffic amplification when exploited in DoS attacks.

State Decorrelation The asynchronous nature of publish/subscribe content delivery places the enhanced burden of assuring consistency among distributed data states. Data states that require correlation are situated in distributed mapping systems, which also need to consistently reflect actual content placements, and in forwarding states at routers that define the paths hop-by-hop from a supplier to the requester. Failures in state coherence lead to service disruptions or unwanted traffic flows.

Path & Name Infiltration The infrastructure relies on the integrity and correctness of content routing and is therefore threatened by poisonous injections of paths and names, in particular. The replicative ICN environment distributes content copies to many, commonly untrusted locations and thereby makes it particularly hard to authenticate valid origins of state insertion requests.

All of these threats bear the potential to seriously degrade the ICN service and lead to insufficient or erroneous data dissemination. A major risk for the ICN infrastructure—and from a general perspective for the ICN concept—results from the power that an end user gains over an ICN distribution backbone.

C. Related work

Content Suppliers Related work on ICN security has primarily focused on validating content correctness and authenticity. Commonly, self-certifying security credentials are included in 'secure names' that facilitate mechanisms for verifying authors, origins, and content integrity [8], [9], [10], [11]. Thus a receiver can be sure to obtain the correct content and an intermediate cache can validate the correctness of the security credentials, which prevents traditional DoS on the ICN system [12]. Nevertheless, having created (or learned) a valid name, any ICN member can re-announce this in the route resolution service, thereby injecting poisonous routes or artificial names into the system.[1] Similar vulnerabilities of DNS and BGP are known from today's Internet infrastructure [13], but remain restricted to (topology) *providers*. ICN opens the liberty of route injection to the group of content suppliers (i.e., *end users*). We will discuss threats unique to ICN in Section III-A.

Content Consumers Little attention has been given to the effects of state management in ICN. Arianfar *et al.* [14] discuss design choices for an ICN router. They concentrate on the content cache and explicitly do not consider per request states. Perino and Varvello [15] have evaluated requirements for content routers that hold content information bases in Bloom filters and reverse paths in pending interest tables (PITs). Under the assumptions of *valid* content requests propagated on *homogeneous* network links with a *maximum global* RTT of 80 ms, average PIT sizes are identified in the order of 1 Gbit/s for current line speeds. FIB sizes and lookup complexity were shown to depend nonlinearly on prefix numbers and name lengths. Lauinger [16] explicitly addresses the threat of DoS attacks by filling the available memory of a router with pending interest states.

Such attacks on hardware resources may be mitigated by limiting overall table sizes. However, securing router resources by table limits does degrade network utilization and cannot abandon resource exhaustion problems. In the presence of a

---

[1]As a countermeasure, DONA introduces certificates of publishers on the price of per cache-instance varying names. Content routing then works on wildcarding names, which re-introduces the threat of route poisoning.

table limit, an attacker could initiate massive drops of pending Interests from a router's table and thus disrupt data delivery to regular receivers. The author in [16] proposes to drop Interests at the head of the PIT, which however may easily be misused to DoS-attacking neighbors, or to use Bloom filters instead of PIs. If applied without strict capacity limits, the latter approach is vulnerable to flooding attacks as interface filters degrade their selectivity. In the following section, we will evaluate these effects in detail.

Request state management and related security issues have been recently raised in [17], [18], [19]. Gasti *et al.* [20] address core issues of route hijacking, state overload, and cache pollution in NDN. They propose counter measures by extending interface functions, e.g., for limiting rates and survey content delivery. Without considering protective measures in BGP, the authors compare BGP with NDN security and argue that the NDN approach reduces vulnerability to blackholing, as routers can identify unresolved content requests and rank/re-route per prefix and interface. Authors miss that on the one hand RPKI secures BGP against hijacking attacks in a straight-forward manner [21], while on the other hand proposed countermeasures in ICN cannot prevent attacks of interception and redirection with service degradation.

Intermediate Summary   ICN opens the control plane of backbone routers for content consumers and suppliers on a fine-grained base. Granting end users access to the routing and forwarding subsystems is a fundamental step away from the current Internet design and bears significant risks. Current concerns in the context of routing mainly focus on state explosion due to the large amount of content items. One might argue that those resource exhaustions will be solved by more powerful hardware in the future. We will discuss options and limitations of related core aspects in Section III. Still, binding the integrity of the routing infrastructure to the courtesy of *all users* is intrinsic to current ICN approaches—and presumably to the overall ICN concept.

III. Basic Threats to Stability

In this section, we theoretically examine the implications at the control plane for the different data operations and discuss resulting threats that inherently arise at the infrastructure level.

A. Routing or Mapping Resources

The common view on routing is that of a topological resolution service: Routing guides the paths to hosts. As ICN abandons the host-centric paradigm to address content objects directly, routes to content items attain the role of traditional topological directives.

State and Update Complexity   In ICN, each content item (file) needs retrieval and therefore must be accessible via some resolution service. This may either be implemented by a distributed routing system, or by a mapping service that provides an indirection to topological locators of publishers or content caches. The average complexity of the corresponding management operations reads $\langle \# \ of \ content \ items \rangle \cdot \langle \# \ cached \ replica \rangle \cdot \langle update \ frequency \rangle$ ($\langle \cdot \rangle$ denots average

values) and must be considered a severe challenge.[2] As a consequence, the request routing/mapping system is stressed by adding and updating name or cache entries at overwhelming frequency, the details of which depend on the implementation of the service.

Cache Announcements   Route maintenance in ICN consists of propagating content publishers (i.e., default paths) as well as cache instances. While the first task is known to generate a high volume of data and frequent updates, caching is expected to largely exceed default announcements in number and update frequency. As a countermeasure, data replication may be limited to caching along default paths, which remarkably reduces the complexity for the routing system. On-path cache replica are met implicitly when requests are routed towards the source. They need not be advertised in the routing or mapping service. On the downside, restricting the caching to default paths will drastically reduce its effectiveness, and a corresponding strategy falls behind today's CDN solutions. Ghodsi *et al.* [12] discussed the caching problems in detail. The authors came to the conclusion that on-path caching is merely a warm-up of traditional web proxies.

Route Integrity   ICN, like the current Internet, relies on the integrity of its routing system. A bogus route may block or degrade services, lead to incorrect content delivery, or violate privacy. These core concerns are well-known from BGP [13], where effective countermeasures exist. However, in addition to those vulnerabilities known from BGP routing, threats uniquely arise from data-driven state management in content-centric routing.

The first issue is inherited from universal caching. An explicit authorization of caches as common in the CDN market is in conflict with open publication and not applicable in general ICN approaches. Rather any node in the network can cache and thus announce any (forged) name, while origin validation measures such as RPKI [22] or [23] cannot be applied. The second issue emerges directly from state maintenance at routers. As the routing infrastructure is vulnerable to increased delays and delay variations in content supply (see Section III-B), route redirections may be applied to slow down content delivery or to jitter response times. Following the first argument, any intermediate cache can—purposefully or accidentally—threaten its neighborhood.

B. Forwarding Resources

Traditional routers in the Internet consist of a central processing unit and main memory that are available to the control plane, mainly to learn and determine new routes, as well as FIB memory that is fed by the route selection process. Data forwarding remains bound to FIB lookup and packet processing at line-cards. This design choice purposefully decouples forwarding capacities from control processing and—with equal importance—protects control states from (bogus) data packets.

---

[2]A global request routing system will need to host at least the amount of the Google index base ($\mathcal{O}(10^{12})$) at a much enhanced update frequency (by caching). For comparison, today's DNS subsumes $\mathcal{O}(10^8)$ names at a very low change rate of $\approx 10^5$ alterations per day.

R_i	The i-th Router
C_i	Capacity of the link between R_i and R_{i+1}
U_i	Utilization of the link between R_i and R_{i+1}
S_i	# of content request states of R_i at its interface towards R_{i+1}
α_i	Content request rate at interface $R_i \rightarrow R_{i+1}$
ω_i	Content arrival rate at interface $R_i \leftarrow R_{i+1}$
T_i	Request timeout at interface $R_i \rightarrow R_{i+1}$
l	Packet length
$\langle \cdot \rangle$	Average value of $\cdot$
$\sigma(\cdot)$	Standard deviation of $\cdot$

Table I
GLOSSARY OF NOTATIONS

Figure 1. Distribution of forwarding states at routers with a 1 Gbit/s link, covering global RTTs [25] in March 2012

Current concepts of content-centric data forwarding break with this separation paradigm, and introduce—similar to IP multicast—an additional reverse path forwarding table, also called PIT. Unlike in multicast, this table is updated *packet-wise* on line speed by data-driven events. In the following subsections, we concentrate on the consequences for routing resources in detail. We will consider a chain of routers R_i along a data path and use the notation summarized in Table I.

1) Content Request States versus Content Request Rates versus Network Utilization: Content request states are the essential building block to control flows in a content-centric distribution system that operates hop-by-hop. Each request state will trigger a data packet on return, why the number of open request states corresponds to data arrival at this interface after the transmission time.

Consider a point-to-point interface at routers R_i in steady operation and in the presence of a (per interface) state timeout T_i. In the absence of request retransmissions, packet loss, and state dismissal, we first want to derive the relation between routing request states at time t and network utilization. The total amount of state increases linearly by newly arriving requests α_i and decreases by content arrivals ω_i. Hence, the basic rate equation reads

$$
\begin{aligned}
S_i(t) &= S_i(t - T_i) + \int_{t-T_i}^{t} \{\alpha_i(\tau) - \omega_i(\tau)\}\, d\tau \\
&= S_i(t - T_i) + \int_{t-T_i}^{t} \{\alpha_i(\tau) - \alpha_i(\pi(\tau))\}\, d\tau \\
&= \langle \alpha_i \rangle \cdot \min(\langle RTT \rangle, T_i) + \\
&\qquad \mathcal{O}\left(\sigma(\alpha_i) \cdot \sigma(\min(RTT, T_i))\right),
\end{aligned}
\tag{1}
$$

where $\pi(\cdot)$ denotes the time delay of the packet arrival process and RTT the random variable of packet round trip times, which is assumed independent of the requests and packet rates.

From Equation (1), we can immediately deduce that timeout values below the (varying) $RTTs$ limit the number of request states, but at the same time will block data forwarding. A second view reveals the strong dependence of routing state on the RTT variation. A similar phenomenon is well-known from TCP [24], but has been overlooked in corresponding previous work on ICN resource considerations [15], [19], [20].

Henceforth we will address the case of data flowing unhindered by the state timeout T_i and assume T_i large enough.
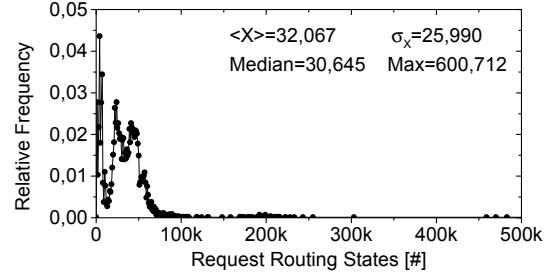
Furthermore—for a steady-state scenario—it is assumed that the content request rate fluctuates on a stationary scale. Equation (1) then simplifies to

$$
\begin{aligned}
S_i(t) &\approx \langle \alpha_i \rangle \cdot (\langle RTT \rangle + \kappa\, \sigma(RTT)) \tag{2} \\
&\approx U_i(t)/\langle l \rangle\, (\langle RTT \rangle + \kappa\, \sigma(RTT)), \tag{3}
\end{aligned}
$$

with an estimating parameter κ for the mean deviation. The well-known term $(\langle RTT \rangle + \kappa\, \sigma(RTT))$ represents a retransmission timeout.[3] For the last step, we roughly assumed that content requests and content arrival are in stationary equilibrium.

Approximation (3) yields the desired coupling of the link utilization U_i and the state management resources at a router: On a single point-to-point link without state retransmissions and in flow balance, state requirements are proportional to the network utilization, enhanced by a factor of a *global retransmission timeout.* At switched interconnects or in bursty communication scenarios, conditions are expected to grow much worse.

The following observations are noteworthy.

1) Unlike in TCP that estimates a single end-to-end connection, content request states at routers subsume various prefixes and numerous flows. Moreover, content items (prefixes) are explicitly not bound to end points. Thus rapidly varying RTTs are characteristic to interfaces and even to individual flows in content-centric routing. The presence of chunk caching may further increase the RTT variation. Hence, no convergent estimator for a round trip time can be reasonably given.

2) In the current Internet, the variation of RTT is commonly larger than its average. End-to-end delays are known to approximately follow a heavy-tailed Gamma distribution [26]. PingER [25] reports means and standard deviations of about 250 ms, with maxima up to 5,000 ms. For a constant content request rate of 125k packets/s these RTTs generate the state distribution visualized in Figure 1.

3) Limiting the absolute size of the content request table imposes a strict bound on network utilization. However,

---

[3]The corresponding (over-)estimator in TCP is commonly set to 4. However, it is well known that standard TCP algorithms and parameters are inefficient at rapidly changing round trip times, which are characteristic for interface conditions in content-centric routing.

the sustained rates are mainly determined by actual RTTs and are hardly predictable. Similar arguments hold for defining timeout values.

4) Applying rate limits to content requests does not change the picture. For an 'on average' optimal limit $C_i \cdot \langle RTT \rangle / \langle l \rangle$, the variation of content replies in time may lead to large over- and under-utilization of network resources that goes along with large fluctuations in request table sizes.

2) Memory Requirements: A content-centric router that is designed to fully utilize its link capacities, requires sufficient table space for content requests under varying network conditions. Equation (3) approximates the corresponding resources when applied to the maximum link capacity C_i. Using the conservative value of $\kappa = 4$ as for TCP, a packet length $l = 1,000$ bytes, and RTT values from PingER as cited in the previous section, we derive

$$S_i = 1,25\,s/8.000\,bit \cdot C_i \approx 1,6 \cdot 10^{-4} s/bit \cdot C_i \quad (4)$$

For a line-speed of 1 to 100 Gbit/s, 160k to 16,000k content request entries then need to be installed per interface at minimum. Due to the more accurate consideration of RTT variation terms, these findings differ from previous results [15], [19] by more than an order of magnitude. Still they are merely a rough *lower estimate*, as larger fluctuations of round trip times may significantly increase resource demands.

It is noteworthy that Equation (4) holds for any router in a content-centric Internet. Unlike today, where full BGP tables are only required at AS border routers, and interior devices operate on a very small routing table, ICN access routers already demand for a full table memory, the size of which is determined by its interface capacities. In practice, this significantly increases router costs, as any fast interface must co-locate a large block of fast memory.

3) CPU Load from Table Management: An ICN router maintains states according to user data requests. For any content request, it needs at line speed to (1) insert a state in its request table. On the arrival of any data packet, it needs to (2) search and (3) delete on success in the same table. In addition, a router has to (4) maintain timers of all (soft) states in its request table. To guarantee robustness, an implementation of the huge request table not only needs to perform dictionary operations very efficiently *on average* but also in *worst-case*. With today's hash table implementations in software or hardware this is impossible to achieve [27], [28].

IV. EXPERIMENTS ON STATE-BASED FORWARDING

In this section, we present the results of straight-forward experiments that show the outcome of the core threats as theoretically discussed in Section III. In particular, we concentrate on system and performance implications of the data-driven state management at infrastructure devices. Even though the measurements mainly relate to the NDN implementation `ccnd`, we should emphasize that we do not evaluate the implementation itself, but use it as one real-world instance of


Figure 2. Topology of the experimental setting

the information-centric network deployment to illustrate the routing protocol mechanisms. Following this spirit, we do not interpret or discuss absolute performance values, which surely can be improved by optimized software and hardware in the future, but focus on structural and asymptotic analysis.

A. Core Measurement Setup

In our measurement study, we intentionally deploy *simple* communication scenarios between one content requester and one publisher. The network topology is represented by a Daisy chain of directly interlinked CCNx routers with 100 Mbit/s, one end connects the content consumer and the other the content repository (see Fig. 2). The *basic topology* consists of two hops and the *extended topology* of five nodes. It is noteworthy that more complex settings, e.g., a Dumbbell topology popular to represent backbone network effects, would enforce the effects, which we already see in our simpler and more transparent examples.

We use the CCNx implementation version 0.5.1 [29], i.e., the client library to announce content Interests, the content repository to store data, and the `ccnd` to forward subscription and data. The following analysis focuses on the effects on the router side. For obtaining a fine-grained view, we concentrate on the local system as well as inter-router dependencies.

We keep default values for all CCNx parameters. In particular, routers do not follow a specific strategy layer, as this would twist robustness towards specific limits as discussed in Section III-B. CCNx routers communicate via TCP (preserving packet order in the basic experiments) or UDP (extended experiments).

B. Basic Experiments: Resource Consumption

1) A Fast Path to Resource Exhaustion: An elementary threat intrinsic to data-driven state management arises from the overloading of routers by Interest requests. This is most easily provoked by initiating requests for content that does *not* exist. In our scenario, the consumer issues 2,000 Interest messages for *non*-existing content, waits 6 seconds, and repeats these steps until overall 150,000 Interests have been sent.

Figure 3 shows the local resource consumptions on the first hop of the content receiver. The number of entries in the Pending Interest Table (PIT), the CPU load, and the required memory increase linearly with subsequent bulks of Interest messages until the system is saturated. In this case, the router reaches its limits of processing and memory resources when

Figure 3. Load at the designated router of the receiver while requesting non-existing content

storing $\approx 120,000$ PIT entries. While sending Interests, the initiating node retransmits previous announcements to keep states fresh at the router. Even though the retransmission timer is below the expiration timer and network delays are very short, the PIT size fluctuates as entries drop due to overloading. After all initial Interest messages have been distributed, the content consumer only retransmits subscriptions.

Our experiment illustrates several problems: A router may easily exhaust PIT space, when content arrives late or not at all. However, even if it was able to store all entries, it would suffer from a *retransmission only* phase. The retransmissions agglomerate over time and create a continuous stream of signaling that consumes CPU cycles. When the update rate is higher than the processing capabilities permit, retransmissions require buffering, which leads to additional memory overhead (cf., Figure 3). A high system load increases the probability of dropping a PIT entry even if its refresh message has been signaled in time. This again causes additional refreshs of the PIT data structure (add/delete calls) and fosters load.

In a recent publication, Yi *et al.* [30] propose to mitigate this threat by signaling content unavailability back to the original requester. Such `NACK` will cure the Interest retransmission effects discussed above for truly unavailable content. However, this workaround has limited effect, as `NACK` suppression introduces a new attack vector at the content supplier side, while a bogus requester can still harm the routing infrastructure (in particular its designated router) by iterating Interest messages over various names of unavailable content.

2) Chunk-based State Multiplication: To analyze the performance of content consumption, we conduct a bulk file transfer. At this, the content receiver initiates the parallel download of multiple 10 Mbit files over a constant time. We consider three scenarios, the request of 2 files, 10 files, and 100 files per second, which correspond to an underutilized, a fully loaded, and an overloaded link. Figure 4 shows the start and

completion time of the download per file (top graph), as well as the PIT size, the effective number of Interest retransmissions, and the traffic load including the mean goodput at the first hop. For visibility reasons, we rescaled the y-axis of PI in Figure 4(a).

With an increasing number of parallel downloads, not only the download times increase significantly, but also the interval of the request and receive phase grows in the scenarios of (over-)load. While the download time is almost constant for two files per second (cf., Fig. 4(a)), the time-to-completion grows non-linearly for the downloads in cases of excessive parallelism (cf., Fig. 4(b),(c)). 150 s are needed to download *each* single file in the worst case (Fig. 4(c)), while the link capacity would permit to retrieve *all* files in about 10 s.

The reason for this performance flaw is visualized in the subjacent graphs. A higher download frequency leads to an increasing number of simultaneous PIT entries, which require coordination with the data plane. Each file request will be split into requests of multiple chunks, in which the generation of corresponding Interest messages will be pipelined. In contrast to Section IV-B1, content exists. As soon as the content traverses, Interest states dissolve and thus release memory. These operations cause a simultaneous burst in CPU load (not shown) and result in growing Interest retransmits after droppings or timeouts (shown in second lowest graphs). This also leads to retransmissions of data chunks. As an overall net effect, the network utilization fluctuates significantly, but does not adapt to actual user demands: Even though data requests could fill the links easily, the average load remains about constant at 30 % of the total network capacity.

In this example we demonstrated that insufficient processing and memory resources will strictly prevent a proper link utilization. This problem cannot be mitigated by rate limiting, as reduced Interest transmission rates will simultaneously reduce network utilization even further (see Section III-B1).[4]

The only visible way to assure proper utilization of network resources requires appropriate routing resources, i.e., a PI table implementation that is sufficiently large and reliably operates at line speed. As we learned from the analysis in Section III-B, corresponding solutions are not available today. At the current state of the art, an attacker can always reproduce the performance degradations by either blowing up RTT and its variation, or by injecting states that degrade the performance of the PI hash table of the routers, i.e, complexity attacks.

C. Extended Experiments: State Propagation and Correlation

In our extended experiments, we take a closer look at hop-by-hop routing performance using the five node routing chain displayed in the lower part of Figure 2. Intermediate nodes are numbered from the designated router of the content receiver (first hop) to the router of the content repository (fifth hop). In the following three experiments, we specifically concentrate

---

[4]We should remind that applying Interest rates in NDN is a mechanism of flow control, and *not* for system resource protection. Intermingling these two aspects is likely to produce unwanted performance flaws and leads to new attacks (cf., Section V).

(a) 2 files per second

(b) 10 files per second

(c) 100 files per second

Figure 4. Parallel download of 10 Mbit files: Start and stop time of the download per file at the receiver & resource consumption at its designated router [Pending Interests (PI), Interest Retransmits (IR), and Network Load (NL) including the mean goodput]

(a) Memory Consumption
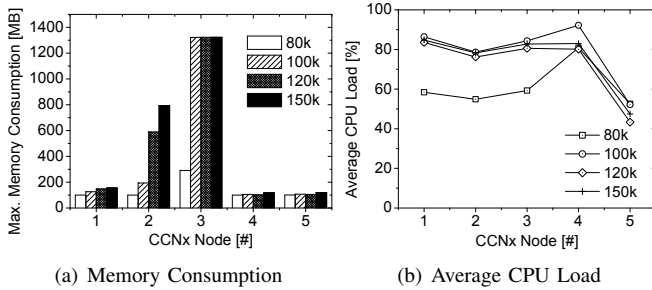
(b) Average CPU Load

Figure 5. Load per hop for a chain of 5 routers while initiating a 80k, 100k, 120k, and 150k different Interests for non-existing content

on correlation effects of the routing resources by controlling the environment using parametrizable virtual machines.

1) A Homogeneous Network: In this first extended experiment, we simply move our previous picture to the larger topology. All forwarding nodes offer the same resources, two cores@2.4 GHz, 3 GB RAM, and link capacities of 100 Mbit/s. A content requester downloads 500 files of size 10 Mbit at an average rate of 100 files per second. We observe a flattening of Interest propagation towards the source, as states resolve earlier from faster packet delivery (cf., Fig. 7(a)).

2) A Single Point of Weakness: It is a valid assumption that the content distribution system will consist of heterogeneous devices in terms of all performance metrics. In this second experiment, we introduce device heterogeneity by weakening a single router, the 4th hop (CCNx4), in a controlled way. We want to study the reaction of state management and network performance to this well-defined degradation.

For an initial observation of the dependency on the weakest node, we reduce the CPU capacity of CCNx4 to 25 % (600 MHz) and recap the scenario from Section IV-B1 for 80k, 100k, 120k, and 150k subscriptions of non-existing content. Independent of the capacity of the network infrastructure, the consumer initiates content subscriptions and continuously refreshes its Interests, which then propagate towards the content repository.

Figure 5 shows the maximal memory consumption and the average CPU load per hop during the measurement period. It is clearly visible that the required memory mainly depends on the position of the node within the topology. Memory requirements on the single path fluctuate by two orders of magnitude. The predecessor of the node with the lowest processing capacities (i.e., the 3rd hop) needs 50% – 500% more memory than any other nodes.

We now take a closer look on gradual effects of routing heterogeneity. We observe corrective mechanisms of the network (i.e., Interest retransmissions) depending on router asymmetry. Interest retransmissions serve as the key indicator for timeouts due to router overload. For this task, we configure CCNx4 with four different processing capacities related to the other CCNx routers: 2,400 MHz (homogeneous), 1,200 MHz (50 % capacity), and 600 MHz (25 % capacity).

Surprising results are shown in Figure 6. Evidently we see an instability in the forwarding behaviour of the network. The characteristic picture of a balanced network is a steady decay of Interest retransmits towards the source, as data delivery gets faster and more reliable in proximity to the publisher. However, at the first occasion of a 'bottleneck'—independent of its strength—the picture flips. Interest retransmission drastically increases and all routers except for the bottleneck equally see about the maximal rate of retransmissions in this scenario. State retransmissions at the weak forwarder (CCNx4) instantaneously doubles to the maximal level of managed states this router can cope with.

This experiment clearly shows how sensitive content-centric routing reacts to varying network resources. A light disturbance of the state propagation process reveals the instability of a steady-state flow by immediately turning content transport into a significantly different condition of maximal error management.

3) Complex Inhomogeneities: In our final experiment concerned with content routing, we explore situations of largely decorrelated network conditions. Therefore we configure all routers to admit fast changing resources occurring in anti-

Figure 6.   Effect of routing heterogeneity on Interest trading

cycles. In detail, each router (CCNx1, . . . , CCNx5) is forced into a 10 s periodic CPU reduction by 90 %. Resource reduction periods were shifted between routers at a rate of 10 s so that at least one of the three routers in the forwarding chain was kept in challenged conditions. The objective of this repelling setup, which similarly may well occur from different side traffics in a meshed backbone, is to analyse the vulnerability of hop-by-hop state maintenance in ICN routing.

A comparative result of the different scenarios in our experimentally-driven analysis is presented in Figure 7. We contrast the load imposed onto the infrastructure by Interest states with the average network performance in the three experimental scenarios, homogeneous network, single point of weakness, and alternating resources at routers. The striking picture in all three settings is that the efficiency of network utilization is low on the overall, but drastically drops whenever inhomogeneities occur. The hop-by-hop forwarding performance thus appears rather fragile. In contrast, network state propagation attains various patterns, but always remains at compatible level at the router of maximal load.

These observations suggest the following rule of thumb for CCN routing performance: State maintenance always follows the maximal requirements, while forwarding performance will adapt to the weakest resource in place. This overall picture is clearly inefficient and future work on ICN solutions would largely benefit from improving this behaviour.

V. EXAMPLES OF ATTACK SCENARIOS

In this section, we briefly introduce an attack scenario for each threat enumerated in Section II-B (see [31] for further attack vectors). Some attacks are unique to ICN, others—even though known from the Internet—gain a new level of severity by exploiting ICN intrinsics.

Resource Exhaustion: Mobile Blockade   A mobile node may issue a large number of invalid (or slow) Interests that block the state table of the access router for the period of state timeout. In a shared link-layer environment that cannot easily detect its departure, the mobile adversary can traverse neighboring networks on circular routes and continue to offload its interest bundle with the effect of a blockade of the regionally available networks. Initial countermeasures are difficult to apply, as the retransmission of Interests is part of the regular mobility pattern in ICN.

State Decorrelation: Heterogeneity Attack   An attacker that controls several machines (e.g., a botnet) may direct requests to accumulate at a specific router in the network and generate a point of performance degradation in the core. Heterogeneity will cause a significant service depletion for all crossing flows (see Section IV-C), if the network does not reroute. In the presence of rerouting, the adversary may use the same attack to trigger route flipping with corresponding jitter enhancements, which—in contrast to the Internet—will degrade access router performance for consumers.

Path and Name Infiltration: Route-to-Death   An adversary that controls a cache system may redirect routes to it and slow down content delivery or jitter response times. As the routing infrastructure is vulnerable to increased delays and delay variations, resource exhaustion threats apply to the requesting infrastructure (see Section III-B). In the presence of universal caching, reasonable counter measures to using a valid, but alienating cache are difficult to define.

VI. DISCUSSIONS AND CONCLUSIONS

In this paper, we have analyzed network instabilities and threats in information-centric networks that are caused by (a) backbone control states initiated by end users and (b) data-driven state management.

Some threats are easy to anticipate (e.g., resource exhaustion), others are more intricate due to the complex interplay of distributed management (e.g., state decorrelation). For the latter previous practical insights in the design of (conceptually related) multicast protocols already revealed good and bad design options. One of the major design goals of Bidirectional PIM [32], for example, was „eliminating the requirement for data-driven protocol events"—after the operating experiences with data-driven DVMRP or PIM-SM. With this paper, we want to stimulate the discussion about basic security in content-centric backbone routing.

Today, (D)DoS attacks are usually directed towards end hosts. In this paper, we have shown that ICN extends these threats to the backbone by design, and that existing countermeasures against both, DDoS and incorrect distribution states fail in the ICN field.

Defending from DDoS is already complicated in the Internet and becomes more intricate in ICN. From the conceptual perspective, the core challenge is not in deploying accountability (e.g., [33], [34]) but identifying an attack. Attack detection approaches [35] usually make application specific assumptions about traffic patterns, which cannot be applied to a generic Internet service for content delivery. We showed that the very fluctuating Internet delay space challenges resource provision in ICN (cf., § III-B). As content states will accumulate in the network (cf., § IV), and inter-provider deployment almost surely will lead to a heterogeneous, unbalanced design, rate limiting may milden, but cannot effectively prevent the resource exhaustion problems discussed in this paper.

Current CDN deployments remain agnostic of these infringements by running under proprietary regimes. Present ICN proposals do not seem to have taken up the battle of

Figure 7. Comparison of state management and forwarding performance in different network scenarios (mean and standard variation)

standing in the wild. In an open Internet, threats are built on the worst scenarios, not on average cases. If we want an information-centric Internet to remain open and reliable, a major redesign of its core architecture appears inevitable.

REFERENCES

[1] B. Ahlgren, C. Dannewitz, C. Imbrenda, D. Kutscher, and B. Ohlman, "A Survey of Information-Centric Networking," *IEEE Communications Magazine*, vol. 50, no. 7, pp. 26–36, July 2012.

[2] M. Gritter and D. R. Cheriton, "An Architecture for Content Routing Support in the Internet," in *Proc. USITS'01*. Berkeley, CA, USA: USENIX Association, 2001, pp. 4–4.

[3] T. Koponen, M. Chawla, B.-G. Chun, A. Ermolinskiy, K. H. Kim, S. Shenker, and I. Stoica, "A Data-Oriented (and beyond) Network Architecture," *SIGCOMM CCR*, vol. 37, no. 4, pp. 181–192, 2007.

[4] L. Zhang, D. Estrin, J. Burke, V. Jacobson, and J. D. Thornton, "Named Data Networking (NDN) Project," PARC, Tech.report ndn-0001, 2010.

[5] V. Jacobson, D. K. Smetters, J. D. Thornton, and M. F. Plass, "Networking Named Content," in *Proc. of 5th ACM CoNEXT*. New York, NY, USA: ACM, Dec. 2009, pp. 1–12.

[6] P. Jokela, A. Zahemszky, C. E. Rothenberg, S. Arianfar, and P. Nikander, "LIPSIN: Line Speed Publish/Subscribe Inter-networking," in *Proc. of ACM SIGCOMM 2009*. NY, USA: ACM, 2009, pp. 195–206.

[7] B. Ahlgren *et al.*, "Second NetInf Architecture Description," 4Ward EU FP7 Project, Tech.report D-6.2 v2.0, 2010.

[8] W. Wong and P. Nikander, "Secure Naming in Information-centric Networks," in *Proc. of Re-Architecting the Internet Workshop (ReARCH '10)*. New York, NY, USA: ACM, 2010, pp. 12:1–12:6.

[9] C. Dannewitz, J. Goliólic, B. Ohlman, and B. Ahlgren, "Secure Naming for a Network of Information," in *Proc. of the IEEE Global Internet Symposium*. Piscataway, NJ, USA: IEEE, 2010.

[10] A. Ghodsi, T. Koponen, J. Rajahalme, P. Sarolahti, and S. Shenker, "Naming in Content-oriented Architectures," in *Proceedings of the ACM SIGCOMM workshop on Information-centric networking*, ser. ICN '11. New York, NY, USA: ACM, 2011, pp. 1–6.

[11] N. Fotiou, G. F. Marias, and G. C. Polyzos, "Publish–Subscribe Internetworking Security Aspects," in *Trustworthy Internet*, N. Blefari-Melazii, G. Bianchi, and L. Salgarelli, Eds. Springer, 2011, pp. 3–15.

[12] A. Ghodsi, S. Shenker, T. Koponen, A. Singla, B. Raghavan, and J. Wilcox, "Information-Centric networking: Seeing the Forest for the Trees," in *Proc. of the 10th ACM HotNets Workshop*, ser. HotNets-X. New York, NY, USA: ACM, 2011.

[13] K. Butler, T. Farley, P. McDaniel, and J. Rexford, "A Survey of BGP Security Issues and Solutions," *Proc. of the IEEE*, vol. 98, no. 1, pp. 100–122, January 2010.

[14] S. Arianfar, P. Nikander, and J. Ott, "On Content-Centric Router Design and Implications," in *Proc. of ReARCH workshop*. New York, NY, USA: ACM, 2010.

[15] D. Perino and M. Varvello, "A Reality Check for Content Centric Networking," in *Proc. of the ACM SIGCOMM WS on Information-centric Networking (ICN '11)*. NY, USA: ACM, 2011, pp. 44–49.

[16] T. Lauinger, "Security & Scalability of Content-Centric Networking," Master's thesis, TU Darmstadt, Darmstadt, Germany, 2010.

[17] Y. Chung, "Distributed Denial of Service is a Scalability Problem," *ACM SIGCOMM CCR*, vol. 42, no. 1, pp. 69–71, 2012.

[18] M. Wählisch, T. C. Schmidt, and M. Vahlenkamp, "Bulk of Interest: Performance Measurement of Content-Centric Routing," in *Proc. of ACM SIGCOMM, Poster Session*. New York: ACM, August 2012, pp. 99–100.

[19] C. Yi, A. Afanasyev, I. Moiseenko, L. Wang, B. Zhang, and L. Zhang, "A Case for Stateful Forwarding Plane," PARC, Tech. Rep. NDN-0002, July 2012.

[20] P. Gasti, G. Tsudik, E. Uzun, and L. Zhang, "DoS and DDoS in Named-Data Networking," ArXiv e-prints, Tech. Rep. 1208.0952, August 2012.

[21] M. Lepinski and S. Kent, "An Infrastructure to Support Secure Internet Routing," IETF, RFC 6480, February 2012.

[22] P. Mohapatra, J. Scudder, D. Ward, R. Bush, and R. Austein, "BGP Prefix Origin Validation," IETF, RFC 6811, January 2013.

[23] A. Li, X. Liu, and X. Yang, "Bootstrapping Accountability in the Internet We Have," in *Proc. of the 8th NSDI*. Berkeley, CA, USA: USENIX Association, 2011.

[24] V. Jacobson, "Congestion Avoidance and Control," *SIGCOMM Comput. Commun. Rev.*, vol. 18, no. 4, pp. 314–329, August 1988.

[25] "PingER. Ping end–to–end reporting," http://www-iepm.slac.stanford.edu/pinger/, 2012.

[26] C. J. Bovy, H. T. Mertodimedjo, G. Hooghiemstra, H. Uijterwaal, and P. V. Mieghem, "Analysis of End to end Delay Measurements in Internet," in *Proc. of the Passive and Active Measurement Workshop-PAM*, March 2002.

[27] S. A. Crosby and D. S. Wallach, "Denial of Service via Algorithmic Complexity Attacks," in *Proc. of USENIX Security Symposium*. Berkeley, CA, USA: USENIX Assoc., 2003, pp. 29–44.

[28] U. Ben-Porat, A. Bremler-Barr, H. Levy, and B. Plattner, "On the Vulnerability of Hardware Hash Tables to Sophisticated Attacks," in *Proc. of IFIP Networking*, ser. LNCS, vol. 7289. Berlin, Heidelberg: Springer–Verlag, 2012, pp. 135–148.

[29] PARC, "The CCNx Homepage," http://www.ccnx.org/, 2012.

[30] C. Yi, A. Afanasyev, L. Wang, B. Zhang, and L. Zhang, "Adaptive Forwarding in Named Data Networking," *SIGCOMM Comput. Commun. Rev.*, vol. 42, no. 3, pp. 62–67, 2012.

[31] M. Wählisch, T. C. Schmidt, and M. Vahlenkamp, "Backscatter from the Data Plane — Threats to Stability and Security in Information-Centric Networking," Open Archive: arXiv.org, Technical Report arXivarXiv:1205.4778, 2012. [Online]. Available: http://arxiv.org/abs/1205.4778

[32] M. Handley, I. Kouvelas, T. Speakman, and L. Vicisano, "Bidirectional Protocol Independent Multicast (BIDIR-PIM)," IETF, RFC 5015, October 2007.

[33] D. R. Simon, S. Agarwal, and D. A. Maltz, "AS-Based Accountability as a Cost-effective DDoS Defense," in *Proc. of Workshop on Hot Topics in Understanding Botnets*. Berkeley, CA, USA: USENIX Association, 2007.

[34] D. G. Andersen, H. Balakrishnan, N. Feamster, T. Koponen, D. Moon, and S. Shenker, "Accountable Internet Protocol (AIP)," in *Proc. of the ACM SIGCOMM*. New York, NY, USA: ACM, 2008, pp. 339–350.

[35] T. Peng, C. Leckie, and K. Ramamohanarao, "Survey of Network-Based Defense Mechanisms Countering the DoS and DDoS Problems," *ACM Comput. Surv.*, vol. 39, no. 1, 2007.