Technical Report

UCAM-CL-TR-841
ISSN 1476-2986

Number 841

UNIVERSITY OF
CAMBRIDGE

Computer Laboratory

Information centric delay tolerant networking: an internet architecture for the challenged

Arjuna Sathiaseelan, Dirk Trossen,
Ioannis Komnios, Joerg Ott, Jon Crowcroft

September 2013

15 JJ Thomson Avenue
Cambridge CB3 0FD
United Kingdom
phone +44 1223 763500

http://www.cl.cam.ac.uk/

Technical reports published by the University of Cambridge
Computer Laboratory are freely available via the Internet:

http://www.cl.cam.ac.uk/techreports/

ISSN 1476-2986

Information Centric Delay Tolerant Networking: An Internet Architecture for the Challenged

Arjuna Sathiaseelan, Dirk Trossen, Ioannis Komnios, Joerg Ott, Jon Crowcroft

Abstract

Enabling universal Internet access is one of the key issues that is currently being addressed globally. However the existing Internet architecture is seriously challenged to ensure universal service provisioning. This technical report puts forth our vision to make the Internet more accessible by architecting a universal communication architectural framework combining two emerging architecture and connectivity approaches: Information Centric Networking (ICN) and Delay/Disruption Tolerant Networking (DTN). Such an unified architecture will aggressively seek to widen the connectivity options and provide flexible service models beyond what is currently pursued in the game around universal service provisioning.

1 Introduction

The Internet Societys recent global Internet survey reveals that the Internet should be considered as a basic human birth right like clean water, public roads, work/school etc., because of its societal benefits [1]. However, in the reality of todays Internet, the vision of universal service provisioning faces the challenge of a growing digital divide, i.e., a growing disparity between those with sufficient access to the Internet and those who cannot afford access to universal services.

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 to these areas. In addition to the physical limitations of terrestrial infrastructures (mainly due to distance) to provide last mile access, remote communities also incur higher costs for connection between the exchange and backbone network when using wired technologies. A large exchange may accommodate many users and allow for competition between service operators; in contrast, a rural/remote broadband often does not offer economies of scale, raising the costs per user. Thus, although service requirements for customers in rural/remote areas and cities are identical, the delivery mechanism needs to be different.

This calls for the need for an Internet architecture that seamlessly integrates multiple transmission technologies to reduce transmission costs, increase efficiency, flexibility and dependability.

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 [2]. This problem is even encountered in developed countries, where many individuals find themselves unable to pass a necessary credit check or live in circumstances that are too unstable to commit to lengthy broadband contracts.

The current economic models for accessing the Internet build on the basic Best-Effort model (which would be a paid users basic Service Level Agreement (SLA)), and the transport protocols that govern the transmission of data were adapted to suit the Best Effort nature of the Internet to contend for available resources. A potential solution for providing free Internet access to all would be for network operators to offer a very basic less-than-best-effort service (scavenger class) utilising the spare capacity in their networks while providing flexible Quality of Service (QoS) allowing enhanced services. Enhanced services such as micropayments or reverse payment models can be enabled by providing higher QoS for a period of time such services would allow remote doctors (receivers) to pay for additional capacity (in the return link) for video conferencing with patients. Such methods will provide better utilization of infrastructure resources and new revenue streams for operators, which may in turn reduce operating expenditures, thus providing an incentive for operators to provide low cost/free Internet access services to under privileged communities.

This calls for the need for an Internet architecture that can inherently provide flexible QoS architectural support that can enable the network operator to introduce a diverse set of new low cost services.

The current Internet is architected in such a way that end-to-end always on connectivity to transmit information is required. This situation creates a barrier for implementing time-windowed or time-shifted access, which could bring in new lower cost access opportunities. [3] 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.

Moreover, 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 devices network interface controller in idle or sleep mode as long as possible without violating the applications' time constraints.

These set of requirements clearly depict the need for the current Internet architecture to inherently support delay tolerance.

With the growing need for accessing more content, the current host-centric model of the Internet architecture leads to wastage of resources, such as redundant transmission, leading to congestion, waste of energy etc., and underutilization of opportunities to cache content on- as well as off-path, for e.g. utilising unused capacity [3]. By enabling the Internet architecture to deliver content from locations closer to the end-user, better service quality can be provided at lower costs, increasing the competitiveness of network operators. Also by enabling a framework for mapping subscribers to interests, could allow shifting demands in time and space. This increased content delivery efficiency can also result in significant energy savings for the network operator.

This calls for the need for the Internet architecture to inherently support an information-

centric model that places information at the heart of communication.

This widening range of requirements imposed on the Internet architecture leads to a proliferation of uncoordinated solutions [4]. Each of these solutions addresses only a limited set of requirements, a subset of the total problem, while accelerating the technological fragmentation.

Taking all of these limitations into account, we propose in this position paper, an Internet architecture that can ensure universal coverage by traversing the entire range of connectivity options through a single unifying communication architecture with a single set of abstractions. Such an architecture not only spurs innovation for a wide range of new services and applications, but also encompasses existing successful Internet services. We make use of advances in the area of Information Centric Networking (ICN) [5] and its inherent ability to push content to the edges, providing more localised access to important content and reducing access cost per bit through the enablement of a transmit-when-needed policy. The concept of overarching ICN enables us to pursue multiple complementary connectivity options, specifically including Delay Tolerant Networking (DTN) [6], through the notion of a dissemination strategy that constitutes an optimal set of protocols that efficiently utilise the local resources. The integration of multiple concurrent dissemination strategies enables the utilisation of both connected and disconnected modes of access under a single architectural abstraction. With this inclusion, we can accommodate a pure IP-based world, a Content Delivery Network (CDN)-enhanced web world or a connectivity-challenged DTN world, all within a single unified architectural framework.

2 Key Technologies

Information Centric Networking (ICN) has been increasingly attracting attention in the wider research community, fuelled by research efforts in various parts of the world. DONA [7] was one of the first clean-slate ICN proposals. DONA uses flat, self-identifying and unique names for information objects and binds the act of resolving requests for information to locating and retrieving information. Content-Centric Networking (CCN) [8] proposes a name-based routing system for locating and delivering named data packets. The fundamental entities in CCN are Interest and Data packets. CCN uses names to identify content objects only; that is, there is no notion of host name, point of attachment or path identifier. COMET [9] closely follows CCN, differentiated in two ways:

1) it modifies the IPv4 packet header in order to make the architecture backwards compatible and

2) it reduces the state maintained by routers via caching: a routing table has a finite size and a router stores the most recently used entries.

In PURSUIT [10][5] and PSIRP [11], a publish/subscribe system is proposed based on dividing the architecture into the core functions of rendezvous, topology management and forwarding. CONIC [12] is a network architecture designed for efficient data dissemination using storage and bandwidth resources in end-systems (i.e., available storage in end hosts is used for caching).

As we move from the core towards the edges of the network, many opportunities for performance and efficiency improvements can be exploited through ICN. One such opportunity is caching in the last mile. With the huge number of devices with spare memory and bandwidth existing at the periphery right next to the users, the ICN paradigm provides the opportunity of caching inside the network, because it names content instead of end-hosts and, thus, abstracts the identity of the content from the location where the content (or its copies) is kept.

Delay Tolerant Networking (DTN) [6, 13] is an emerging technology for a new era in interoperable communications. Like IP, DTN operates on top of existing link layer and network protocols and technologies, creating a DTN overlay network. The key advantage over IP is that DTN effectively copes with long delays, high error rates and prolonged link disruptions, thus allowing the interconnection of networks with very diverse characteristics. In particular, DTN extends internetworking in the time domain: rather than relying on contemporaneous connectivity on all segments of an end-to-end path as IP networks do, DTN operates in a store-and-forward fashion: intermediate nodes assume temporary responsibility for messages and keep them until an opportunity arises to forward them to the next hop. While stored, messages may even be physically carried within a node as the node is transported: the model is sometimes termed store-carry-forward to reach the next hop. This inherently deals with temporary disconnections or disruptions and allows for eventual connection, over time, among nodes that would be disconnected in space at any point in time, by exploiting time-space paths.

In essence, DTN technology enables seamless communication between diverse devices by hiding the complexity, the diversity, and the potential discontinuity of the heterogeneous end-to-end communication fabric from the communication service. This is accomplished by an inherently asynchronous interaction service offered to applications together with the use of numerous underlying convergence layer protocols to map to link/network layers. These convergence layer protocols offer a communication framework independent of the type of device, allowing for communication among a great variety of devices that range from common sensors and smart (mobile) devices to deep space sensors and embedded routers.

3 Information Centric Delay Tolerant Network (I-DTN)

In this section, we propose our Information Centric Delay Tolerant Network (I-DTN) architectural framework. The main goal of the proposed unified I-DTN architectural framework is to efficiently exploit all possible communication opportunities, from fixed or mobile broadband networks to disruptive networks and satellite links, while providing a unified abstraction to application developers for supporting current Internet-based services and enabling innovative future solutions.

The framework of our I-DTN architectural framework combines IP, ICN and DTN solutions into a novel system architecture, exposing a common information-centric abstraction to applications, while supporting a range of networking protocols over different transport networks. Figure 1 presents our architectural framework. ICNs publish/subscribe paradigm allows for shifting demands in time and space, since expressions of interest (subscriptions) can be satisfied long after they have been issued, from any entity that has a copy of an object that matches the interest. This is in tandem to the philosophy of DTN. DTN and ICN are complementary in many aspects. For instance, DTN deals with

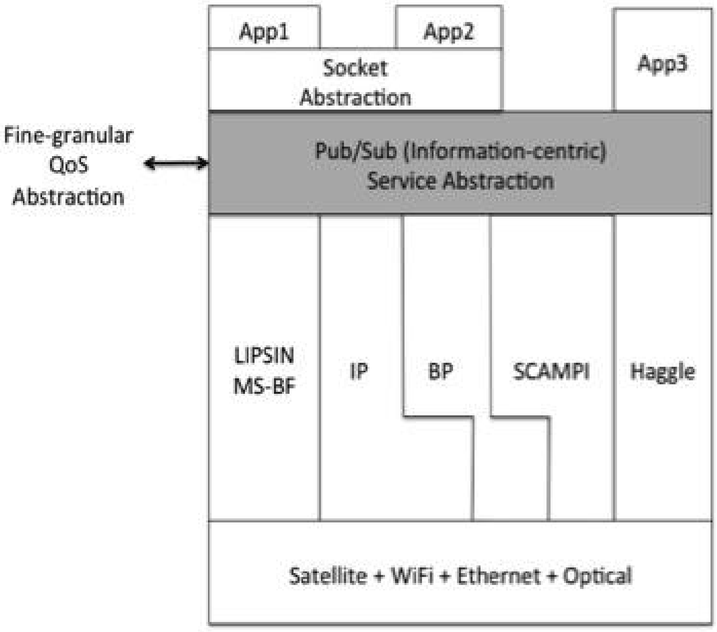

Figure 1: I-DTN Architecture

discrete data items and operates hop-by-hop, optionally requiring intermediate nodes to retain content until reception has been confirmed. This allows for time and space shifting of when and where resources are used. This 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 framework 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.

3.1 Design Tenets

The design principles of I-DTN are described below:

(1) We design the service abstraction that is provided to applications by defining an information model, as well as a service model, that is exposed to them. We utilise existing DTN and ICN solutions (see Section 3.2) as a basis for this common abstraction, providing an object-level graph-based information abstraction. Information is split into several items or objects and each such object is associated with a context (also known as scoping). Scope represents sets of information. Both information objects and scopes are represented as directed acyclic graphs (DAG) manipulated through a set of publish/subscribe operations. While we expect applications to natively utilise this common information-centric interface of the architectural framework, we also foresee interfaces being defined that allow, for example, socket emulation [15] that would enable backward compatibility.

(2) We functionally decompose the network components using PURSUIT ICN [5] and existing DTN (Bundle Protocol [14]), into three core functions (see Section 3.2), namely

rendezvous, topology management and forwarding. The functional decomposition also addresses the interaction with the underlying networks, such as satellite, cellular, WiFi or optical networks. This is accomplished mainly through the topology management function, which manages the resources available in the form of links, spectrum, wavelength but also storage and computational capability.

(3) Based on our decomposition, we define the interfaces between the core components of our architectural framework, e.g., for initiating discovery requests, assembling network resources for store-and-forward operations or forwarding information objects over paths that were assembled through the topology management function. These interfaces are realised through various dissemination strategies that enable traversal across the various connectivity options, e.g., over challenged and opportunistic network environments (using DTN), IP-based backhauls (IP being used as a 'framing' (link layer) protocol) or using native ICN (LIPSIN [17]) for high speed optical links.

Our concept of dissemination strategies within our architectural framework is the underlying foundation for enabling autonomous operation within each dissemination strategy (e.g., across a DTN-enabled network), while preserving the common end-to-end nature of communication that is exposed to the application through the unified information-centric service interface.

3.2 Choosing PURSUIT ICN

Our I-DTN architectural framework is based on the PURSUIT ICN architecture [5] for delivering information centric services, for the following reasons: Both PURSUIT ICN and DTN (Bundle Protocol) [14] share an object approach to information (while, for instance, CCN operates at packet level). PURSUIT ICN separates core network functions into three:

1. *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.

2. *Topology management and formation*: realizes the management of the overall delivery topology and the formation of specific delivery graphs.

3. *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.

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.

In I-DTN, DTN represents a particular dissemination strategy (in this case, for 'challenged networks'), which coexists with other strategies that are, for instance, highly optimised for optical high-speed networks. This strategy concept is crucial for spanning the expected spectrum of connectivity options, in order to aggressively acquire and utilise any connectivity option available. I-DTN complements these dissemination strategies through solutions for caching and replicating information and content, in particular in the edge networks. Replication is done using set reconciliation and network coding to proactively push and update information toward network nodes that may later need, while caching uses local algorithms to opportunistically store received data. Due to the fine-grain granularity of object representation of information, several objects can be fetched from several sources over multiple paths. In this way, we improve discovery times for information requests and utilise replication of information to improve 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.

3.3 Enabling flexible QoS

In broadband access networks, access is normally shared among a pool of users. Elastic provisioning of resources is vital due to the high costs of fixed assigned capacity for sporadic and variable communication demands. Our architectural framework satisfies this requirement providing more fine-grained QoS. Such a method presents immense benefits by enabling a wider breadth of services than those directly available from the standard Service Level Agreement (SLA). In order to support fine-grained, i.e. information object-specific QoS, I-DTN provide interfaces that allow for defining and manipulating the QoS parameters in alignment with the service model that is exposed to the applications. This is achieved through the modeling of QoS as an algorithmic relation to the information object it is associated with. This allows the same information abstraction and service interfaces that are used for the original information object to be used for its QoS. The core network function that is responsible for network resource management, i.e. the topology management function, can utilise the algorithmic relation for a fast lookup of QoS parameters at the time of a transfer request providing late-binding. Such algorithmic relations between information object and QoS policy can dynamically change based on changing environmental conditions (such as faults in the network or availability of spare resources), changed business arrangements (such as maxing out a particular resource quota), or changing application conditions (such as the detection of an abnormal condition in the application context).

Our architectural framework provides extensions to support applications that want to preserve IP abstractions while becoming content-aware. An example would be an IP-based camera that exercises Expedited Forwarding (EF) QoS policy when its image recognition detects movement. Here, an IP abstraction over ICN is utilised for backward compatibility, while providing a QoS control interface for IP-flow-based QoS; this control interface, however, utilises the ICN QoS subsystem, enabling fine-grained control of QoS policies per IP node.

4 Conclusion

In this technical report, we present I-DTN, an architectural framework combining ICN and DTN solutions into a novel system architecture, exposing a common information-centric abstraction to applications while supporting a range of networking protocols over different transport networks.

Transfer over legacy IP (as a transport network, not an application-facing abstraction) will still be a predominant mode in our architecture. We plan to develop solutions for using IP as a framing protocol to forward packets between links. For this, we intend to investigate the suitability of placing such framing on top of raw IP packets or utilise transport layer framing solution, e.g., over TCP.

The introduction of a new information-centric abstraction to application developers is expected to be a huge driver for new applications. However, we expect that many legacy IP applications will continue to run over our architecture as well. We also envision that, increasingly, such legacy applications will desire to utilise the QoS features that are introduced by our architecture. Hence, we must enable backward compatibility, as well as support for these new features. Our approach is based on enabling both content-awareness and context-awareness in the application. While we will use initial solutions to provide backward compatibility via the socket abstraction, we will develop interfaces and solutions for signaling the requirements of the content- and context-aware IP applications to the stack. These requirements will be then incorporated into the core functions realised by the dissemination strategy and, particularly, the topology management function.

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 through a minimization of unused capacity throughout our system, while the information-centricity of our architectural framework 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. These technical solutions for transport efficiency through economic models can complement current economic models for broadband provisioning with new forms of stakeholder engagements resulting in public and private partnerships that will bring broadband to those who could otherwise not afford it.

We are currently in the process of implementing this architecture incorporating the mechanisms presented in this paper. We intend to evaluate the architecture through both testbed emulations and deployments in-the-wild.

References

[1] Internet Society, Global Internet User Survey Reveals Attitudes, Usage, and Behavior, https://www.internetsociety.org/news/global-internet-user-survey-reveals-attitudes-usage-and-behavior, 2012.

[2] J. Fildes, UN reveals global disparity in broadband access, BBC, http://www.bbc.co.uk/news/technology-11162656, 2010.

[3] N. Laoutaris, P. Rodriguez, Good things come to those who (can) wait of how to handle delay tolerant traffic and make peace on the Internet, ACM HotNets-VII, 2008.

[4] M. Honda, Y. Nishida, C. Raiciu, A. Greenhalgh, M. Handley, H. Tokuda, Is it still possible to extend TCP?, ACM SIGCOMM IMC, 2011.

[5] D. Trossen, G. Parisis, Designing and Realizing an Information-Centric Internet, Communications Magazine, IEEE, vol.50, no.7, pp.60,67, July 2012.

[6] K. Fall, A Delay-Tolerant Network Architecture for Challenged Internets, IRB-TR-03-003, February 2003.

[7] T. Koponen, M. Chawla, B.-G. Chun, A. Ermolinskiy, K. H. Kim, S. Shenker, and I. Stoica, "A Data-Oriented (and Beyond) Network Architecture," In Proc. of SIGCOMM 07.

[8] V. Jacobson, D. K. Smetters, J. D. Thornton, M. F. Plass, N. H. Briggs, R. L. Braynard (PARC), Networking Named Content, CoNEXT 2009, Rome, December 2009.

[9] A. Detti, N. Blefari-Melazzi, S. Salsano, and M. Pomposini, "CONET: A Content Centric Inter-Networking Architecture", in ACM SIGCOMM workshop in Information-centric networking, 2011.

[10] EC PURSUIT, http://www.fp7-pursuit.eu/, 2013.

[11] EC PSIRP project, http://www.psirp.org/.

[12] Y. Zhu, M. Chen, and A. Nakao, "CONIC: Content-Oriented Network with Indexed Caching," INFOCOM IEEE Conference on Computer Communications Workshops, 2010.

[13] V. Cerf, S. Burleigh, A. Hooke, L. Torgerson, R. Durst, K. Scott, K. Fall, and H. Weiss, Delay-tolerant networking architecture, IETF, RFC 4838, Apr. 2007, Available: http://www.ietf.org/rfc/rfc4838.txt.

[14] K. Scott, S. Burleigh, "Bundle Protocol Specification", IETF RFC 5050, experimental, November 2007, http://www.ietf.org/rfc/rfc5050.txt.

[15] G. Xylomenos, B. Cici, Design and Evaluation of a Socket Emulator for Publish/Subscribe Networks, Proc. of the Future Internet Symposium, 2010.

[16] N. Laoutaris , G. Smaragdakis , P. Rodriguez , R. Sundaram, Delay Tolerant Bulk Data Transfers in the Internet, Proc of ACM Sigmetrics, 2009.

[17] P. Jokela, A. Zahemszky, C. Rothenberg, S. Arianfar, P. Nikander, LIPSIN: line speed publish/subscribe inter-networking, ACM SIGCOMM, 2009.