

# Towards an Information Centric Network Architecture for Universal Internet Access

Dirk Trossen  
InterDigital Europe, Ltd  
first.last@interdigital.com

Arjuna Sathiaseelan  
University of Cambridge  
first.last@cl.cam.ac.uk

Joerg Ott  
Technische Universitaet Muenchen  
first.last@tum.de

This article is an editorial note submitted to CCR. It has NOT been peer reviewed.

The authors take full responsibility for this article's technical content. Comments can be posted through CCR Online.

## ABSTRACT

Enabling universal Internet access has been recognized as a key issue to enabling sustained economic prosperity, evidenced by the myriad of initiatives in this space. However, the existing Internet architecture is seriously challenged to ensure universal service provisioning at economically sustainable price points, largely due to the costs associated with providing services in a perceived always-on manner. This paper puts forth our vision to provide global access to the Internet through a universal communication architecture that combines two emerging paradigms, namely that of Information Centric Networking (ICN) and Delay Tolerant Networking (DTN). The decoupling in space and time, achieved through these underlying paradigms, is key to aggressively widen the connectivity options and provide flexible service models beyond what is currently pursued in the game around universal service provisioning. In this paper, we provide an outlook on the main concepts underlying our universal architecture and the opportunities arising from it. We also offer some insight into ongoing work to realize our vision in a concrete test bed and trial setting.

## 1 Introduction

The Internet has played a great role in fuelling the transition of our society from the industrial age to the information age and is seen as a fundamental driver of today's knowledge economy. The Internet's impact is imprinted on all spheres of human life—personal, societal, political, economical, and educational—in both developing and developed countries. However, only 40% of the world's population, mostly in developed economies, has access to the Internet [2]. Emerging economies such as Africa, Asia and Latin America have some of the lowest Internet penetration rates. Some of the reasons cited for lack of Internet access are affordability, lack of infrastructure, perceived lack of need, linguistic barriers etc.

Addressing digital exclusion due to socio-economic barriers is extremely important. The United Nations revealed the global disparity in fixed broadband access, showing that access to fixed broadband mainly in less-developed countries costs almost 40 times their national average income [6]. It is to be noted that there is an intertwine between socio-economics and geography when it comes to defining the barriers for universal Internet access. Access problems often result from sparsely spread populations living in physically remote locations, since it is simply not cost effective for Internet Service Providers (ISPs) to install the required infrastructure for broadband Internet access in those areas. In addition to the physical limitations of terrestrial infrastructures to provide last mile access (mainly due to distance), remote communities also incur higher costs for connecting the local exchanges to the backbone networks when using wired technologies. A large exchange may accommodate many users

and allow for competition between service operators; in contrast, rural/remote broadband often does not offer economies of scale, raising the costs per user. Most importantly, in many developing countries, poor connectivity between ISPs is so prevalent that local traffic is routed over expensive international links in an effort to ensure that it successfully reaches even destinations within the country of origin. The notion of economic barriers directly relates to the need for regulatory/policy changes as witnessed in some of the recent developments in Africa, Asia and Latin America. These kinds of challenges promote questioning established wisdom on the net: Should we insist on end-to-end delivery, or rather promote (more than standard Internet developments) localized communication—in contrast to what CDNs do, which is focussed mainly on pushing mainstream content closer to the users for latency reduction and load balancing. The result is a socioeconomic obstacle: mainstream business models don't work [14].

**Challenges stemming from the Always-on Internet:** The end-to-end always-on connectivity required for the current Internet creates a substantial barrier not only for implementing new flexible access/economic models [15, 16] but also induces increased wastage of both capacity and energy—which are the fundamental drivers for reducing the cost of Internet access.

1. The scheduling uncertainty of the end-to-end nature of the Internet creates a substantial barrier to implement time-shifted access (without, of course, implementing costly middleboxes), which could bring in new lower cost access opportunities (for e.g., exploiting under-utilised infrastructures [15, 16]). [11] has shown that, by taking advantage of already-paid-for off-peak bandwidth resulting from diurnal traffic patterns and percentile pricing, delay tolerant asynchronous bulk data (on the order of several terabytes) can be transferred effectively without incurring any transmission cost to the ISP.
2. The end-to-end always-on nature of the current Internet architecture introduces scheduling uncertainty forcing a receiver to continuously wait for packets, inevitably enforcing an energy-wasting policy. Energy is a scarce resource in many developing/less developed countries and hence technologies that can save energy is of paramount importance. We believe that we do not have to optimize the system for “always-on” connectivity, either for routing or access. What we need are mechanisms that enable a certain degree of delay tolerance that can keep a device's network interface controller in idle or sleep mode as long as possible without violating the applications' time constraints. Simple variants of these ideas were, e.g., presented for creating aggregating downstream traffic to mobile devices into bursts using in-network proxies to reduce mobile energy consumption and prolong device battery life [23].

3. With the growing need for accessing more content, the current host-centric model of the Internet architecture leads to wasting resources. Examples for such wastage include redundant transmissions, which in turn might lead to unnecessary congestion and therefore waste of energy. Furthermore, popularity as well as relevance relations between content and its viewer populations show that there are many missed opportunities to cache content on- as well as off-path, even in the ever-increasing presence of CDNs [7, 20]. By enabling the Internet architecture to deliver both the content and service from locations closer to the end-user, better service quality can be provided at lower costs, increasing the competitiveness of network operators. This increased content delivery efficiency can also result in significant energy savings for the network operator. Although Content Delivery Networks (CDN) do currently try and deliver content from geographical locations closer to the users, these solutions are proprietary and costly, thus requiring the operators to deploy CDNs in geographical areas where the costs of deploying such CDNs have an economic benefit while the others are left out. Moreover, the current CDNs are deployed as overlays and hence are not efficient enough to utilize the multicast and broadcast capabilities of the underlying network, sub-optimal routing etc [10].

These challenges inherently point out to a potential solution, too: could the key to enabling universal Internet access lie within an architecture that has been touted as a possible replacement for IP: Information Centric Networking (ICN)? This question was first raised in [17][16]. In this paper, we probe this further: Can ICN be a potential solution for enabling universal Internet access? The rest of the paper is as follows: we first discuss ICN and its associated benefits, and what are the current deployment obstacles in Section 2. In Section 3, we take into consideration the challenges discussed in Section 2 and propose a unified ICN architecture. Section 4 discusses some key benefits of our proposed architecture, and we finally present our conclusions and perspectives in Section 5.

## 2 Information Centric Networking (ICN): A Saviour That Came to Rescue?

Solving the problem of universal Internet access requires understanding the problem itself in the first place: what do users actually expect from Internet access? They need access to crucial Internet services and their associated content. These services differ from region to region and local services and interactions may be more important than central ones. This is where we believe Information Centric Networking (ICN) could play a key role towards universal Internet access [17]. It may change how users communicate and access information, moving from the traditional host-centric access paradigm—where access to a piece of content is mapped to its (fixed) location—to an information centric model, which eliminates this mapping content and supports access irrespective of the location where the content is held. ICN enables efficient resource management allowing the joint optimisation of network capacity, storage, and computation resources. We expect this in turn to support our ambitious goals of efficiency and economic sustainability.

1. Decoupling the content from the location removes the need for the current end-to-end client server model. Thus, the services and content can be served directly by any node that is able to offer the service or content at a given point in time. This inherently addresses the issues of mobility and reliability.
2. ICN supports a framework for mapping interests to publishers. This enables shifting demands in time and space and thus allows devising flexible access/economic models that could

make Internet access more affordable and offer better utilization of the capacity and energy (for e.g. transmitting during offpeak hours). The framework for mapping interests to publishers also removes access to SPAM content (e.g. Nigerian widows)—thus reducing the amount of unnecessary traffic that passes through the expensive backhaul.

3. ICN further improves the notion of CDNs by integrating the provisioning of content with the locationless notion of information delivery in ICN. This allows for different flavours of caching, from on-path caching to edge caching through a farm of surrogate servers. Any such caches can be quickly integrated into the overall (ICN) routing fabric without the need for DNS redirection or other cumbersome solutions of the current Internet. This increases the flexibility of the caching functionality with the ultimate goal to provide tradeoffs between timeliness and cost of delivery, which is particularly suitable for remote community deployments.
4. Although ICN has not been primarily designed for real-time (user-to-user) communication, the ability to localize traffic in ICN (e.g., through serving content through a local cache) alleviates pressure on the overall network, which in turn can benefit largely non-cacheable real-time communication through the increase in available aggregate bandwidth throughout the network. Furthermore, the support for native multicast in an ICN architecture also directly improves on real-time multiparty communication, particularly for ICN architectures that provide stateless multicast in an advance beyond the current heavily stateful IP multicast.

Although the benefits of ICN clearly exhibit a very positive perspective with a potential to address the aforementioned challenges of the always-on Internet through its spatial and temporal decoupling, ICN has also received criticism for privacy related issues, inter-domain policy issues, and most importantly not being able to scale [8]. Hence, for any attempt to integrate ICN into an architecture for universal Internet access, it is important to also address the aforementioned critical points for its successful adoption. Another important aspect is maintaining the Internet experience we currently enjoy. Backward compatibility not only fosters adoption due to the existing device basis that can be exploited but also reduces the disruption by not relying solely on novel ICN applications.

## 3 RIFE: An Unified Architecture for an Internet for Everybody

Following our discussion so far, we envision that an architecture for an Internet for everybody encompasses a range of connectivity options. This is to ensure universal coverage through a single unifying communication architecture, with a single set of abstractions that not only spurs innovation for a wide range of new services and applications but also embraces existing successful Internet services. In the following, we discuss such an architecture in more detail.

For such an architecture, ICN is a cornerstone, since its publish/subscribe paradigm allows for shifting demands in time and space. That is, expressions of interest (subscriptions) can be satisfied long after they have been issued, and may be served from any entity that has a copy of an object that matches the interest. This is in tandem to the philosophy of delay-tolerant networking (DTN) [5]. For instance, DTN deals with discrete data items and operates hop-by-hop, where each hop may store an item for an extended period of time until the next hop(s) become available for forwarding (store-carry-forward). Optionally, intermediate nodes may be required to retain content until reception has been confirmed. This allows for time and space shifting of when and where resources are used. It

is similar to ICN designs that involve hop-by-hop communication between the ICN architectural elements.

By its nature, hop-by-hop communication does not require global reachability per se. As a consequence, one could run any of these designs either over IP or as its replacement. What we propose is an integrative architectural platform that brings IP, ICN and DTN together into a single framework, in which DTN complements current IP and ICN solutions as an ideal candidate for communication in network environments, where added delay and disruption tolerance allows operation in disconnected environments.

We will capitalize on this observed behavior by defining a joint ICN/DTN architecture, which creates communication opportunities acquired by reclaiming unused capacity from commercial carriers in order to provide a low cost network service to the socio-economically disadvantaged [17][16]. The concept of an overarching ICN architecture—that integrates both with DTN and IP—enables us to pursue connected and disconnected modes of access under a single architectural abstraction. The adoption of ICN and DTN could develop enhanced QoS mechanisms for the backhaul infrastructure, utilising the generally information-centric nature of its edge network deployments, most specifically the opportunistic nature of DTN. This will increase the overall utilization of the network by exploiting under-utilised infrastructures in a Less-than-Best Effort (LBE) class while providing QoS enhanced services as a differentiation [15, 16]. The information-centricity of our architecture will allow for further reduction of transfers through caching at the edge of the network. This providing more localised access to important content enables a *transmit-when-needed* policy, thus reducing cost per bit access.

An Internet for Everybody cannot only compete, however, with the promise for better applications built on a new architecture, such as the one we propose in the following. Instead, any solution for bringing the Internet to everybody must also support the experiences that exist in the current Internet. With that, the Internet as we know it is the primary killer application for any solution without which we cannot succeed in bringing the Internet to a wider audience. Novel applications, heavily utilizing the new capabilities of the proposed new architecture, particularly its support for locality and the temporal decoupling of operations, will ultimately progress beyond the current experiences and accelerate the adoption of the architecture. With that in mind, we must pay crucial attention to supporting legacy IP-based applications, possibly even with an operational efficiency that exceeds that of today's IP Internet solutions.

Backward compatibility is not only important when it comes to the experiences that users have been used to. It also touches upon the important aspect of minimizing the need for equipment renewal, particularly for end users and edge equipment providers such as local community organisations. Hence, any system solution must take into account this aspect for a successful adoption in the real world.

### 3.1 Our Chosen ICN Starting Point

ICN has been a flourishing research area for many years now with a variety of flavours in terms of architectural and implementation choices [9, 21, 7]. For our architecture, we heavily rely on the concepts introduced by the PURSUIT flavour of ICN [3, 21] for the following reasons:

The PURSUIT flavour of ICN provides a graph-based object model which integrates well with the DTN (Bundle Protocol) [18] object approach to information

PURSUIT ICN functionality is provided by its three core network functions:

- **Rendezvous:** matches supply of information to demand for it. This process results in some form of (location) information

that is used for binding the information delivery to a network location.

- **Topology management and formation:** realizes the management of the overall delivery topology and the formation of specific delivery graphs.
- **Forwarding:** receives publications and forwards them to the network and/or to the local node. With this control and data plane separation, routing and forwarding are decoupled, enabling to trade off options in state management between various network components. This separation aligns very well with DTN as well as concepts in software-defined networking (SDN).

PURSUIT also functionally scopes the dissemination of information, i.e., it allows different strategies for the three core network functions to co-exist. This is the key to integrating DTN with ICN since, from an ICN perspective, DTN represents a particular dissemination strategy—in this case, for “challenged networks”. This strategy coexists with other strategies that are, for instance, highly optimised for optical high-speed networks. The strategy concept is crucial for spanning the expected spectrum of connectivity options, in order to aggressively acquire and utilise any connectivity option available.

RIFE complements these dissemination strategies through solutions for caching and replicating information and content, in particular in the edge networks. Replication can be achieved through set reconciliation [19] and network coding [12] to proactively push and update information toward network nodes, while caching uses local algorithms to opportunistically store received data. Due to the fine granularity of object representation of information, several objects can be fetched from several sources over multiple paths. In this way, retrieval times for information can be improved, positively impacting the overall user experience. Edge caching will also aid mobility solutions, where handover mechanisms are supported by replicated content in the delay-tolerant edge network. Moreover, autonomous operation based upon cached content objects will enable object manipulation even in the absence of instant connectivity to a cloud infrastructure.

### 3.2 RIFE System Architecture

The socket emulation [24] capabilities of our architecture allow running legacy IP applications on RIFE-enabled end devices. However, to avoid modifying the end devices and to preserve the IP interfaces towards the user equipment (UE) and apps, RIFE uses a gateway approach: the bridging between IP and ICN is performed in the network attachment points (NAPs) [22], i.e., the access gateways from customers to the network, and the RIFE border gateway (RIFE BGW), i.e., the access from and to the general Internet. As can be seen in Figure 1, we largely assume an IP-based interface from the UE to the NAP, following the gateway approach. The interface represents a collection of all supported abstractions with IP being supported as a bare minimum.

In our architecture, (a) the NAP serves as an IP-ICN gateway (handling all offered abstractions, i.e., HTTP, TCP, COAP, IP). Moreover, (b) ICN might be exposed towards the UE in selected environments, such as IoT, due to the fluidity that still exists in some environments; note that the strategy used in these selected environments is likely going to differ from those used in the ISP's network. Finally, (c) all communication beyond the ICN BGW complies with standard IP technologies, making the ISP network appear like a standard IP AS.

The figure also shows the interfaces between the logical ICN core functions, i.e., rendezvous (RV), topology management (TM), and forwarding (FN). Here, the communication between publisher (and subscriber) and rendezvous is captured as ICN(PR), invoking the

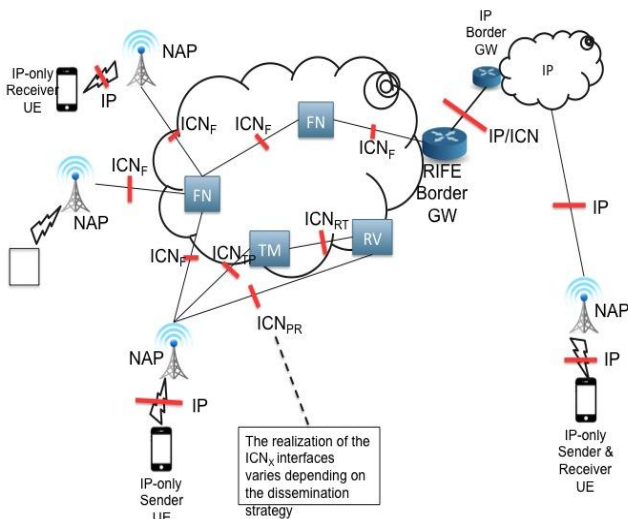


Figure 1: RIFE System Architecture.

ICN(RT) interface in the case of a positive match and finally invoking the ICN(TP) for delivering the appropriate forwarding information to the publisher. Finally, the basic ICN(F) interface is used for forwarding operations.

Although the functions and interfaces in Figure 1 seem to imply a centralized implementation, e.g., for the RV and the TM functions, their distributed realization is not only doable, in alignment with the architecture of Figure 2, but also desirable in many deployment scenarios. For instance, work in [13] has outlined a hierarchical intra- and inter-domain rendezvous system, which scales to the size of the Internet, while allowing for several local options for rendezvous. Also, the TM function would usually be implemented for a region of the network, often defined through geography or through network technology being used (e.g., core network resources being managed by one TM realization while access network resources are managed by another). For such distributed realization of the functions, the corresponding interfaces would implement a protocol that is distributed across a number of participating entities, e.g., RV implementation.

### 3.3 RIFE Node Platform

The main goal of the proposed unified RIFE architectural platform is to efficiently exploit all possible communication opportunities, from fixed or mobile broadband networks to (partly) disconnected networks and satellite links, while providing a unified abstraction to application developers for supporting current Internet-based services and enabling innovative future solutions.

Our RIFE architectural platform combines IP, ICN, and DTN solutions into an unified system architecture, exposing a common information-centric abstraction to applications, while supporting a range of networking protocols over different transport networks. Figure 2 presents our architectural platform.

The platform encompasses a number of different dissemination strategies, security functions, and application support functions. Our platform is highly modular, comprising a platform core with a number of connectors for plug-ins to be extensible (even at runtime). Figure 1 depicts this principal design in a simplified fashion: the core essentially offers the forwarding and scheduling logic and interconnects the services, support functions, network interfaces, and local resources as per the supplied system configuration. The platform supports both physical network interfaces (NICs), offering support for direct link layer access, and logical network interfaces to implement overlay functions. The dissemination strategies define

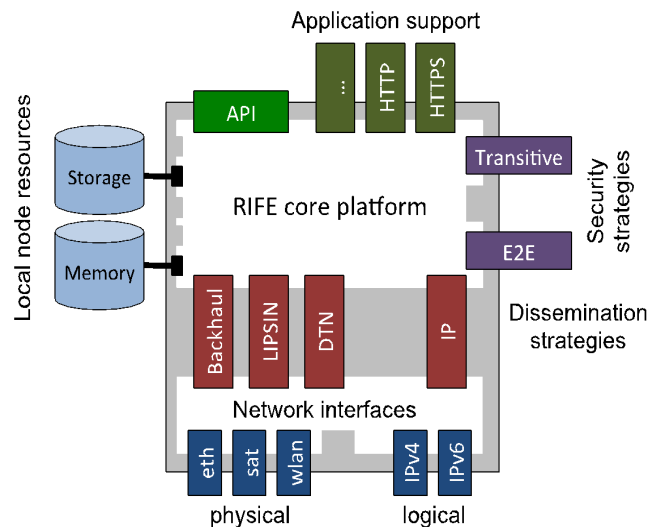


Figure 2: RIFE Node Platform.

the switching fabric between the interfaces and, along with routing tables, govern the data flow inside and between nodes, thus creating the network. Security strategies provide supplementary security support for strategies that no longer can rely on end-to-end connectivity (as in IP), whereas local resources offer computation, storage, and other local support functions. For example, local multi-tier storage may comprise persistent storage (for larger data volumes at slower access rates, e.g., for pre-distribution of content) and fast memory (for smaller data volumes at faster access rates, e.g., for media repair or in-memory processing). The platform offers an API towards applications to create novel applications that are aware of the extended functions of the RIFE platform on end systems as well as plug-ins for in-network support of legacy as well as novel applications.

## 4 Opportunities

The architecture we outlined above creates a number of opportunities, both at the level of technology and its enabling services as well as at the socio-economic level, some of which we will discuss in the following.

### 4.1 Providing IP-based services a lower price points

Postulating a new architecture for global access is ambitious, particularly when such new architecture and its capabilities only come to fruition through new services and a changed way of designing solution. Both communities, the ICN and DTN community, have suffered from outlining a new way of networking that rather explicitly requires the design and development of applications and services in order to reap the benefits that the new architecture provides.

In order to avoid that the adoption of any proposed new architecture is hampered by the lack of suitable applications and services, we argue that our architecture is not only capable of providing IP-based services as a migration strategy but such IP-based services on top of our new architecture can be provided at least at a comparable, possibly better price point than existing network deployments.

While the next sub-sections provide some insight into the improved capabilities and opportunities of native ICN/DTN as well as IP-based services on top of our architecture, we now briefly outline how IP-based services would be provided on top of our architecture. For this, we rely on the underlying paradigm in ICN, namely that everything is information and information is everything, in that we interpret any IP-based communication (over protocols like HTTP,

CoAP, TCP, or plain IP) as the exchange of information pertaining to a specific endpoint address. This endpoint address constitutes the name of the exchanged information. For illustration, consider a device with IP address A, which desires to send a packet to a device with IP address B. From an ICN perspective, such operation can be realized as device A publishing (an IP packet) information to the name B. Device B, in turn, can receive any such information by subscribing to (the name) B.

This idea was first presented in [24] with a working prototype that allowed for IP-based communication within a single ICN network. This initial work has now been extended towards any form of IP-based communication, presented in [22] and realized in an internationally funded research project called POINT [1]. Recently demonstrated in a first prototype, this work supports the realization of various IP-based protocols, such as HTTP, CoAP, TCP and basic IP datagram exchange, over the same ICN architecture that has been outlined in Section 3. This is achieved by mapping the various data structures of the underlying protocols onto suitably named objects within the ICN network, while also providing suitable network border functions to ensure connectivity to anywhere in the Internet. With such approach, IP-based services will be possible to realize over our proposed ICN architecture, providing a natural migration path towards purely native ICN solutions while providing qualitative and quantitative improvements, such as which we will discuss in the following sub-sections.

## 4.2 Simpler and fairer content distribution

In today's Internet, Content Delivery Networks (CDNs) enable content delivery with delays acceptable to end users. This is achieved by placing popular content at nearby (surrogate) servers. Such CDN solutions are primarily used by content aggregation services, such as YouTube or Vimeo, or network providers. However, CDN-based content placement creates a barrier of entry for smaller content providers and, in particular, individuals, due to the lack of exposure of the publishing APIs, which require agreements (and the economic buying power that comes with such agreements) with the CDN provider. Furthermore, while providing content through CDNs might lead to improved end user experience, non-popular content delivery might not be improved at all since it will not find entry into the CDNs because the latter usually focusses on popular content in order to make CDNs economically feasible. Hence, while certain content might well be relevant to a community, it will not be delivered through a CDN due to its (global) non-popularity. Furthermore, the dedicated placement in CDNs possibly wastes (caching) resources due to the inherent inflexibility of that approach.

We see at least two opportunities for improving today's situation for potential content distributors. Firstly, utilizing the explicit cache-aware resource management in ICN, e.g., as shown in [20], can increase the flexibility of providing the best quality to the content that is most popular within a given resource management regime as well as within a given (localized) context, i.e., removing the focus on highly aggregated, popular content in today's CDNs. Secondly, we also identify the opportunity to simplify CDN-like deployments. In our environment, such deployments can be based on smaller content providers utilizing localized computing resources in order to flexibly spin up HTTP-based content servers and connect them to an HTTP-over-ICN network, utilizing possible IOverICN solutions as outlined in Section 4.1. Specific CDN mechanisms for capturing the content requests are not necessary for these solutions since the surrogate (web) servers are mere copies of the original content server, flexibly dimensioned as virtual machines, with the traffic being directed to them through the shortest-path routing of the underlying ICN mechanisms from a particular client to the (surrogate) server.

The additional support for DTN in our architecture will further enhance the content distribution with opportunities to deliver content as slower, time-shifted, content access. For instance, we envision the seamless support for so-called content mules, which provide content storage in a DTN-type scenario, while being able to inject their content into a well-connected network upon attaching to such network. Such content mules can be utilized to remove pressure from backhaul connectivity by distributing content between, e.g., local villages, without the need to utilize often expensive backhaul connectivity but instead rely on physical movement between the settlements, such as through local travellers or public transportation.

## 4.3 New forms of resilience

Resilience is an essential part of any operator's network planning and management. Traditional forms of resilience have been enabled through IP routing protocols and physical/link layer protection and redundancy mechanisms. These mechanisms have provided resilience recovery with a response time in the orders of minutes, or seconds, for IP routing protocols down to the order of tens of milliseconds for physical layer protection mechanisms. While these mechanisms have proved to be effective in protecting the underlying network delivery, they do not provide mechanisms for host or server recovery. Server network resilience is provided through, relatively, newer mechanisms such as server duplication, load balancing systems, or DNS and/or HTTP redirection.

ICN provides a new opportunity for resilience through its inherent anycast delivery: multiple publishers can be providers for the same information thus creating information resilience [4]. If one publisher fails another can naturally take over without specialist intervention. In this paper, we extend this notion by considering that the endpoint IP address is simply an ICN name. This is unlike IP, where the endpoint address is uniquely bound to the routing function such that anycast (multiple senders with the same address) is difficult to achieve. Although DNS redirection does provide some limited anycast capability for IP networks, it is a clumsy tool that relies upon all DNS servers respecting short DNS TTL values and is not designed for application layer recovery.

In contrast, ICN information resilience allows a network designer to work at an appropriate level of abstraction. For example, it might be that a whole IP address prefix needs to be duplicated for resilience of a data center; alternatively it might be that a particular URL needs to be duplicated. In either case, the naming can be handled through an appropriate entry in the ICN namespace and a suitable (working) publisher can be chosen for the delivery. Clearly, handling resilience information for a whole IP address prefix requires less failure detection overhead than detecting a failure at the URL level. However, with ICN the choice about the level of detail for failure recovery rests with the network designer, rather than being limited by the routing or naming resolution, as it is with IP.

## 4.4 Improved Quality-of-Service capabilities at lower operational costs

We foresee enhanced QoS mechanisms utilising the generally information-centric nature of its edge network deployments, most specifically the opportunistic nature of DTN. This allows for utilising spare capacity in a less-than-best-effort service class, while providing QoS enhanced services as a differentiation. With this, the overall utilisation of the network can be increased by minimizing unused capacity throughout our system. Moreover, the information-centricity of our architecture will allow for further reduction of transfers through caching at the edge of the network, down even to individual mobile devices that operate in an opportunistic setting. Utilizing our ICN mechanisms to localise traffic or offset peak traffic against preloading content at offpeak times will further allow

network operators to provide backhaul connectivity to these local authorities at decreasing costs per bandwidth unit.

Utilizing otherwise unused capacity will also enable new models for revenue creation, such as offsetting provisioning of content with caching for future usages, therefore optimizing the usage of underutilized parts of the network etc. This in turn allows current as well as new business stakeholders to expand their revenue. For instance, we foresee the entrance of local government stakeholders into the last-mile access market, bringing Internet connectivity to users driven by social objectives. One example could be connecting local pupils to school-provided communication resources for educational purposes. Although seemingly competitive, we argue that such entrance of local stakeholder will positively complement the revenue streams of larger network operators with the cost savings that can be achieved by these local authorities through the reselling of their provisioned bandwidth and the fulfilment of given socio-economic incentives. It creates a win-win situation for both network operator and local authority, while bringing Internet connectivity to those who would likely stay disconnected.

## 5 Conclusion

The current Internet has driven a massive innovation in digital services and applications. However, many have been left out, who would like to participate in this new digital society and improve their lives. This paper not only outlined the challenges from the current Internet towards global access but also presented our vision of how to address these challenges. To this end, we presented an architectural framework that combines ICN and DTN concepts and solutions into a novel system architecture. It exposes a common information centric abstraction to applications while supporting a range of networking protocols over different transport networks, including the existing IP-based protocol suite of the current Internet.

The combination of ICN and DTN primarily opens up opportunities in temporarily and spatially decoupling the producer and the consumer of services (and the content created within these services). This in turn creates opportunities to disseminate information and services more efficiently through surrogate providers, caches and data mules, all of which makes local connectivity less dependent on an (often expensive) always-on connectivity to the wider Internet.

The support for legacy IP-based services makes our architecture compelling for existing application developers and solution providers. The support for new QoS mechanisms will drive new applications and innovations, not only striving for high-bandwidth services but also for utilising lower-than-best-effort QoS to reduce the pricing at the often so expensive backhaul in order to bring more content and more services to remote areas and communities. In this, context and content-awareness is specifically supported by the information-centricity of our architecture, which in turn leads to novel solutions for QoS management that are implemented through the dedicated resource and topology management function of our architecture.

Through such support for new QoS solutions as well as the stronger decoupling in the time and space domain, our architecture will be able to complement existing economic models for Internet access pricing with new ones that target the resale of always-on connectivity for providing lower-than-best-effort services, which can in turn be used to fill highly distributed caches and surrogate providers at significantly lower cost points. We expect such new economic models will lead to new forms of stakeholder engagements resulting in public and private partnerships that will bring broadband to those who could otherwise not afford it.

Beyond the presented architecture and its core concepts, we are currently in the process of not only implementing this architecture but also evaluating many of our claimed benefits through scaled

evaluations in an international testbed as well as in a real-life trial deployment. The latter is particularly key to shedding more light on acceptance at the user level as well on economic viability, particularly through an engagement with local authorities.

## Acknowledgments

This work is partly supported by the RIFE project, funded within the EC Horizon 2020 programme under grant agreement no. 644663.

## 6 References

- [1] Point webpage. <http://www.point-h2020.eu/>.
- [2] Internet usage, internet live stats. <http://www.internetlivestats.com/internet-users/>, 2015.
- [3] Pursuit webpage. <http://www.fp7-pursuit.eu/PursuitWeb/>, 2015.
- [4] M. Al-Naday, M. Reed, D. Trossen, and K. Yang. Information resilience: source recovery in an information-centric network. *Network, IEEE*, 28(3):36–42, May 2014.
- [5] K. Fall. A delay-tolerant network architecture for challenged internets. In *Proceedings of the 2003 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications, SIGCOMM '03*, pages 27–34, New York, NY, USA, 2003. ACM.
- [6] J. Fildes. Un reveals global disparity in broadband access. <http://www.bbc.co.uk/news/technology-11162656>, 2010.
- [7] G. Garcia and et al. Comet: Content mediator architecture for content-aware networks. *Future Network and Mobile Summit*, June 2011.
- [8] A. Ghodsi, S. Shenker, T. Koponen, A. Singla, B. Raghavan, and J. Wilcox. Information-centric networking: Seeing the forest for the trees. In *Proceedings of the 10th ACM Workshop on Hot Topics in Networks, HotNets-X*, pages 1:1–1:6, New York, NY, USA, 2011. ACM.
- [9] V. Jacobson, D. K. Smetters, J. D. Thornton, M. F. Plass, N. H. Briggs, and R. L. Braynard. Networking named content. In *Proceedings of the 5th International Conference on Emerging Networking Experiments and Technologies, CoNEXT '09*, pages 1–12, New York, NY, USA, 2009. ACM.
- [10] D. Kutscher and et al. Icn research challenges. Technical Report draft-irtf-icnrg-challenges-01, IRTF, February 2015.
- [11] N. Laoutaris, G. Smaragdakis, R. Stanojevic, P. Rodriguez, and R. Sundaram. Delay-tolerant bulk data transfers on the internet. *Networking, IEEE/ACM Transactions on*, 21(6):1852–1865, Dec 2013.
- [12] G. Parisi and D. Trossen. Filling the gaps of unused capacity through a fountain coded dissemination of information. In *Proceedings of the 2013 ACM MobiCom Workshop on Lowest Cost Denominator Networking for Universal Access, LCDNet '13*, pages 35–40, New York, NY, USA, 2013. ACM.
- [13] J. Rajahalme, M. Särelä, K. Visala, and J. Riihijärvi. On name-based inter-domain routing. *Comput. Netw.*, 55(4):975–986, Mar. 2011.
- [14] A. Sathiseelan. Researching the global access to the internet for all (gaia). *IETF Journal*, July 2014.
- [15] A. Sathiseelan and J. Crowcroft. Lcd-net: Lowest cost denominator networking. *SIGCOMM Comput. Commun. Rev.*, 43(2):52–57, Apr. 2013.
- [16] A. Sathiseelan, D. Trossen, I. Komnios, J. Ott, and J. Crowcroft. Information centric delay tolerant networking: An internet architecture for the challenged. Technical Report UCAM-CL-TR-841, Univ of Cambridge, 2014.
- [17] A. Sathiseelan, D. Trossen, J. Ott, and J. Crowcroft. An internet architecture for the challenged. *IAB ITAT*, December 2013.
- [18] K. Scott and S. Burleigh. Bundle protocol specification, ietf rfc 5050. Technical Report <http://www.ietf.org/rfc/rfc5050.txt>, IETF, 2007.
- [19] M. Skjogstad and T. Maseng. Low complexity set reconciliation using bloom filters. In *Proceedings of the 7th ACM SIGACT/SIGMOBILE International Workshop on Foundations of Mobile Computing, FOMC '11*, pages 33–41, New York, NY, USA, 2011. ACM.
- [20] V. Sourlas, G. S. Paschos, P. Mannersalo, P. Flegkas, and L. Tassioulas. Modeling the dynamics of caching in content-based publish/subscribe systems. In *Proceedings of the 2011 ACM Symposium on Applied Computing, SAC '11*, pages 478–485, New York, NY, USA, 2011. ACM.
- [21] D. Trossen and G. Parisi. Designing and realizing an information-centric internet. *Communications Magazine, IEEE*, 50(7):60–67, July 2012.
- [22] D. Trossen, M. Reed, J. Riihijärvi, M. Georgiades, G. Xylomenos, and N. Fotiou. Ip over icn – the better ip? *EuCNC (European Conference on Networks and Communications)*, June/July 2015.
- [23] L. Wang, B. Yu, and J. Manner. Proxies for energy-efficient web access revisited. In *Proceedings of the 2nd International Conference on Energy-Efficient Computing and Networking, e-Energy '11*, pages 55–58, New York, NY, USA, 2011. ACM.
- [24] G. Xylomenos and B. Cici. Design and evaluation of a socket emulator for publish/subscribe networks. In *Proceedings of the Third Future Internet Conference on Future Internet, FIS'10*, pages 11–19, Berlin, Heidelberg, 2010. Springer-Verlag.