An Internet Architecture for the Challenged

Arjuna Sathiaseelan[1], Dirk Trossen[1], Ioannis Komnios[2], Joerg Ott[3], Jon Crowcroft[1]
[1]Computer Laboratory, University of Cambridge
[2]DUTH, Greece
[3]Aalto University, Finland
{arjuna.sathiaseelan,dirk.trossen,jon.crowcroft}@cl.cam.ac.uk, ikomnios@spice-center.org, jo@netlab.tkk.fi

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 short paper 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

Enabling universal Internet access is a key issue that is currently being addressed globally. However the existing Internet architecture is seriously "challenged" (geographically, socio-economically and technically) to ensure universal coverage.

Geographic Challenge: Access problems often result from sparsely spread populations living in physically remote locations – it is simply not cost effective for Internet Service Providers (ISPs) to install the required infrastructure for broadband Internet access to these areas. This problem is widely and publically recognized. In the UK alone, a third of the population lack broadband access – typically those living in rural or remote locations.

Socio-Economic Challenge: Addressing digital exclusion due to socio-economic barriers is also important. The United Nations revealed the global disparity in fixed broadband access, showing that access to fixed broadband in some countries costs almost 40 times the national average income. This problem is also applicable to developed countries where many individuals find themselves unable to pass a necessary credit check, or living in circumstances that are too unstable to commit to lengthy broadband contracts. Indeed, Internet services are increasingly accessed on the move and so current models of "roaming" access provision drive this economic exclusion to a new level, not currently addressed by the push to deploy broadband. We also believe that leaving connectivity for all to be governed by market economics is a major impediment to achieving the full benefits of the Internet, and that basic Internet access should be made freely available to all due to its societal benefits, a sentiment recently expressed by Berners-Lee.

Technological Challenge: The current Internet architecture is progressively reaching a saturation point in meeting increasing user's expectations, and it is progressively showing inability to efficiently respond to new challenges:

(1) The current economic models for accessing the Internet build on the basic Best-Effort model (which would be a paid user's 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. This makes it impossible to support service models that could lower the cost of Internet access, e.g., through adaptive QoS. This is vital to the notion of Universal access reducing the economic barriers for access.

(2) The current Internet is architected in such a way that it requires end-to-end always-on connectivity to transmit information; receivers are required to be prepared to accept data at any time, never knowing when transmissions will be infrequent and when more energy-efficient operating modes may be safely entered. This scheduling uncertainty therefore wastes energy, raising a barrier to implementation of time-windowed – or time-shifted – access which could bring in new lower cost access opportunities for, e.g., utilising unused capacity [1].

(3) With the growing need for accessing more content, the host-centric model of the Internet leads to waste of resources, such as redundant transmission leading to congestion, while wasting opportunities to cache content on- as well as off-path.

This widening range of requirements imposed on the Internet architecture leads to a growing collection of solutions (the fragmentation of the Internet). Each such solution addressed some set of requirements, a subset of the total problem, while accelerating the fragmentation.

Through this workshop and paper, we hope to initiate a discussion on understanding these challenges and discuss a potential architectural solution, which we briefly describe in the next section. We also hope to understand through the workshop the challenges such an architecture would face in terms of wider adoption.

2. Proposed Solution: Lowest Cost Denominator Networking (LCDNet)

We propose to address these challenges with the vision of making the Internet ubiquitous, accessible and energy-efficient. We do this by traversing a range of connectivity options that ensure universal coverage, while providing 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 encompasses existing successful Internet services. We utilise advances in information-centric networking (ICN) [2] to provide this abstraction - an abstraction driven by access to and provisioning of information rather than the connection to explicitly identified endpoints.

Through this abstraction, we accommodate today's web-based services while providing a path to future immersive and sensor-rich applications, such as those envisioned by the Internet of Things. The focus on information enables inherent support for rich caching policies that ultimately increase the efficiency of the network across different technologies by providing information from the most efficient provider rather than the original source.

The concept of overarching ICN enables us to pursue multiple complementary connectivity options, specifically including Delay Tolerant Networking (DTN) [3], as distinct *dissemination strategies, each of which* constitutes a set of protocols that optimally utilise local resources. Integrating multiple concurrent dissemination strategies enables the utilisation of connected and disconnected modes of access under a single architectural (information-centric) abstraction. This enables us to accommodate a pure IP-based world as much as a challenged connectivity DTN world, all within a single architecture, while exploiting all possible communication opportunities that particular network deployments provide, ranging from fixed, all-optical deployments of wireless and mobile networks to satellite-supported deployments that cover difficult-to-reach environments. Figure 1 summarizes our architecture.

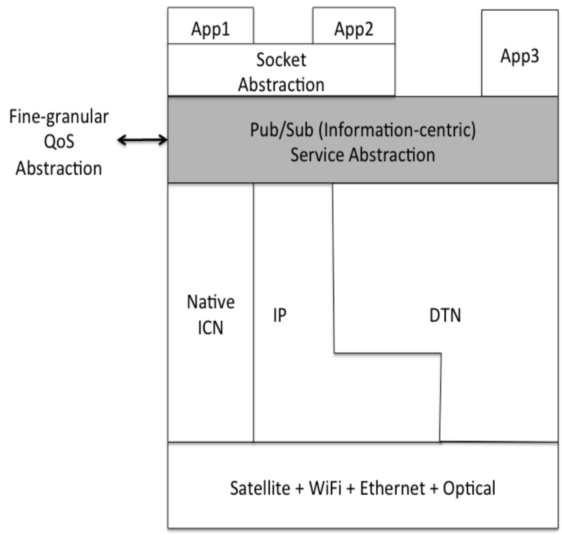

Figure 1: LCDNet Architecture

The design tenets of our LCDNet architecture 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 will utilise existing DTN and ICN solutions as a basis for this common abstraction, providing an object-level graph-based information abstraction, manipulated through a set of publish-subscribe operations. While we expect applications to natively utilise this common information-centric interface of the architecture, we also foresee defining interfaces allowing, for example, socket emulation that would enable backward compatibility.

(2) We functionally decompose the network components, using the PURSUIT ICN work [2] and the existing DTN work, into three core functions, 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, and wavelength but also storage and computational capability.

(3) Based on our decomposition, we define the interfaces between the core components of our architecture, 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 will be realised through various dissemination strategies that enable traversal across the various connectivity options, e.g., over challenged and opportunistic network environments, high-speed networks or IP-based backhauls.

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. Envisioned benefits of our architecture

(1) Our architecture will enable new models for revenue creation from currently underutilized infrastructure (offsetting provisioning of content with caching for future usages, therefore optimizing the usage of underutilized parts of the network etc.), allowing current business stakeholders to expand their revenues. By providing tools to manage and deliver content 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 is expected to result in significant energy savings for the network operator.

(2) Our architecture's ability to couple DTN with more localized Internet access using ICN and other opportunistic communication approaches will enable economically disadvantaged users to exploit such lower-quality time-windowed or time-shifted services. This will also introduce new access models for the Internet of Things.

(3) The information-centricity of our architecture is key to extending from the current host-centric to the future Internet that connects people, content, computational resources and things under a single abstraction on top of which developers can provide solutions. We foresee our basic architecture to be extended with solution enablers that take advantage of our rich information abstraction and its underlying mechanisms that achieve context- and content-awareness across all connectivity options.

(4) 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 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. 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.

(5) Elastic provisioning of resources is vital due to the high costs of fixed assigned capacity for sporadic and variable communication demands. Our architecture will satisfy this requirement to provide the ability to vary QoS parameters and user capacity on demand or as needed. In order to support fine-grained, i.e. information object-specific QoS, our architecture will 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). The flexible QoS support provided by the architecture will enable the network operator to introduce a diverse set of new services including the introduction of micropayments, reverse payment models that can enable higher QoS for a period of time and will allow businesses to transfer files, enable remote doctors to pay for additional capacity for video conferencing with patients, etc. Our architecture 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.

4. Acknowledgements

We would like to thank Scott Burleigh for providing us with valuable feedback and helping us shape our thoughts. This work was partly supported by the EPSRC Grant EP/K012703/1.

5. References

[1] A. Sathiaseelan, J. Crowcroft, LCD-Net: Lowest Cost Denominator Networking. ACM Computer Communication Review, April 2013.
[2] D. Trossen, G. Parisis, "Designing and Realizing an Information-Centric Internet", Communications Magazine, IEEE , vol.50, no.7, pp.60,67, July 2012.
[3] K. Fall, A Delay-Tolerant Network Architecture for Challenged Internets, IRB-TR-03-003, February 2003.