Feasibility of IP-over-ICN

Jaume Benseny¹ · Dmitrij Lagutin¹ · Heikki Hämmäinen¹ · Dirk Trossen²

Published online: 31 July 2019
© The Author(s) 2019

Abstract

The IP-over-ICN strategy intends to establish islands of networks that internally route packets based on Information-Centric Networking (ICN) while maintaining IP-based protocols at the ingress and egress of the network. This strategy aims at benefits from the use of ICN-based routing while maintaining backward compatibility with IP-based services. In the long run, an ICN-based Internet architecture may emerge from the interconnection of these ICN-based islands. We assess the feasibility of this strategy by discussing the willingness of Internet stakeholders to adopt one particular IP-over-ICN implementation based on the Publish-Subscribe Internet Technologies (PURSUIT) for flow-based routing, multicast routing, and service routing. We suggest that the IP-over-PURSUIT solution offers viable mechanisms for IP interoperability and routing scalability as well as potential advantages in comparison to substitutes, including IP-based solutions, such as IPv6; Multiprotocol Level Switching; and hybrid ICN; as well as other IP-over-ICN implementations based on Content-Centric Networking. We indicate that triple play operators and micro-operators have a greater incentive to adopt IP-over-PURSUIT since they can maximize the utilization of the multicast and service routing, respectively. However, we argue that IP-over-PURSUIT requires new exterior inter-stakeholder interfaces for significant operator traffic to be delivered through its new and cost-efficient routing capabilities, thus increasing the likelihood of operator adoption. Finally, we suggest that the advent of an ICN-based Internet architecture might be delayed until Internet stakeholders can trustworthily delegate the delivery of valuable content and services via information-based exchange points.

Keywords IP-over-ICN · PURSUIT · Internet architecture · Interior gateway protocol · Feasibility

1 Introduction

The Internet Protocol (IP) was designed in the 70s, implementing a host-centric communication model. As Internet usage has shifted towards information retrieval, alternative Internet architectures have been proposed imple-

menting data-centric models, including Data-Oriented Network Architecture (DONA) [1], Content-Centric Networking (CCN) [2], The Publish Subscribe Internet Technology (PURSUIT) [3], and Network of Information (NetInf) [4]. However, these architectures have not been deployed since they require the globally synchronized substitution of IP and endanger the market position of dominant stakeholders [5–8].

The feasibility of Internet architectures implementing Information-Centric Networking (ICN) has been analyzed by anticipating stakeholder conflicts or tussles [8]. Tussle analyses have identified conflicts based on trust, economics, security, among others. More importantly, key design aspects remain unanswered, including the management of global content identifiers or the inter-domain routing of multi-source content [9–11]. To address these challenges, two deployment strategies have been devised aiming for gradual ICN deployments that remain IP interoperable, namely ICN-over-IP and IP-over-ICN [12].

This work was supported by the EC H2020 RIFE Project Grant No. 644663 and by the EC H2020 POINT Project Grant No. 643990.

✉ Jaume Benseny
jaume.benseny@aalto.fi

Dmitrij Lagutin
dmitrij.lagutin@aalto.fi

Heikki Hämmäinen
heikki.hammainen@aalto.fi

Dirk Trossen
dirk.trossen@interdigital.com

¹ Communications and Networking Department, Aalto University, Konemiehentie 2, Espoo, Finland

² InterDigital Europe, 64 Great Eastern St, London, UK

Although the ICN-over-IP strategy enables rapid deployment of ICN islands, these islands have IP/ICN overlapping routing functions, which may increase operator costs and avoid protocol competition.¹ In contrast, the IP-over-ICN strategy deploys native ICN networks that remain interoperable with existing IP networks/devices via gateways. Within the ICN network, data-centric routing features, such as multicast delivery; mobility support; fast indirection; in-network computing, can be exploited to handle IP-based protocols [12]. Thus, IP-over-ICN enables operators to increase the efficiency of the underlying transport network while keeping synergies with the IP ecosystem.

This article evaluates the feasibility of IP-over-ICN employing a conceptual framework created to analyze Internet protocols that emerge from consensus-based forums, such as the Internet Engineering Task Force (IETF) [15]. We study stakeholder adoption by (1) analyzing the IP-over-PURSUIT architecture,² (2) comparing stakeholder-specific benefits across alternative value networks, and (3) comparing its technical implementation against substitute solutions, including IP-over-CCN, IPv6, Multiprotocol Level Switching (MPLS), and hybrid-ICN (hICN). The study includes three use cases of significant economic attractiveness, including flow-based routing, multicast routing, and service routing. We address technical and business feasibility challenges by conducting new technical evaluations, referring to simulation results from key publications, suggesting design modifications, and providing evidence from a field trial. Finally, we specifically discuss operator adoption depending on service offering, market conditions, and stakeholder collaboration.

The rest of this article is structured as follows. Section 2 describes the architecture and functions forming the IP-over-PURSUIT routing solution. Section 3 presents the framework for the feasibility analysis of Internet protocols and reports feasibility results for IP-over-PURSUIT as an IGP solution. Section 4 discusses the feasibility of the IP-over-ICN deployment strategy. Finally, Sect. 5 summarizes conclusions.

2 Background on the IP-over-PURSUIT routing solution

2.1 PURSUIT ICN network

In contrast to the destination-based routing of IP, the PURSUIT ICN network implements path-based routing through

¹ Recently, the Hybrid Information-Centric Networking (hICN) has been proposed, inserting CCN into IPv6 and reducing function overlap. However, it may require an inconvenient mapping of entire IPv6 address families to content-named prefixes [13,14].

² We study the IP-over-PURSUIT architecture because it provides wider IP interoperability than IP-over-CCN, allowing the simultaneous translation of IP multicast, HTTP, CoAP, and general IP into PURSUIT.

three ICN functions developed in the FP7 PURSUIT project [3,17]. The rendezvous function (RV) manages the namespace of the information items being published and subscribed to, allowing for matching demand and supply for information. The topology manager function (TM) constructs suitable forwarding identifiers (FIDs) connecting sources and sinks of information. Ultimately, the forwarding function (FWD) delivers the information according to the match performed by the RV following the delivery path indicated by the FID. One efficient implementation of the FWD utilizes path-based rather than flow-based forwarding. For this, the TM labels every link in the network with a specific bit position in a bit-field, similarly to previous Bloom-filter-based approaches [18]. For this, each link within the network is identified by a constant length link identifier (LID). Thus, FIDs are computed by the TM by executing a binary OR operation over the LIDs along the delivery path between the matched publishers and subscribers. The FID is then included in the packet header. Switches along the path will perform a binary AND operation between the FID and the LID for each of the switch's output ports. Thus, forwarding decisions are based on bit-field matching. A positive match means that the FID contains the specific LID, in which case the packet is sent to that port. Multiple matches between the LID and output ports can occur at once, in which case the packet is forwarded to multiple ports, enabling multicast. In this work, we assume that positions in the bit-field are not used for more than one link, preventing false positives.

At the deployment level, the RV and the TM functions might be located in central locations within the network, while the forwarding must be present in all intermediary switches. One significant advantage of the path-based forwarding via bit-field matching is the ability to create any delivery paths, including point-to-multipoint (P2MP) multicast while keeping the required state at the switches constant [19]. Such delivery paths can also be created in an ad-hoc manner by combining existing point-to-point (P2P) paths into a new P2MP path through a binary OR combination of the original ones. However, the number of identifiable links within a network is limited by the FID length (e.g., four-bit identifiers can only store four non-overlapping link names) [20].

2.2 IP-over-PURSUIT architecture

The IP-over-PURSUIT solution follows a gateway-based architecture in which an interior ICN network is bordered by Network Attachment Points (NAPs), as shown in Fig. 1. Thus, NAPs act as IP gateways, translating IP-based protocols into ICN traffic, and vice versa. More importantly, NAPs are purposefully designed to execute multiple applications each of which handles a specific IP-based protocol. These applications, known as protocol handlers, can optimize the

Fig. 1 IP-over-PURSUIT gateway-based architecture [16]

transmission of IP data taking advantage of the ICN functions, including multicast, caching, and fast indirection of higher layer protocols [16,21].

At the forwarding level, standard SDN switches are used to convey the traffic. Since SDN switches support forwarding based on IPv6 addresses, the size of FIDs is exactly the same as the combined size of the IPv6 source and destination addresses (256 bits), FIDs can be embedded in the packet header instead of the IPv6 addresses for forwarding decisions. At the same time, the node LIDs are stored in flow-tables of SDN switches. Support for forwarding based on bit-field matching (Arbitrary Bit Match, ABM) has been added to the OpenFlow 1.3 release [16,19]. Therefore, any OpenFlow 1.3 compatible SDN equipment can be used with the IP-over-PURSUIT solution.

2.3 IP interoperability

The IP-over-PURSUIT solution is IP interoperable thanks to the interplay between the *IP handler* running in NAPs, and the *IP namespace* stored in the RV. This namespace organizes the IP address space according to address prefixes, as shown in Fig. 2. Thus, when an IP packet arrives at an ingress NAP, either a publication or a subscription is added to the scope with the longest prefix-match based on the packet IP destination. In the case of an IP packet traveling from A to B, the A-facing NAP publishes the packet to the address prefix of B in behalf of A. Upon receiving this publication, the RV function matches it with a previous subscription of the B-facing NAP to the address prefix of B. Then, the RV

Fig. 2 IP namespace [22]

triggers the A-facing NAP to send the packet to B through the B-facing NAP. For this, the A-facing NAP looks up its local database for the appropriate FID to reach the B-facing NAP. The first packet sent between A and B includes an implicit subscription to the prefix address of A in behalf of B, thus ensuring that the B-facing NAP immediately sends responses produced by B back to the A-facing NAP, without consulting the RV. Since IP gateways keep a partial copy of the *IP namespace*, the RV is not consulted for each path request. This procedure is only required for the first data packet exchange between two IP endpoints. Subsequent data packets can be sent using the allocated FIDs. To facilitate forwarding inwards and outwards the ICN network, the prefixes are split into user-facing NAPs (I_i), and Internet-facing NAPs (O_i), respectively. An additional namespace is required to support subnet-based addressing, enabling NAPs to subscribe to entire subnets [16].

2.4 Handling specific IP-based protocols

The IP-over-PURSUIT solution currently provides protocol handlers for IP multicast, HTTP, and CoAP. These handlers optimize the transmission of IP data in the ICN network by taking advantage of the ICN functions, including multicast, caching, and fast indirection of higher layer protocols. In this section, we describe how these protocols utilize ICN-based operations and namespaces to handle IP multicast and HTTP.

For IP multicast, NAPs execute the *IGMP handler* enabling the addition of and subscription to IP multicast addresses in the *IGMP namespace*. In addition, NAPs also execute the *IP multicast handler* to compute P2MP delivery paths with as many destinations as client-facing NAPs belong to the multicast group. This computation, known as ad-hoc multipath creation, can be autonomously executed by the server-facing NAP by applying a logical OR operation over the individual P2P FIDs. When the data is finally received at the client-facing NAPs, they generate as many IP multicast packets as served IP clients [16].

For HTTP, the *HTTP handler* in client-facing NAPs can aggregate quasi-synchronous HTTP requests, which are directed to the same Uniform Resource Locator (URL) but originated from different local IP clients, into a single HTTP request that is transported over the ICN network. In turn, the resulting HTTP response is delivered back to the client-facing NAPs via a P2MP delivery path. This P2MP path is again created by the server-facing NAP via ad-hoc multipath creation, i.e., the binary OR of the incoming request path information. Finally, the client-facing NAPs generate as many individual unicast HTTP responses as requesting IP clients. We refer to this functionality as the HTTP coincidental multicast. For this functionality to work, *HTTP handlers* register the relevant Fully Qualified Domain Name (FQDN) of a service in the *HTTP namespace* by assigning Content Identifiers (CID), as shown in Fig. 3. Reverse-CIDs (rCIDs) are also created for client-facing NAPs to generate HTTP responses that follow the original unicast semantic. For this, rCIDs include not only the URL but other HTTP requests parameters [22].

The performance of HTTP-based services can be further improved through the notion of surrogate services. These services are IP-level endpoints exposing FQDNs that are already registered in the local NAP, resulting in several subscriptions to the same CID. Thus, for a given request, the RV can choose between one or more service endpoints which will be included in the final path computation by the TM. Furthermore, this path computation can be constrained by policies, including FQDN-based policies and delay-based policies. In the latter case, the TM would need to realize monitoring capabilities to weigh the path computation algorithm with delay constraints.

Fig. 3 HTTP namespace [22]

2.5 Mobility management and service indirection

The gateway-based architecture of IP-over-PURSUIT facilitates both the handover of User Equipment (UE) between NAPs and the switchover of HTTP connections between surrogate services.

When a UE detects a NAP providing a better signal reception than the NAP to which it is currently attached, it starts a handover from the old NAP to the new NAP. First, the old NAP directly informs of the de-attachment to all other NAPs which had established packet exchange with the UE, i.e., it directly requests the un-subscription to the address prefix of the UE. After the UE is physically attached to the new NAP, the new NAP publishes incoming packets to the RV according to the address prefix of their IP destination. At this point, the RV re-matches the previously existing publications and subscriptions for the UE, thus triggering the new NAP to re-establish communications. On the other end of the packet exchange, other NAPs realize there is a new NAP acting as the IP gateway for the UE since the new NAP always includes an implicit subscription within the first packet after an RV update. Note that FIDs do not need to be re-calculated by the TM, IP gateways only need to switch the old FID by the new one, which is either provided by the RV or looked up from a local copy of the *IP namespace*. As a pre-condition for mobility, NAPs need to execute the *DHCP handler* assigning the same IP address to UE when they re-attach during a handover. For this, the unequivocal assignment of IP addresses to UE is centrally stored in the *DHCP namespace* of the RV. Summarizing, in IP-over-PURSUIT, handovers are based on the substitution of NAPs as publishers and/or subscribers for the IP addresses of UE thanks to the decoupling of IP addresses from the UE location. More sophisticated handover procedures have been specified for IP-over-PURSUIT establishing delivery paths with multiple destinations, simultaneously reaching those NAPs with the highest probability of re-attachment. For this, a handover handler and namespace are required [16,23].

As mentioned in the previous section, surrogate services enable the RV to select the most suitable source of content for HTTP requests. Therefore, they allow for the dynamic switchover from one HTTP surrogate service instance to another. For this, the TM invalidates forwarding information upon new registrations or failover indications, leading to a new path computation at the next service request at those NAPs that serve initiating clients. We refer to this functionality as HTTP service indirection [24]. Hence, IP-over-PURSUIT can sequentially execute a handover followed by the switchover of HTTP connections, thus connecting the re-attached UE to the surrogate server that offers the best service quality [23]. The switchover is even possible during session runtime since details on individual HTTP requests are centrally stored in the *HTTP namespace*. For this, the RV needs to update state information across namespaces, given that the handover and switchover based on the *IP namespace* and the *HTTP namespace*, respectively. In case no viable surrogate server exists after a handover (or after a new attachment), the availability of emerging surrogate services can be directly notified to client-facing NAPs. Although the service indirection functionality works for stateless HTTP applications, surrogate services need to synchronize state information for applications with user context. Application-level synchronization is out of the scope of this study.

2.6 Routing scalability

The routing scalability of the IP-over-PURSUIT solution is limited by the scalability of its network functions. For scaling the FWD, the IP-over-PURSUIT solution adopts a zoning mechanism that divides the ICN network into zones with a maximum number of links [20]. Due to the integration with SDN switches, the FID length is constrained by the size of the IPv6 header, thus limiting the number of links that can be included in a FID. In this work, we assume that positions in the FID are only used to identify one link per zone, thus preventing false positives. The proposed zoning mechanism requires the IPv6 packet header to include not only a zone-specific FID but also a zone identifier (ZID). In addition, the packet payload piggybacks FIDs of future traversed zones. Therefore, inter-zone routers are also required in the IP-over-PURSUIT architecture to update the packet header with the correct (ZID + zone-specific FID) according to the next zone the packet is forwarded to.

For scaling the RV and TM, virtual instances of these functions are replicated throughout the network using Network Function Virtualization (NFV). Thus, the workload generated by NAPs is adequately balanced between RV/TM replicas. Hereafter, we refer to this scalability mechanism as the NFV-based RV/TM. The state information in RV/TM replicas can be updated through a distributed database mechanism. For the RV, the state information includes the

namespaces, the active subscriptions, and the publications (the latter are only used in reading mode). For the TM, this information includes the state of links which is only updated after an inventory change, i.e., when a forwarding node enters or leaves the topology. Regarding the workload generated by NAPs, path computation requests only occur for initial service requests, i.e., only the first time that a specific URL is being requested at the client-facing NAP. If any other future request is issued at the same NAP (from any of its locally attached clients), no path computation request is needed and therefore neither RV nor TM is used. Server-facing NAPs do not generate any load on the RV/TM since the return path for packet equals that of the forward path, unless weighted policies are used that differentiate forward from return paths.

3 Feasibility analysis

3.1 Framework for the feasibility analysis of Internet protocols

This article employs a conceptual framework created to analyze the feasibility of new Internet protocols that are developed in consensus-based forums, such as the IETF [15]. The framework helps to identify the most attractive use cases, to compare the protocol implementation against its substitutes, and to assess the incentives of stakeholders to participate in its deployment. The framework establishes a sequential and iterative process, as shown in Fig. 4. Each step in the process contains clearly defined objectives which are listed in Table 1.

The framework has been used to analyze protocols with a moderate level of uncertainty regarding their adoption by stakeholders. For example, it has been used on CoAP [25], Multipath-TCP [9], and Host Identity Protocol [26].

3.2 Use case analysis

We identify three use cases in which Internet stakeholders can obtain significant benefits from the IP-over-PURSUIT solution.

The Multiprotocol Label Switching (MPLS) has been widely adopted by operators to implement traffic engineering and VPN management via Label Switched Paths (LSPs). Recently, flow-based solutions have been developed enabled by IPv6 and SDN to provide similar routing features, albeit with their specific challenges in terms of scalability and costs [27]. Given the significant market opportunity derived from MPLS phase out, we decide to study the ability of IP-over-PURSUIT to perform flow-based routing and traffic engineering through path-based routing and associated benefits regarding SDN integration.

Fig. 4 Framework for the feasibility analysis of Internet protocols [15]

Table 1 Steps in the feasibility analysis

	Analysis steps	Objective
1	Use case analysis	To identify use cases in which the IP-over-PURSUIT solution can provide potential benefits to Internet stakeholders
2	Technical architecture	To identify the architecture elements that serve the use case needs To identify deployment actions for the identified architecture elements
3	Value network analysis	To identify the relevant stakeholders for the use case To assign deployment actions and business roles to stakeholders To identify alternative value networks for the use case
4	Deployment environmental analysis	To compare the protocol implementation against substitute solutions and identify advantages and drawbacks for each use case To identify external factors affecting the feasibility of the protocol
5	Feasibility analysis	To identify and evaluate the key technical and business factors that can dispute the potential benefits of IP-over-PURSUIT
6	Solution analysis	To identify positive net benefits for all relevant stakeholders as well as significant deployment challenges of the protocol for each use case

The IP multicast protocol has been deployed by operators offering linear TV services. However, consumer surveys report consistent media consumption growth through Over The Top (OTT) services including Video-on-Demand (VoD). For instance, the US market is experiencing cable-to-VoD substitution driven by lower prices of VoD services which reached 50% market penetration in 2016 [28,29]. Given that OTT services are provided by Content Delivery Networks (CDNs) via best-effort IP, IPTV operators are not able to exploit the IPTV infrastructure nor differentiate from connectivity oriented operators, thus losing content-related revenues. As a reaction, European operators following a service convergence strategy are acquiring media assets.³ Hence, we decide to study the ability of the IP-over-PURSUIT solution

to implement scalable IP multicast and improve the distribution of HTTP-transported media content via multicast routing.

New mobile edge services with low-latency requirements, such as vehicle-to-network services, are expected to emerge driven by cost-efficiency gains [31]. To this end, server-side application logic and data need to move closer to end users and follow them across distributed computing resources (e.g., across an edge cloud infrastructure). In this context, future routing solutions are expected to support the deployment of virtual service instances on the edge of the network and seam-

Footnote 3 continued

operator also acquired distribution rights over the Finnish ice hockey league. Similarly, the Spanish operator Telefonica acquired rights over the UEFA Champions League and UEFA Europa League for its residential market.

³ The Swedish operator Telia bought the Bonnier Broadcasting company, thus taking control over 3 linear TV channels [30]. The same

Table 2 Identified use cases

Use cases	Purpose
Flow-based routing	To evaluate the ability of IP-over-PURSUIT to perform flow-based routing employing its path-based routing functionality
Multicast routing	To evaluate the improved implementation of IP multicast and the improved delivery of HTTP-transported media via HTTP coincidental multicast
Service routing	To evaluate the ability of IP-over-PURSUIT to enable mobile edge services via UE handover and HTTP service indirection

lessly switch connections between instances. Therefore, we decide to study the ability of the IP-over-PURSUIT solution to enable mobile edge services via service routing, given its ability to handover UE between NAPs and reroute services via HTTP service indirection.

Results from the use case analysis are summarized in Table 2.

3.3 Technical architecture analysis

IP-over-PURSUIT can serve the needs of the three use cases by adopting different architecture configurations, as shown in Table 3. In each configuration, the ICN functions provide different functionalities. For example, as the routing complexity of use cases increases, NAPs need to gradually execute additional protocol handlers.

3.3.1 FWD function

The current IP-over-PURSUIT design keeps a minimal state in routers and requires low controller-to-switch SDN interaction. A key feature of the IP-over-PURSUIT architecture

is that the FWD remains as a bit-field matching operation across use cases. Therefore, flow-tables in SDN switches only need to account for the switch own links, keeping the state at its minimum and independent from traffic demand. Furthermore, flow rules can be inserted proactively by SDN controllers during the network bootstrapping, removing flow setup latency caused by reactive flow insertion. As a result, controller-to-switch interaction is very low since link-state-updates are only required when the topology changes.

The current IP-over-PURSUIT design states that a zoning mechanism is required to ensure FWD scalability. However, no guidelines exist on how zones should be created, i.e., providing a maximum number of links per zone. Therefore, the zone-based routing scalability mechanism remains as an open feasibility challenge. Moreover, inter-zone routers need to store network state mapping links to zones, thus breaking the IP-over-PURSUIT promise of stateless routing.

3.3.2 RV/TM functions

The NFV-based RV/TM guarantees that the RV and the TM act as a logically centralized function, while their execution can be distributed, as described in Sect. 2.6. From the TM's viewpoint, link-state information needs to be logically centralized maintaining a network-wide scope, thus enabling the computation of FIDs connecting the opposite extremes of the network. From the RV's viewpoint, the logically centralized management of server-side subscriptions is also necessary to guarantee network-wide content availability. Based on these observations, the current design of the NFV-enabled RV/TM resembles a distributed flat SDN control plane that additionally manages pub/sub information and resolves content requests through the user plane [27]. Although the centralization of link-state and pub/sub information can significantly simplify network management, it also creates a single point of failure if the network function replication is not conducted correctly.

Table 3 Architecture configurations for the use cases

Use cases	In NAPs	In the RV	In the TM	In the FWD
Flow-based routing (Architecture config. 1)	<i>IP handler</i>	<i>IP namespace</i>	Link-state data Path computation	Bit-field match
Multicast routing (Architecture config. 2)	<i>IP handler</i> <i>IGMP handler</i> <i>IP multicast handler</i> <i>HTTP handler</i>	<i>IP namespace</i> <i>IGMP namespace</i> <i>HTTP namespace</i>	Link-state data Path computation	Bit-field match
Service routing (Architecture config. 3)	<i>IP handler</i> <i>HTTP handler</i> <i>DHCP handler</i> <i>Handover handler</i>	<i>IP namespace</i> <i>HTTP namespace</i> <i>DHCP namespace</i> <i>Handover namespace</i>	Link-state data Path computation Monitor FQDN connections	Bit-field match

Although the current IP-over-PURSUIT design establishes that the workload generated by NAPs is adequately balanced among RV/TM replicas, no guidelines exist indicating how zones should be created, i.e., providing a maximum number of NAPs or aggregated active subscriptions per replica.

3.3.3 Function deployment

Since the majority of elements in the IP-over-PURSUIT architecture are software-based, they can be replicated and deployed on the strategically located equipment in the network. While RV/TM replicas, SDN controllers, and inter-zone routers can be deployed in central locations; NAPs should be deployed on the network edge. The deployment of IP-over-PURSUIT can be facilitated by Infrastructure as a Service (IaaS) solutions, given that virtual resources can be dynamically assigned to network functions. IaaS solutions can also facilitate the deployment of third-party services in the network, i.e., HTTP surrogate services. In this context, IP-over-PURSUIT can be deployed either following the traditional IGP approach, i.e., as a single network routing solution, or it can also be deployed in a part of the network.

3.4 Value network analysis

We assess the willingness of Internet stakeholders to adopt IP-over-PURSUIT by comparing alternative value networks, aiming to identify those networks that better balance benefits among stakeholders [32,33]. We generate value networks by assigning architectural elements and related business roles to stakeholders. Finally, actor-specific benefits are derived from interfaces and contracts that connect technical elements and business actors, respectively.

Use case 1: flow-based routing Operators that provide IP connectivity services can deploy IP-over-PURSUIT as an IGP solution to perform flow-based routing and, they can keep control over all elements in the architecture. In other words, operators do not need to concede the control of any architectural element to stakeholders. Therefore, operators do not need to expose any additional technical interface apart from the existing inter-domain interfaces, which remain IP interoperable. Hence, connectivity operators can adopt IP-over-PURSUIT in its architecture configuration 1 without changing current value networks for the provision of IP connectivity services.

Use case 2: multicast routing Connectivity operators that also provide content-based services, e.g., triple play operators, can also deploy the IP-over-PURSUIT solution as an IGP and keep control over all architectural elements. However, the use case benefits obtained by the operator, i.e., the release of link capacity due to multicast, depend on the frac-

tion of traffic that can be delivered via IP multicast and HTTP coincidental multicast. In both cases, operators need to sign business contracts with content owners to deliver a substantial amount of traffic through these protocols. Note that the multicast of content that is not protected by copyright might not release enough capacity to justify adoption.

Triple play operators (and IPTV operators) already perform contract-based content provision via the IP multicast protocol. Further, they have complete control over the delivery process since they typically manage both the content servers and the network. As a result, they can enjoy from an improved implementation of IP multicast without any additional inter-stakeholder technical interface or business contract.

However, triple play operators do not control the provision of OTT content, i.e., they distribute OTT content as any other HTTPS communication because of the payload encryption. To address this technical issue, business contracts should be signed to allow the exchange of certificates, enabling operators to publish OTT content on the RV.⁴ However, OTTs may lack incentives to sign these contracts since their content is already delivered successfully by CDNs and they do not obtain an immediate benefit from IP-over-PURSUIT. More precisely, the release of operator link capacity via HTTP coincidental multicast does not benefit OTTs. Moreover, the current IP-over-PURSUIT design does not yet provide to OTTs the same level of information on consumer received quality or consumer behaviour in comparison to existing CDN solutions. In the worst case scenario, some implementations could deliberately prevent OTTs from obtaining this information. Since in HTTP coincidental multicast quasi-synchronous HTTP requests are combined into one single request, the number of requesting users could be hidden from content servers, limiting personalized advertisements, etc. Hence, for operators to fully enjoy the benefits of HTTP coincidental multicast, new inter-stakeholder technical and business interfaces are required for certificate exchange and consumer information sharing, at least.

Use case 3: service routing use case Although mobile edge services do not yet exist, we generate plausible value networks by adapting business roles and stakeholders that exist in the provision of cloud services. We envision that value networks for the provision of mobile edge services will include a business role responsible for the provision of surrogate services and, they will also include a stakeholder that provides edge applications.

Operators providing mobility services could deploy IP-over-PURSUIT as an IGP, albeit they could only provide mobility via Wi-Fi, given that a pure IP-over-PURSUIT system is not currently able to attach UE via 3GPP protocols.

⁴ Technically, certificates are needed to generate rCIDs and store them in the *HTTP namespace*.

Table 4 Value network analysis results

Use cases	Leading actor (deployment type)	Required stakeholders	Disputed business roles	Business contracts (Technical interfaces)
Flow-based routing	Connectivity operator (IGP)	None	None	None
Multicast routing	Triple play operator (IGP)	OTT provider	Provision of OTT content	Content distribution contract (RV access)
Service routing	Micro-operator (IGP)	None	None	None
	MNO (Edge cloud)	Edge application provider	Provision of surrogate services	Service surrogacy contract (RV access) (3GPP signaling)

However, Mobile Network Operators (MNOs) might be interested in deploying IP-over-PURSUIT on subsystems of the mobile network.⁵ MNOs could deploy IP-over-PURSUIT on the edge cloud since specifications of 5th generation mobile systems (5G) allows local routing and traffic steering [34], thus possibly enabling HTTP service indirection for edge applications requiring mobility. However, technical interfaces integrating 3GPP signaling protocols with the RV/TM do not exist yet. Nevertheless, some equipment vendors are interested in this integration, as shown in recent IETF drafts [35].

Although the IGP deployment option might not be attractive to MNOs, it can work for micro-operators serving a limited coverage area, e.g., a campus or a city area. In this micro-operator-driven value network, the micro-operator is likely to control the provision of surrogate services, given that edge application providers have little incentive to serve such a limited coverage area and, they most likely are subcontracted by the micro-operator [36]. As a result, the publication of surrogate services to the RV is greatly simplified.

When evaluating the assignment of business roles, it is not clear which stakeholder will take control over the provision of surrogate services. While MNOs can impose their signaling protocols within the access network, edge application providers can leverage their pre-existing business relationships with end users, forcing some remote control over the surrogates (i.e., large service providers, such as Google). Therefore, business contracts for service surrogacy need to be signed between MNOs and edge application providers, thus facing similar challenges than use case 2. We conclude that the adoption of IP-over-PURSUIT to serve the use case 3 is more likely in a micro-operator-driven value network than in an MNO-driven value network.

In summary, although the operator is solely responsible for the deployment of IP-over-PURSUIT across all use cases, benefits from multicast and service routing are only realized

in full when technical interfaces and business contracts are established, as summarized in Table 4.

3.5 Deployment environment analysis

The willingness of Internet stakeholders to adopt the IP-over-PURSUIT solution depends on the advantages and drawbacks provided by this solution in contrast to substitutes. Therefore, we compare the implementation of substitute solutions against the one of IP-over-PURSUIT for the three use cases under study. Substitute solutions include IP-based solutions, such as IPv6; MPLS; and hICN as well as other IP-over-ICN implementations, such as IP-over-CCN.

Since both hICN and IP-over-CCN rely on CCN, we briefly describe this ICN architecture. In a CCN network, users request content items by broadcasting *Interest* messages containing the name of the content, which follows a hierarchical and aggregate-able structure. In fact, content names have a key role in the transport of *Interest* messages since CCN routers forward these messages according to a name-based prefix match. Similarly to IP forwarding, routers perform a table lookup, matching the content name in packets against the Forward Information Base (FIB), in which content sources (in name format) are associated to one or more outbound interfaces. For each *Interest* message that is forwarded, routers update the Pending Interest Table (PIT), thus establishing a reverse path. When *Interest* messages are finally received by content sources, content items are sent to requesters encapsulated in *Data* messages. Along the data path, routers delete *Interest* entries from the PIT, enforcing the principle that one *Interest* messages retrieves at most one *Data* packet. This forwarding system has been called soft stateful [2].

3.5.1 Advantages and drawbacks compared to IP-based solutions

Use case 1: flow-based routing MPLS, IPv6, and, hICN allow flow-based routing. In MPLS, IP packets are assigned to Forward Equivalence Classes (FECs) by ingress routers

⁵ IP-over-PURSUIT could implement the control plane of 5th generation mobile systems (5G) since surrogate services could implement a service-based architecture. However, this implementation is limited to user plane routing thus outside the scope of use case 3.

according to the longest prefix match between the packet IP destination and the IP addresses in the routing table. Then, MPLS packets are routed according to a new header known as label [37]. In IPv6, packet headers already include a dedicated bit-field for flow assignment. Since hICN reuses IPv6 routing protocols, an hICN network can perform flow-based routing like any other IPv6 network [14]. In comparison, the IP-over-PURSUIT solution also allows flow-based routing by assigning delivery paths to packets based on the longest address prefix match. Furthermore, the match is done against the *IP namespace* of the RV in a centralized manner, facilitating flow management. In addition to flow-based routing, IP-over-PURSUIT allows path-based routing through which new/updated delivery paths can be included in packet headers without waiting for flow rule updates to disseminate through the network.

For MPLS to forward packets along Label Switched Paths (LSPs), routers need to maintain tables matching labels against next-hops, known as Incoming Label Maps (ILMs). Moreover, routers also need to swap packet-labels on a hop-by-hop basis [19,37]. On top of these, MPLS routers typically execute link-based routing protocols such as OSPF to monitor the state of links and IP routing tables. Consequently, the required memory and processing power within MPLS routers increase with the number of links [38]. In contrast, SDN switches in IP-over-PURSUIT are simple forwarders since link-state information is managed by the RV/TM [16,39]. Furthermore, these switches require fewer resources than in the majority of flow-based IPv6/hICN solutions, given that flow tables only need to account for the switch own links and routing decisions are based on bit-field matching [19]. From an SDN viewpoint, this implies that controller-to-switch interaction is limited to topology changes and the cost of integrating SDN switches is constant regardless of the conveyed traffic.

For the control plane to scale up, MPLS networks are divided into levels, thus reducing the size of ILMs in individual routers [37]. Similarly, in IP-over-PURSUIT the scalability of the FWD requires zones due to the limited length of FIDs. While MPLS packets can traverse several levels by stacking multiple labels in the same packet, PURSUIT packets carry as many FIDs as traversed zones. Thus, MPLS routers in-between levels behave similarly than IP-over-PURSUIT inter-zone routers since both swap labels and FIDs, respectively [20,38]. As with other SDN-based solutions, IP-over-PURSUIT also needs to divide the network to balance the load between SDN-controllers, albeit IP-over-PURSUIT enjoys a lower switch-to-controller interaction. In contrast to other NFV-based solutions, the NFV-based RV/TM needs to additionally manage pub/sub information. Nevertheless, in the use case 1, pub/sub information is limited to the subscriptions on IP prefixes and subnets.

MPLS is widely used in IP networks to facilitate traffic engineering via LSP tunnels. In contrast, the IP-over-PURSUIT solution provides comparable, or even superior features since network-wide information on link-state is centralized in the RV/TM. This might not be the case for MPLS since LSPs are established between ingress and egress routers with partial network information [38]. Other flow-based IPv6/hICN solutions that implement control via SDN can also provide similar traffic engineering capabilities.

IP-over-PURSUIT advantages and drawbacks:

- + Forwarding decisions are more efficient and require fewer resources than substitutes.
- + Comparable or even superior traffic engineering capabilities than substitutes.
- + Unlike flow-based IPv6 solutions, SDN switch-controller interaction is low.
- Substitutes based on MPLS/IPv6 are market-ready.

Use case 2: multicast routing The IP multicast protocol adds a tree-building protocol (e.g., PIM) and a rendezvous function (e.g., IGMP) on top of IP, requiring individual routers to maintain per-flow state. Conversely, IP multicast can be provided in an IP-over-PURSUIT network while maintaining routers almost stateless, i.e., except for inter-zone routers [19]. This advantage is maintained against hICN and MPLS. hICN relies on IPv6 to perform multicast and Data packets are routed using IPv6 multicast group addresses instead of name prefixes [14]. In MPLS, additional memory and processing are required in routers, given that P2MPLSPs are established by an additional protocol named Multipoint-LDP [40].

HTTP has been adopted as the main delivery protocol to transfer video chunks from content servers to end users. In IP-over-PURSUIT, HTTP can enjoy coincidental multicast thus responses to the same chunk request can be bundled into a single multicast response over the ICN network. Since gateways restore unicast responses towards clients, traffic load can be reduced in proportion to the number of concurrent viewers.⁶

New multicast solutions are being developed in the IETF Bit Index Explicit Replication (BIER) group. To be precise, BIER encapsulates IP packets with an additional header indicating multiple egress routers through which the packet should exit the network [41]. Although the BIER solution might be able to eliminate the per-flow state in routers and explicit tree-building protocols, it still acts as an extension to the IP architecture. Therefore, routers in the BIER network

⁶ Concurrency is defined within an interval related to the chunk length, often several seconds, rather than requiring strict concurrency in the millisecond range.

might need to manage BIER-specific forwarding tables and share information with link-based routing protocols, such as OSPF. Unless BIER networks adopt a gateway-based architecture, they cannot apply the multicast to IP-based protocols, e.g., HTTP coincidental multicast, nor enforce routing policies, e.g., FQDN-based or delay-based policies. Such addition of gateways for multicast delivery of HTTP responses has been proposed in [42].

IP-over-PURSUIT advantages and drawbacks:

- + IP multicast can be implemented without storing per-flow state in routers.
- + Routing policies can be established based on FQDN and delays.
- + Unlike IPv6, the distribution of HTTP data can be optimized via HTTP coincidental multicast.

Use case 3: service routing To enable IP mobility, current 3GPP-based mobile solutions channel UE-generated traffic towards the core network via GPRS Tunneling Protocol (GTP) tunnels, impeding access to edge services. To address this issue, the 3GPP specification for 5G mobile systems enable local routing and traffic steering [34]. Thus, the 5G core network will support UE access to the so-called “local area data network” in which edge applications will be deployed. Service continuity will be guaranteed for certain mobility scenarios assigning multiple anchor points to a single session, i.e., one anchor point for each data network, including the telco cloud and the edge cloud [43]. In contrast to 3GPP-based solutions, operators deploying IP-over-PURSUIT can enjoy anchorless mobility, allowing delivery paths to follow the optimal route and avoiding network bottlenecks. Furthermore, IP-over-PURSUIT enables the continuity of HTTP sessions after a handover via fast service indirection [23].

Furthermore, the availability of emerging surrogate services can be directly notified to client-facing NAPs, avoiding DNS propagation delays across caches on access points and clients. In this context, the IP-over-PURSUIT solution aims to remain compatible with the current specification MEC [44]. In fact, MEC edge applications are managed through a pub/sub system like IP-over-PURSUIT service routing [45]. If we extend our comparison to include hICN, this routing technology also allows anchorless mobility since mobile UE can redirect in-flight Data packets to its new location by re-sending *Interest* messages through its new attachment point, thus creating a new reverse path. For this, hICN does not require rendezvous-based rematch of pub/sub information. Further, hICN also allows service indirection since routers can associate different sources of the same content to different outbound interfaces of

routers. However, switchovers from old to new emerging content sources might be slow due to packet propagation delays.

IP-over-PURSUIT advantages and drawbacks:

- + Edge content is made directly available to NAPs avoiding DNS/CCN propagation delays.
- + Unlike IPv6, connections can be seamlessly switched between HTTP surrogate services.
- + Unlike IPv6, anchorless mobility and decoupling of IP addresses and UE location.
- Unlike hICN, rendezvous-based pub/sub matching is required for mobility.

3.5.2 Advantages and drawbacks compared to IP-over-CCN implementations

Use case 1: flow-based routing An IETF Internet draft exists describing how IP can be tunneled over CCN [14]. According to this draft, IP gateways need to become content sources and disseminate their associated content name, i.e., the IP address space or prefix they serve, through the CCN network. Therefore, IP traffic generated by UE can be encapsulated in *Interest* messages that have the egress gateway as the content name. After the egress gateway receives the IP response from the Internet, it forwards this back encapsulated in *Data* messages. To ensure that a reverse path exists for as many *Data* messages as the IP response may require, UE may need to send additional *Interest* messages, thus keeping enough active entries in routers’ PIT. At the same time, IP gateways have to maintain an ordered queue of content names that have been requested by each client, namely the Client Interest Table (CIT). Thus, UE and IP gateways collaborate by adjusting the number of in-flight messages according to PIT entry lifetimes, in what may resemble TCP window advertisements. This IP tunneling solution creates pub/sub state in IP gateways since content requested by clients is stored in CITs. While this approach may create a dispersed pub/sub state that is difficult to manage, IP-over-PURSUIT provides a rendezvous-based solution with a scalable design. Further, it does not dedicate gateway resources to artificially maintain reverse paths.

None of the existing IP-over-CCN implementations define how flow-based routing could be performed. Intuitively, it is clear that IP packets that enter a CCN network can be aggregated to become a single flow by encapsulating them in *Interest/Data* packets the content name of which includes the same egress gateway. For this, ingress routers need to assign the content name of egress gateways to *Interest/Data* packets according to the destination address of IP packets,

similarly to MPLS. Based on this assumption, we suggest that IP-over-PURSUIT enjoys, at least, the same advantages compared to IP-over-CCN than to IP-based solutions. First, IP-over-PURSUIT implements more efficient routing decisions using significantly fewer resources than CCN routers, given its stateless forwarding. Second, flow management can be centrally controlled through the *IP namespace* in the RV. Third, IP-over-PURSUIT can potentially perform traffic engineering features via the TM.

IP-over-PURSUIT advantages and drawbacks:

- + Forwarding decisions are more efficient and require fewer resources than substitutes.
- + Pub/sub state is centrally managed via a rendezvous
- + Reverse paths do not need to be artificially maintained.
- + More mature specification of the generic IP handling application

Use case 2: multicast routing Two implementations exist describing how HTTP could be transported over CCN, improving the delivery of heavy content. In the first implementation [46], CDN caches are substituted by CCN routers equipped with Content Storage. These routers need CCN-to-HTTP translation to request content from higher level caches (or other HTTP sources) and, they also need HTTP-to-CCN translation to understand client requests. The main technical difficulty is found in the CCN-to-HTTP translation since HTTP gateways need to generate CCN manifests, which are included in *Data* messages providing a list of content objects and associated signatures. Since these manifests cannot be constructed until the full HTTP response is received, CCN-based communication is stalled. Although multicast routing is mentioned in [46], no detailed implementation is described. Moreover, HTTP gateways are placed immediately after content sources, shortening potentially shared data paths and minimizing multicast gain. The second HTTP-over-CCN implementation adopts a generic architecture positioning HTTP gateways at the border of any CCN network [47]. These gateways translate GET and POST methods into a sequence of Interest/Data packet flows. In contrast to the first implementation, IP responses are split into small chunks which are immediately exchanged as independent data objects. Although this second implementation could potentially implement a coincidental multicast mechanism by aggregating HTTP requests in ingress gateways, this is not described in the specifications document. Moreover, none of the two cited implementations describe how IP gateways keep track of HTTP requests when sent by different UE.

IP-over-PURSUIT advantages and drawbacks:

- + A more mature specification of the HTTP handler.
- + HTTP handler implements HTTP coincidental multicast.

Use case 3: service routing The authors have not found any scientific publications or technical specifications that define how the indirection of IP-based services can be performed over CCN. Nevertheless, if a new source of already existing content would be made available to the CCN network by an IP gateway, this information would need to propagate through CCN routers until other IP gateways become aware. In contrast, IP-over-PURSUIT allows the direct notification of those IP gateways currently serving that content item.

IP-over-PURSUIT advantages and drawbacks:

- + Edge content is made directly available to NAPs avoiding CCN propagation delays.
- Rendezvous-based pub/sub matching is required for mobility.

3.6 Feasibility analysis

This section identifies and evaluates the key technical and business factors that can dispute the potential benefits of IP-over-PURSUIT and potentially discourage stakeholder adoption. Table 5 presents the feasibility challenges identified in previous sections and evaluates the ability of IP-over-PURSUIT to address them for each use case. Based on the evaluation results, we further study the challenges with an unknown feasibility outcome and we suggest possible solutions for the challenges that dispute the feasibility.

3.6.1 Challenge 2: routing scalability

The routing scalability of the IP-over-PURSUIT solution is constrained by the scalability of its network functions. As described in Sect. 2.6, while the FWD is supported by the zoning mechanism, the RV and the TM are supported by the NFV-based RV/TM. Both mechanisms divide the network into zones and assign resources to the management of these zones. Since the current specification of the IP-over-PURSUIT solution does not provide guidelines on how the network should be divided, we decide to evaluate the mechanism that addresses the most challenging scalability bottleneck. The scalability of the FWD is the most critical since it equally affects all three use cases under study. Note that the RV/TM might not present an immediate scalability challenge for use case 1, given that the RV only needs to store the *IP namespace*. Although the scalability of rendezvous

Table 5 Feasibility challenges

	Feasibility challenge	Flow-based routing	Multicast routing	Service routing	Evaluated in
1	Can the architecture design serve the use case needs?	✓	✓	✓	Technical architecture analysis (see Table 3)
2	Can the scalability mechanisms enable routing for nation-sized topologies?	?	?	— ^a	Technical architecture analysis
3	Can the architecture design provide technical advantages in comparison to substitute solutions?	✓	✓	✓	Deployment environment analysis
4	Can the architecture design provide enough benefits for OTT video players to sign a content distribution contract?	—	× ^b	—	Value network analysis (see Table 4)
5	Can the architecture design provide enough benefits for edge application providers to sign a service surrogacy contract?	—	—	× ^c	Value network analysis
6	Is the implemented solution IP interoperable?	?	?	?	Value network analysis

✓ = yes; ? = not known yet; × = disputed; — = the use case is not affected

^aService routing is not affected by the challenge 2 since adoption is limited to micro-operators or MNOs on the edge cloud

^bThe release of operator link capacity via HTTP coincidental multicast does not immediately benefit OTTs

^cEdge application providers cannot remotely monitor service quality or control switchover between surrogates

functions implementing an Internet-scale model have been studied in the past, results do not apply here [48].

The technical feasibility of the FWD has been extensively studied via simulations as an enabler for stateless multicast solutions [17,18]. More recently, it has been demonstrated that the FWD can be implemented with existing SDN switches with less state than existing unicast switching solutions [19]. In addition, to guarantee the FWD scalability, a zoning mechanism has been proposed implementing the Extensible Bloom Filter (XBF) approach [20]. In the context of XBF, researchers have studied the delay introduced by inter-zone routers [49] and, they have devised a segmentation method that minimizes the inter-zone forwarding [20]. Although this segmentation method reduces the number of inter-zone operations, it does so by increasing the number of inter-zone routers, increasing deployment and management costs and challenging economic scalability. Therefore, we decide to study a new segmentation method for the XBF approach that minimizes the number of inter-zone routers. To this end, (a) we demonstrate how the XBF zoning mechanism enables the creation of zones, (b) we present our segmentation method, and (c) we finally evaluate the scalability of the resulting FWD function.

- (a) The XBF approach enables the creation of zones in an IP-over-PURSUIT network. Table 6 shows how the number of zones can be increased by dividing the space reserved for FIDs (256 bits, the combined size of the IPv6 source and destination addresses) into ZIDs

(zones) and zone-specific FIDs [50]. Each bit added to the ZID length doubles the numbers of zones, while it decreases the number of links per zone by one. As a result, the number of total unidirectional links is almost doubled with each new zone. Note that positions in the FID are only used to identify one link per zone, thus preventing false positives in packet forwarding.

- (b) The main idea behind our segmentation method is to add links to a zone until the maximum number of links per zone is reached. At this point, we create a new zone. This process starts from the bottom of the topology and is recursively repeated until core links are reached. Given that links between NAPs and clients remain IP, the segmentation method starts by grouping NAPs in the access layer.⁷ At this point, if the node degree of the distribution layer is smaller than the maximum number links per zone, then all NAPs hanging from a node in the distribution layer belong to the same zone, including the node in the distribution layer. Otherwise, multiple zones need to be created for each node in the distribution layer, and this node becomes an inter-zone router.
- (c) We assess the scalability of the FWD function for a network topology that approximates a mobile operator serving approximately 2.4 million users, which is

⁷ Recent work in 3GPP Release 16 on vertical, cellular LAN specifies the use of path-based forwarding between core and access nodes only, i.e., the user plane functions in cellular networks, while utilizing LAN-based communication to reach UE.

Table 6 XBF zoning segmentation

ZID (bits)	sFID (bits)	Zones	Unid. links	UE
2	254	4	1016	
3	253	8	2024	
4	252	16	4032	1 10 ⁶
5	251	32	8032	2 10 ⁶
6	250	64	16,000	3 10 ⁶

derived from Mobile Network Operator (MNO) interviews in Finland [51]. Although this topology follows a simple hierarchical model, it captures the deployment structure of commercial access operators with a nationwide presence, in contrast to topologies from research networks or topology generators. This decision is coherent with the objectives of the study that aims at assessing market feasibility. More importantly, the selected topology addresses the minimum deployment requirements of stakeholders identified as possible adopters during the value network analysis, i.e., IP-connectivity operators, triple play operators, MNOs, and micro-operators. As shown in Fig. 5, the topology is determined by node degrees of 714, 5, and 161 in the access, distribution and core layers, respectively. In each layer, nodes serve the same number of slaves and are connected to only one master (e.g., nodes in the transport layer have a node degree of 5 meaning that they are connected to 4 slave nodes on the access layer and to 1 master node on the core layer).

To estimate the number of required zones, we simulate a growing topology that gradually adds nodes to access, distribution, and core layers up to their maximum node degree by increasing the number of subscribers. Results show that the studied topology could be served using zones with 250 unidirectional links, as shown in the last row of Table 6. Since the node degree of distribution nodes is 10 (4 bidirectional links to access nodes and 1 bidirectional link with a core node), each core node can handle up to 25 distribution nodes without adding a new zone. This implies that only core nodes need to become inter-zone routers.

The next step in the scalability analysis is to identify how many zones each inter-zone router would need to handle. Based on the same simulation, Fig. 6 shows that although the number of total zones (in circles) increases with subscribers, the zones per core router vary between 1 and 7 (in rectangles). Figure 7 shows this same variation in parallel with the increase in the number of core routers (in triangles). This variation is caused by the assignment of zones to core routers, as new core routers are added to the topology.

Fig. 5 Topology structure of a Finnish mobile operator **Fig. 6** Growth in the number of zones **Fig. 7** Zones per router

This study reveals that inter-zone forwarding is only required in core nodes, similar to the OSPF zone 0 approach. More importantly, as the number of required zones grew

to accommodate new links and subscribers, the number of zones per router remained below 8. However, this zoning solution requires further study since the topology does not account for redundant links. Moreover, the payload of the PURSUIT packet needs to carry not only IP data, but also piggybacked FIDs. As a result, the ICN packet might grow and suffer layer-2 segmentation. At the current state of development, the IP-over-PURSUIT solution presents lighter zoning requirements than the OSPF recommendation by Cisco. Cisco guidelines on the configuration of OSPF recommend 50 routers per area, 3 areas per ABR, and a maximum node degree of 60 [50].

3.6.2 Challenge 4 and 5: stakeholder benefits

According to the value network analysis, operators need to take control over the provision of OTT content and over the provision of surrogate services to fully benefit from multicast routing and service routing, respectively. However, OTT providers and edge application providers might not be willing to concede this control since their role in the provision of content and services would be downgraded to content owners and application developers, respectively. For example, OTT providers are not interested in operators releasing core link capacity via HTTP coincidental multicast if, in return, they lose control over service provision.

Similarly, edge application providers might not be interested in HTTP service indirection if they need to compete with other providers for NAPs to switchover their connections. Therefore, these providers have little incentive to collaborate with an IP-over-PURSUIT operator independently of the operator-specific benefits, given that, in IP-based networks, these providers can maintain their current superior business roles. Without these providers conceding control, the benefits obtained by the operator might not be enough to justify the adoption of IP-over-PURSUIT over substitute solutions, such as IP-based solutions.

However, the potential reductions in costs through HTTP coincidental multicast or latency through HTTP service indirection may eventually yield benefits for OTTs/edge application providers, thus making it desirable to reconcile their need in obtaining useful service delivery information with the need of operators to have more control over the delivery. In operator-OTT negotiations, the latter may enjoy a stronger position, as demonstrated by the remote control that OTTs have over caches which are placed within operator networks, e.g., Google Global Cache, Netflix Open Connect.

To overcome this feasibility challenge, we suggest the implementation of technical interfaces allowing providers to remotely access the RV (and maybe strategic NAPs) to monitor technical and business information about their own content and services, i.e., providers can monitor service quality and account for service usage. By offering these new techni-

cal interfaces, operators create a win-win situation in which their new and cost-effective routing capabilities could be used without challenging the current business roles of providers. For the multicast routing use case, an OTT-operator interface needs to inform in real-time about the quality of the service delivery and usage statistics, matching the current information/control through HTTP/HTTPS. For example, this new interface needs to allow content providers to deliver personalized adverts, retrieve click counts, etc. For the service routing use case, an edge application provider-operator interface needs to allow providers to remotely control the prompt switchover of connections between surrogate services, e.g., according to NAP load. Similar inter-stakeholder interfaces exist in the Internet allowing remote access to MNO systems by Mobile Virtual Network Operators via Policy and Charging and Control interfaces [52].

3.6.3 Challenge 6: IP interoperability

The IP-over-PURSUIT solution was trialed to provide Internet access to 45 households near the city of Barcelona in Spain, using a newly deployed wireless network made of six nodes, as shown in Fig. 8 [53]. The trial network included all the elements in the IP-over-PURSUIT architecture which were deployed in software, including Open vSwitch switches. While all six nodes executed FWD and NAP functions, only one node executed the RV/TM and acted as the only Internet gateway. While the transport of HTTP data over the PURSUIT network was managed by *HTTP handlers*, the rest of the IP data (non-HTTP) has been indistinguishably managed by *IP handlers*.

The field trial has demonstrated for the first time that real IP traffic can be successfully routed over an ICN network without modifying UE or applications [53]. Furthermore, measurements accounting for 15 days of network activ-

Fig. 8 Field trial topology [53]

Table 7 Measured peak throughputs in the Barcelona trial [53]

	Protocol	Peak RX (Mbps)	Peak TX (Mbps)
Client-facing NAP	HTTP	0.187	8
Client-facing NAP	non-HTTP	1.8	8
Internet-facing NAP	HTTP	10	0.063
Internet-facing NAP	non-HTTP	9	2

ity between a client-facing NAP and the Internet-facing NAP revealed peak throughput values within the expected range. As shown in Table 7, data transmission (TX) and reception (RX) achieved peak values between 8 and 10 Mbps, respectively. These values approximately correspond to the maximum throughput achievable with the deployed equipment which was resource constrained. To be precise, laboratory deployments with similar resource-constrained equipment achieved throughputs no larger than 12 Mbps. Note that non-constrained laboratory measurements have achieved NAP-to-NAP throughputs up to 85 Mbps.

In addition to IP interoperability, the trial demonstrated that specific routing behavior can be applied to individual protocols by handling HTTP traffic independently from other IP protocols. However, measurements revealed that HTTP only accounted for a tiny 0,666% of the total traffic. Moreover, only 4 IP addresses were detected accessing HTTP services simultaneously. Under these conditions, the HTTP coincidental multicast could not have provided any efficiency gains to the underlying transport network. This result is consistent with the trend observed in other networks in which HTTP is rapidly substituted by HTTPS and emphasizes the importance of the upcoming feasibility challenge on the Control over the provision of OTT content. Finally, from an operational point of view, the trial also demonstrated that operators can gradually deploy ICN on their IP networks one link at a time [53].

3.7 Solution analysis

In this section, we discuss operator adoption of IP-over-PURSUIT as a single network routing solution (i.e., as IGP) depending on positive net benefits of stakeholders taking into account the operator service offering and the general market conditions.

Any operator that adopts the IP-over-PURSUIT routing solution will remain interoperable with IP networks and UE, as demonstrated by a field trial. More importantly, the IP integration is straightforward since NAPs manage publications and subscriptions based on IP address prefixes. Further, our initial calculations suggest that IP-over-PURSUIT could serve nation-sized network topologies while keeping the number of inter-zone routers under a reasonable upper bound. However, further study is required to frame with precision the

scalability requirements of the RV/TM for content-intensive use cases.

Connectivity operators might be interested in IP-over-PURSUIT since it enables flow-based routing. In contrast to substitute solutions, IP-over-PURSUIT facilitates the frequent update of flow paths through a light SDN integration. Flow paths are easily updated since these are centrally computed by the RV/TM and only disseminated to NAPs when requested by new service queries. Furthermore, operators can reuse preexisting SDN switches, integrate SDN switches more efficiently than current SDN solutions, and enjoy a lower controller-to-switch interaction. Thus, the IP-over-PURSUIT solution can delay and reduce capital expenditure (CAPEX) by reusing existing equipment and by acquiring equipment with less memory and computing power. Further, it can also reduce operational expenditure (OPEX) by simplifying flow management. Regarding stakeholder collaboration, no new or unusual business contracts need to be signed for operators to enjoy the use case benefits. However, IP-over-PURSUIT will encounter fierce competition from market-ready solutions implementing IPv6.

Connectivity operators that also provide content-based services might be interested in IP-over-PURSUIT since it enables the multicast of media content. For example, triple play operators that either own or plan to acquire content distribution rights have a greater incentive to adopt IP-over-PURSUIT routing solutions since they can continue contract-based content distribution via IP multicast. Further, IPTV services would gain in scalability, increasing the maximum number of synchronous viewers, given that P2MP transmissions do not store per-flow state in nodes. Furthermore, triple play operators could spare capacity in core links by delivering OTT traffic via HTTP coincidental multicast if OTT providers facilitate the proxy of HTTPS connections. While small content providers might allow this proxy, large OTTs like Netflix might not since their own delivery solutions guarantee a direct technical and contractual relationship with consumers via HTTPS and they do not obtain an immediate benefit from IP-over-PURSUIT. To address this challenge, operators need to convince OTTs that their business role as content providers is not challenged by IP-over-PURSUIT, e.g., by offering real-time information over the delivery of content through a new OTT-operator technical interface, by offering economic incentives when deploying caches within

operator networks. In terms of OPEX, the ability to centrally manage link-state and pub/sub information through the RV/TM can be a relevant differentiator against substitute solutions. Further, triple play operators might also be interested in introducing FQDN-based and delay-based routing policies.

It is unlikely that MNOs adopt IP-over-PURSUIT as a routing solution for the edge cloud since NAPs are not compatible with 3GPP signaling protocols enabling traffic steering to the edge cloud. As a result, HTTP service indirection between surrogate services in the telco and edge cloud is currently not possible. Nevertheless, recent developments in the release 15 of 3GPP specification for 5G may facilitate this integration, given the service-based architecture. Micro-operators covering a limited geographical area might be interested in adopting IP-over-PURSUIT as an IGP solution. In comparison to small-scale 3GPP-based solutions like MuLTEfire [54], IP-over-PURSUIT micro-operators could provide anchorless mobility and fast HTTP service indirection to low-latency applications via Wi-Fi. However, if micro-operators open the edge cloud to third-party applications, application providers might require real-time information and control over the HTTP service indirection to guarantee service quality.

4 IP-over-ICN deployment strategy

The long-term IP-over-ICN deployment strategy aims to establish a fully ICN-based Internet architecture in which Internet stakeholders interconnect based on information items [11]. This information-based interconnection can potentially reduce inter-domain traffic, although a critical mass of adopters is required to unleash significant network effects. Next, we study the key challenges delaying information-based interconnection between Internet stakeholders.

CDN-to-CDN interconnection is technically feasible since most CDNs implement content management systems, including pub/sub, e.g., Akamai [55]. However, CDNs have little incentive to interconnect since they compete for the same content items/owners. If operators adopt ICN-based routing systems, their willingness to share valuable content via operator-to-operator interconnection is also limited when coverages overlap since they compete for the same consumers. When coverages do not overlap, operators cannot share valuable content since distribution rights are typically exclusive for their coverage area. Moreover, the infrastructure of operators is not typically needed for global content distribution since content providers use CDNs.

Finally, operator-to-CDN interconnection can unleash substantial benefits for operators as studied for the multicast routing use case. However, this interconnection gives oper-

ators control over the provision of content, thus preventing content providers from monitoring content consumption. As a first consequence of this interconnection, CDNs may lose the ability to differentiate based on the Quality of Service (QoS). As a second consequence, depending on the IP-over-ICN implementation, content providers might not be able to establish direct contractual relationships with consumers nor provide personalized ads. Therefore, it is unlikely that content providers accept the CDN-to-operator interconnection, given that content is already delivered successfully via CDNs. In fact, OTT providers enjoy a dominant position on today's provision of content since OTTs keep control over caches which are placed within operator networks, e.g., Google Global Cache, Netflix Open Connect.

Hence, the advent of an ICN-based Internet architecture may be delayed until Internet stakeholders have enough incentives to delegate the delivery of valuable content and services. For this, inter-stakeholder technical and business interfaces need to emerge enabling trustworthy collaboration and compensation in future information-based interconnection points. For example, allowing stakeholders to charge one another for the delivery of copyrighted content through pay-per-delivery schemes, e.g., via smart contracts.

5 Conclusions

This article assesses the feasibility of IP-over-PURSUIT as single network routing solution for flow-based routing, multicast routing, and service routing use cases. First, we suggest that these use cases can be successfully served by the solution, albeit different but compatible architecture configurations are required. Second, we indicate that IP-over-PURSUIT potentially offers technical advantages in comparison to substitute solutions for the three use cases, including efficient forwarding, light SDN integration, almost stateless HTTP coincidental multicast, and fast HTTP service indirection. Third, the routing solution does not present immediate challenges regarding IP interoperability and routing scalability, as indicated by field trial evidence and the zone-based forwarding assessment, respectively. Fourth, we suggest that the IP-over-PURSUIT architecture should adopt new exterior interfaces since its current design may not provide enough stakeholder-specific benefits. Note that OTTs and edge application providers need to concede control over the provision of content and services for a significant fraction of the operator traffic to be delivered through HTTP coincidental multicast and HTTP service indirection.

We suggest that triple play operators owning content distribution rights have incentives to adopt the IP-over-PURSUIT solution since they can continue contract-based content distribution via scalable IP multicast. Further, triple play operators can spare capacity in core links by deliv-

ering OTT traffic via HTTP coincidental multicast if OTT providers are provided with remote service monitoring capabilities. The more traffic triple play operators can deliver through multicast, the larger the motivation to adopt IP-over-PURSUIT. In addition, the ability to centrally manage link-state and pub/sub information through a light SDN integration can be a relevant differentiator for operators against substitute solutions. In addition to triple play operators, micro-operators covering a limited geographical area might also be interested in adopting IP-over-ICN to provide anchorless mobility and fast HTTP service indirection to low-latency applications via Wi-Fi.

Finally, we indicate that the advent of an ICN-based Internet architecture via the IP-over-ICN deployment strategy might be delayed until Internet stakeholders can delegate the delivery of valuable content and services via automatic compensation schemes in information-based exchange points.

The IP-over-PURSUIT routing solution will be further developed in the H2020 FLAME project to provide flexible routing of IP services following a Service as a platform model. In this context, IP-over-PURSUIT has been recently applied to service routing in 5G core networks through efforts on service-based architecture in 3GPP Release 15.

Acknowledgements Open access funding provided by Aalto University. The authors would like to thank Arto Karila and Pekka Nikander for their valuable comments and discussions. This work was supported by the EC H2020 RIFE Project Grant No. 644663 and by the EC H2020 POINT Project Grant No. 643990.

Compliance with ethical standards

Conflict of interest Dr. Dirk Trossen is directly involved in the development of IP-over-PURSUIT solutions at InterDigital. The other authors have no conflict of interest.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Koponen, T., Chawla, M., Chun, B.-G., Ermolinskiy, A., Kim, K. H., Shenker, S., et al. (2007). A data-oriented (and beyond) network architecture. *ACM SIGCOMM Computer Communication Review*, 37(4), 181–192.
- Jacobson, V., Smetters, D. K., Thornton, J. D., Plass, M., Briggs, N., & Braynard, R. (2012). Networking named content. *Communications of the ACM*, 55(1), 117–124.
- Trossen, D., & Parisi, G. (2012). Designing and realizing an information-centric internet. *IEEE Communications Magazine*, 50(7), 60–67.
- Dannewitz, C., Kutscher, D., Ohlman, B., Farrell, S., Ahlgren, B., & Karl, H. (2013). Network of information (NetInf)—An information-centric networking architecture. *Computer Communications*, 36(7), 721–735.
- Handley, M. (2006). Why the Internet only just works. *BT Technology Journal*, 24(3), 119–129.
- Turner, J. S., & Taylor, D. E. (2005). Diversifying the Internet. In *GLOBECOM—IEEE global telecommunications conference*.
- Clark, D. D., Wroclawski, J., Sollins, K. R., & Braden (2002). Tussle in cyberspace. In *Proceedings of the 2002 conference on applications, technologies, architectures, and protocols for computer communications—SIGCOMM '02*, no. 4. New York, NY: ACM Press.
- Clark, D. D., Wroclawski, J., Sollins, K. R., & Braden, R. (2005). Tussle in cyberspace: Defining tomorrow's internet. *IEEE/ACM Transactions on Networking (ToN)*, 13(3), 462–475.
- Kostopoulos, A., Papafili, I., Kalogiros, C., Levä, T., Zhang, N., & Trossen, D. (2012). A tussle analysis for information-centric networking architectures. In F. Álvarez, et al. (Eds.), *The future internet* (pp. 6–17). Berlin, Heidelberg: Springer.
- Agyapong, P. K., & Sirbu, M. (2012). Economic incentives in information-centric networking: Implications for protocol design and public policy. *IEEE Communications Magazine*, 50(12), 18–26.
- Trossen, D., & Kostopoulos, A. (2011). Exploring the tussle space for information-centric networking. In *Telecommunications policy research conference (TPRC)*. Social Science Research Network (SSRN). <https://ssrn.com/abstract=1979924>.
- Rahman, A., Trossen, D., Kutscher, D., & Ravindran, R. (2019). *Deployment considerations for information-centric networking (ICN)*. Internet Engineering Task Force. <https://datatracker.ietf.org/doc/html/draft-irtf-icnrg-deploymentguidelines-06>.
- Muscariello, L., & Engineer, P. (2018). *Hybrid information-centric networking ICN inside the Internet protocol*. <https://datatracker.ietf.org/meeting/interim-2018-icnrg-01/materials/slides-interim-2018-icnrg-01-sessa-hybrid-icn-hicn-lucamuscarillo>.
- Muscariello, L., Carofoglio, G., Auge, J., & Papalini, M. (2018). *Hybrid information-centric networking*. Internet Engineering Task Force. <https://datatracker.ietf.org/doc/html/draft-muscariello-intareahicn-00>.
- Levä, T., & Suomi, H. (2013). Techno-economic feasibility analysis of Internet protocols: Framework and tools. *Computer Standards and Interfaces*, 36(1), 76–88.
- Xylomenos, G., Al-Khalidi, M., Al-Naday, M., Fotiou, N., Karila, A., Petropoulos, G., et al. (2018). *H2020 POINT D2.4 scenarios, requirements, specifications and KPIs*. Technical report 643990.
- Fotiou, N., Trossen, D., & Polyzos, G. C. (2012). Illustrating a publish-subscribe Internet architecture. *Telecommunication Systems*, 51(4), 233–245.
- Jokela, P., Zahemszky, A., Rothenberg, C. E., Arianfar, S., & Nikander, P. (2009). LIPSIN: Line speed publish/subscribe inter-networking. *Design*, 39(4), 195–206.
- Reed, M. J., Al-Naday, M., Thomos, N., Trossen, D., Petropoulos, G., & Spirou, S. (2016). Stateless multicast switching in software defined networks. In *2016 IEEE international conference on communications, ICC 2016*.
- Antikainen, M., Wang, L., Trossen, D., & Sathiaselan, A. (2016). XBF: Scaling up bloom-filter-based source routing. [arXiv:1602.05853](https://arxiv.org/abs/1602.05853).
- Trossen, D., Reed, M. J., Riihijärvi, J., Georgiades, M., Fotiou, N., & Xylomenos, G. (2015). IP over ICN-the better IP? In *2015 European conference on networks and communications (EuCNC)*.
- Xylomenos, G., Al-Naday, M., Fotiou, N., Karila, A., Petropoulos, G., Reed, M. J., et al. (2016). *H2020 POINT D2.3 scenarios, requirements, specifications and KPIs, 3rd version*. Technical report.

23. Al-Khalidi, M. Q., Thomos, N., Martin, R., Al-Naday, M., & Trossen, D. (2018). Anchor free IP mobility. *IEEE Transactions on Mobile Computing*, 18, 56–69.
24. Rahman, A., & Trossen, D. (2018). *Alternative handling of dynamic chaining and service indirection draft-purkayastha-sfc-service-indirection-01*.
25. Levä, T., Mazhelis, O., & Suomi, H. (2014). Comparing the cost-efficiency of CoAP and HTTP in Web of Things applications. *Decision Support Systems*, 63, 23–38.
26. Levä, T., Komu, M., Keränen, A., & Luukkainen, S. (2013). Adoption barriers of network layer protocols: The case of host identity protocol. *Computer Networks*, 57(10), 2218–2232.
27. Karakus, M., & Duresi, A. (2017). A survey: Control plane scalability issues and approaches in software-defined networking (SDN). *Computer Networks*, 112, 279–293.
28. Rich, M. (2017). *What behavioral data tells us about the OTT viewing habits of cord-cutters*. comScore, Inc.
29. Enoch, G. (2016). *The Nielsen total audience report—Q2 2016*. Technical report.
30. Telia Company acquires Bonnier Broadcasting - Telia Company. Retrieved August 2, 2018 from <https://www.teliacompany.com/en/news/press-releases/2018/7/telia-company-acquires-bonnier-broadcasting/>.
31. Bilal, K., Khalid, O., Erbad, A., & Khan, S. U. (2018). Potentials, trends, and prospects in edge technologies: Fog, cloudlet, mobile edge, and micro data centers. *Computer Networks*, 130, 94–120.
32. Casey, T., Smura, T., & Sorri, A. (2010). Value network configurations in wireless local area access. In *2010 9th conference of telecommunication, media and Internet, CTTE 2010*. IEEE.
33. Benseny, J., & Hammainen, H. (2016). Value network analysis in a low-cost and affordable internet. In *EUCNC 2016—European conference on networks and communications*.
34. Kekki, S., Featherstone, W., Fang, Y., Kuure, P., Li, A., Ranjan, A., et al. (2018). *ETSI white paper no. 28 MEC in 5G networks*. Technical report.
35. Suthar, P., Stolic, M., Jangam, A., Trossen, D., & Ravindran, R. (2018). *Native Deployment of ICN in LTE, 4G Mobile Networks*. Internet Engineering Task Force, Internet-Draft draft-irtf-icnrg-icn-lte-4g-02, work in Progress.
36. Walia, J. S., Hammainen, H., & Matinmikko, M. (2018). 5G micro-operators for the future campus: A techno-economic study. In *Joint 13th CTTE and 10th CMI conference on Internet of Things—Business models, users, and networks*. IEEE.
37. Rosen, E., Viswanathan, A., & Callon, R. (2001). *Multiprotocol label switching architecture*. Technical report.
38. Moy, J. (1998). *OSPF Version 2*. Technical report.
39. Petropoulos, G., Katsaros, K. V., & Xezonaki, M. E. (2017). OpenFlow-compliant topology management for SDN-enabled information centric networks. In *Proceedings—IEEE symposium on computers and communications*.
40. Eckert, T., Leymann, N., & Napierala, M. (2013). *Multipoint LDP in-band signaling for point-to-multipoint and multipoint-to-multipoint label switched paths*. <https://doi.org/10.17487/rfc6826>.
41. Dolganow, A., Przygienda, T., & Aldrin, S. (2017). *Multicast using bit index explicit replication (BIER)*. RFC 8279.
42. Purkayastha, D., Rahman, A., Trossen, D., & Eckert, T. (2018). *Applicability of BIER multicast overlay for adaptive streaming services*. Internet Engineering Task Force, Internet-Draft draft-purkayastha-bier-multicast-http-response-01, work in Progress. Retrieved May 2, 2019 from <https://datatracker.ietf.org/doc/html/draft-purkayastha-bier-multicast-http-response-01>.
43. 3GPP TS 33.402. (2011). *3rd generation partnership project; technical specification group services and system aspects; 3GPP system architecture evolution (SAE); Security aspects of non-3GPP accesses*. 3GPP, Technical report v.11.2.0 - Release 11.
44. ETSI. (2018). *GR MEC 017 - V1.1.1—Mobile edge computing (MEC); Deployment of mobile edge computing in an NFV environment*. Technical report.
45. ETSI. (2017). *Mobile edge computing (MEC); Mobile edge platform application enablement*. Technical report.
46. White, G., & Rutz, G. (2016). *Content delivery with content-centric networking*. Cable Television Laboratories, Inc.
47. Wang, S., Bi, J., Wu, J., Yang, X., & Fan, L. (2012). On adapting http protocol to content centric networking. In *Proceedings of the 7th international conference on future internet technologies* (pp. 1–6). ACM.
48. Rajahalme, J., Särelä, M., Visala, K., & Riihijärvi, J. (2009). *Inter-domain rendezvous service architecture PSIRP technical report #TR09-003*. Technical report.
49. Kärkkäinen, T., Baig, R., Sathiseelan, A., Ali, A., Isah, M., Arnal, F., et al. (2017). *H2020 RIFE D3.3: Final platform design and set of dissemination strategies*. Technical report.
50. Tiso, J. (2011). *Designing Cisco network service architectures (ARCH): Developing an optimum design for layer 3 (CCDP)*. <http://www.ciscopress.com/articles/article.asp?p=1763921&seqNum=6>.
51. Zhang, N., & Hammainen, H. (2015). Cost efficiency of SDN in LTE-based mobile networks: Case Finland. In *2015 international conference and workshops on networked systems (NetSys)*. IEEE.
52. 3GPP. (2018). *TS 29.212 V15.3.0 Policy and Charging Control (PCC). Reference points*.
53. Selimi, M., Isah, M., Sathiseelan, A., Baig, R., Trossen, D., Krishna, R., et al. (2018). *H2020 RIFE D4.3: Final report on technology validation*. Technical report.
54. Specification — MulteFire. Retrieved August 17, 2018 from <https://www.multefire.org/specification/>.
55. Nygren, E., Sun, J., Sitaraman, R. K. R., Sun, J., Sitaraman, R. K. R., & Sun, J. (2010). The Akamai network: A platform for high-performance Internet applications. *ACM SIGOPS Operating Systems Review*, 44(3), 2–19.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Jaume Benseny is doctoral researcher in network economics at the Department of Communications and Networking at Aalto University in Finland. He received his M.Sc. in communications ecosystem from Aalto University in 2016 and his degree in telecommunications engineering from Technical University of Catalonia (ETSETB-UPC) in 2006. In between degrees, he worked at i2CAT research center. His research interests include the techno-economic scalability of networking architectures and operator models.

Dmitrij Lagutin received the M.Sc. (Tech.) degree from the Helsinki University of Technology, Finland, in 2005 and the D.Sc. (Tech.) degree from Aalto University, Finland, in 2010. He was a Researcher in several research projects with Helsinki University of Technology and Aalto University, including EU FP7 PSIRP and PURSUIT and an EU H2020 POINT projects. He is currently a Coordinator and Research Fellow in the EU Horizon 2020 SOFIE Project, working at Aalto University, Finland. His

research interests include network security and privacy, Internet of Things, blockchains, and future network technologies.

Heikki Hämmäinen is Professor of Network Economics at Department of Communications and Networking, Aalto University, Finland. He has M.Sc. (1984) and Ph.D. (1991) in Computer Science from Helsinki University of Technology. His main research interests are in techno-economics and regulation of mobile services and networks. Special topics recently include measurement and analysis of mobile usage, value networks of flexible Internet access, and diffusion of Internet

protocols in mobile. He is active in several journals and conference duties.

Dirk Trossen is a Senior Principal Engineer at InterDigital Europe, the European branch of InterDigital Inc. His main responsibility lies in establishing the European presence of InterDigital through engagements within the EU-funded Horizon 2020 work programme as well as within UK-funded efforts. Dirk has more than 25 years of experience in network architectures, services and wireless technology. He is currently technical leading the European efforts FLAME, deploying

and trialing the evolutionary introduction of ICN concepts in upcoming 5G networks. He is also an active contributor to European efforts in the 5G(PPP) space through contributions to ETP and 5GPPP whitepapers. Prior to joining InterDigital, Dirk was co-founder of TecVis LP, a UK-based software solution company in the mobile, context-aware solution space and he held prior positions as a Senior Researcher with Cambridge University, Chief Researcher with BT Research and as a Principal Scientist at Nokia Research. He is also a research affiliate with the Advanced Network Architecture group at MIT CSAIL. He holds a Ph.D. degree in Computer Science from Technical University of Aachen, Germany. He published more than 85 peer-reviewed papers in international conferences and journals and has currently 32 international patents.