JOURNAL OF INFORMATION POLICY 2 (2012): 26-50.

TECHNO-ECONOMIC ASPECTS OF INFORMATION-CENTRIC NETWORKING

BY DIRK TROSSEN[*] AND ALEXANDROS KOSTOPOULOS[†]

Can the Internet be redesigned to reduce future conflicts? The Internet's underlying architecture, Internet Protocol (IP), was introduced in 1974. Since then many ideas have been put forward about how to update and improve it. One branch of these is called "Information-Centric Networking" (ICN). Trossen and Kostopoulos note how ICN could improve the ability of the Internet to resolve conflicts between the various constellations of stakeholder interests, conflicts that they call "tussles." Introducing a "tussle taxonomy," they provide examples of how tussles might be resolved differently in ICN. They believe the ICN model would help rationalize pricing in a three-sided market; reduce congestion and transit costs; provide more transparency; offer more choices and possible outcomes with respect to issues such as privacy, intellectual property, and data protection; and better enable not just present but future business models that actors within the system might strive to establish.

INTRODUCTION

In a global communication system like the Internet, conflicts between adversaries are inevitable. Such conflicts can be driven by economic as well as political interests but also by the desire of individuals to express themselves in the many forums that the Internet provides. It has long been recognized that the nature of these conflicts has a direct impact on the viability of various designs of the Internet in general and many design decisions in particular. Such recognition plays an important part not only in today's Internet but even more so in any effort that aims at designing new Internet architectures. One such exemplary effort is "Information-Centric Networking" (ICN). In this article, we look at the economic aspects of such architecture from the viewpoint of conflicts that may unfold between various parties. We investigate how an information-centric Internet can improve how such conflicts are addressed through an increased modularity of functions. We present our work along a set of use cases, directly inspired by a taxonomy of players' conflicts laid out in the

---

[*] Senior researcher, Computer Laboratory, University of Cambridge. The research of Dirk Trossen is supported by the EU FP7 project "PURSUIT" under grant FP7-INFSO-ICT 257217.
[†] Doctoral candidate, Department of Informatics, Athens University of Economics and Business. The research of Alexandros Kostopoulos has been co-financed by the European Union (European Social Fund – ESF) and Greek national funds through the Operational Program "Education and Lifelong Learning" of the National Strategic Reference Framework (NSRF) – Research Funding Program: Heracleitus II, Investing in Knowledge Society Through the European Social Fund.

next section. Since ICN brings challenges for the existing players as well as possible changes imposed upon them, our analysis provides insights into the future of the Internet.

There has been increasing interest in redesigning the IP (Internet Protocol) layer of the Internet. However, Internet evolution and its effects on participants' interests have triggered the need for redefining its current design. A particular branch of efforts, such as NDN,[1] PURSUIT,[2] and others, declares *information* to be a first-class citizen at the networking level, building transient relationships between providers and consumers of information at any point in time. We group these efforts as *information-centric networking* in this article.

Core to these proposals is the recognition that the *what* within a communications relationship is more important than *who* is communicating. Supported by technological developments in computing and storage resources, these efforts recognize that the *what* of a communication scenario is likely to exist in many more places than the originally addressed *who*. But we can go even further beyond this observation and create a link between information dissemination and realization of distributed computational tasks. We argue that available storage and computing resources within a distributed environment are utilized towards implementing such tasks. Hence, it becomes the role of an information-centric network to facilitate the dissemination of any information pertaining to those tasks, while optimizing the particular implementation of this facilitation within the realm in which the information is disseminated. This makes sub-architecture optimization a crucial aspect of information-centric networking, an issue that will be important in the work presented in this article.

But any such radical change to the design, the provided abstractions, and the resulting implementations at this core layer of the current Internet require careful thinking as to what their potential benefits might possibly be. A set of desirable architectural claims that would motivate such fundamental change to today's Internet has been proposed in the research literature.[3] We focus on improving the delineation of tussles along well-defined boundaries within the resulting architecture. The term "tussle" was initially inspired by Clark et al.[4] Here the Internet is symbolized as a playground, where the involved stakeholders have conflicting interests (tussles). As suggested, proper modularization along crucial lines of delineation within the overall architecture is essential for ensuring viability and adjustability of the architecture to varying socio-economic conditions. It is this improved ability to adjust to changes that is the essence of the claim of "tussles" in information-centric networking.

The separation of functions for identifying information, finding it, and finally delivering it along a suitable delivery graph within an information-centric architecture is at the heart of this claim. Furthermore, the focus on information allows for establishing information boundaries and effective

[1] "Named Data Networking," Dec. 31, 2011, accessed Feb. 16, 2011, http://www.named-data.net/.
[2] "PURSUIT," accessed Feb. 16, 2011, http://www.fp7-pursuit.eu/PursuitWeb/.
[3] Dirk Trossen, Mikko Sarela, and Karen Sollins, "Arguments for an Information-Centric Internetworking Architecture," *ACM SIGCOMM Computer Communication Review* 40, no. 2 (2010): 27-33.
[4] David D. Clark, John Wroclawski, Karen R. Sollins, and Robert Braden, "Tussle in Cyberspace: Defining Tomorrow's Internet," *IEEE/ACM Transactions on Networking* 13, no. 3 (2005): 347-356.

information asymmetries more flexibly, given the exposure of information in a different, more consistent way throughout the architecture.

This article intends to provide some insight into why an information-centric networking architecture improves the ability to accommodate various constellations of stakeholder interests. For this, we utilize an approach that is driven by dedicated use cases from an economics and business point of view. For each of these use cases, we outline the possible conflicts between major players as they exist in an IP-centric world and how these conflicts could play out in an information-centric alternative. We believe that this comparative approach based on concrete examples will provide useful insights into the tussle space analysis in an architectural context.

Before delving into the use cases, we first provide in the next section the architectural backdrop of information-centric networking. We then classify various conflicts in such system through a tussle taxonomy in the following section. This taxonomy will help us better understand the specific use cases in the next two sections, each of which has a specific focus within the architectural context of information-centric networking. We end the discussion with general policy lessons learned before concluding the article.

THE ARCHITECTURAL CONTEXT

Information-centric networking has been touted as a replacement for the traditional endpoint-centric IP networking approach of the current Internet.[5] In order to enable an understanding of the tussle analysis that is introduced later in the article, we first provide a brief introduction of this new architectural context that information-centric networking provides. We omit many of the details necessary to understand the full workings of the various proposed approaches and focus on a better understanding of certain aspects that will follow in our tussle analysis.

The intuitive starting point is that all network operations are based on information as the primary named entity across all layers. We assert that this aids the consistency of concepts across the layers and enables efficiency gains in operating over a single concept, namely that of information, across all layers. We assume that each piece of information has a statistically unique name and that applications can request that the network deliver the named information. Hence, the primary function of the network is to find an appropriate location for an information provider and to deliver information. In other words, the network emphasizes the *what* of a communication scenario, while building transient relations between those *who* might have and want the information at hand. This is significantly different than IP, which places the emphasis on the exchange of opaque bits between specifically identified endpoints, i.e. it helps to locate hosts and arranges communications between them.

[5] Trossen, Sarela, and Sollins.

In order to make the vast amount of information manageable, we introduce a concept called *scope* as a way to group related data together. From the network's perspective, a scope denotes the party responsible for locating a copy of the data. With that, it creates a point of control to implement, e.g. access control and usage policies. Each information item may reside in more than one scope. Treating a set of items as an information item then allows grouping scopes within other scopes as well. With this, the network directly operates on a directed acyclic graph[6] of information with operations that manipulate these graphs. These operations follow a *publish-subscribe* ("pub/sub") model. In other words, information is published by any provider, while it is subscribed to by anybody who is interested in it. A dedicated matching process ensures that data exchange only occurs when a match in information item and scope has been made.

This intuitive introduction into information-centric networking highlights a very important aspect of changing this paradigm of internetworking, namely the change of abstractions that are visible to applications and network nodes alike. These abstractions move from links, sockets, and endpoints to information graphs with operations to manipulate these through a pub/sub model rather than a push-like send/receive model.

Conceptual High-Level Architecture

With this change in abstractions being exposed to application and network developers alike, the conceptual architecture changes in significant ways. In order to implement the abstractions outlined above, the architecture provides the required mapping of the underlying concepts into concrete forwarding relations between endpoints, which are producing and consuming information. While this keeps the network architecture simple (and allows for separately optimizing the realization of parts of the network), it enables a growing complexity of application-level problems to be implemented on top of this simple model.

Figure 1 below represents the main architectural components on a very high level. The pub and sub components at the application level implement applications based on basic publish-subscribe network services, enabling publications and subscriptions towards information items within particular scopes. Transactional services, operating in request-reply mode, can easily be supported through a publish-subscribe model, with the server subscribing to receive requests over identifiers being created for that purpose by the application. The relation of such new API to traditional middleware layers is that it conflates low-level information discovery as well as location determination of publishers and subscribers into a single network service. This is likely to have an impact on middleware developments, an issue left out of the discussions in this article.

[6] Loops in information structures are prevented through the directed acyclic nature of these graphs, allowing for rich structures such as hierarchical trees (i.e. similar to the Internet's domain system). Directed acyclic graphs are used in mathematics and computer science and are constructed to prevent a vertex from cycling back to its starting point.

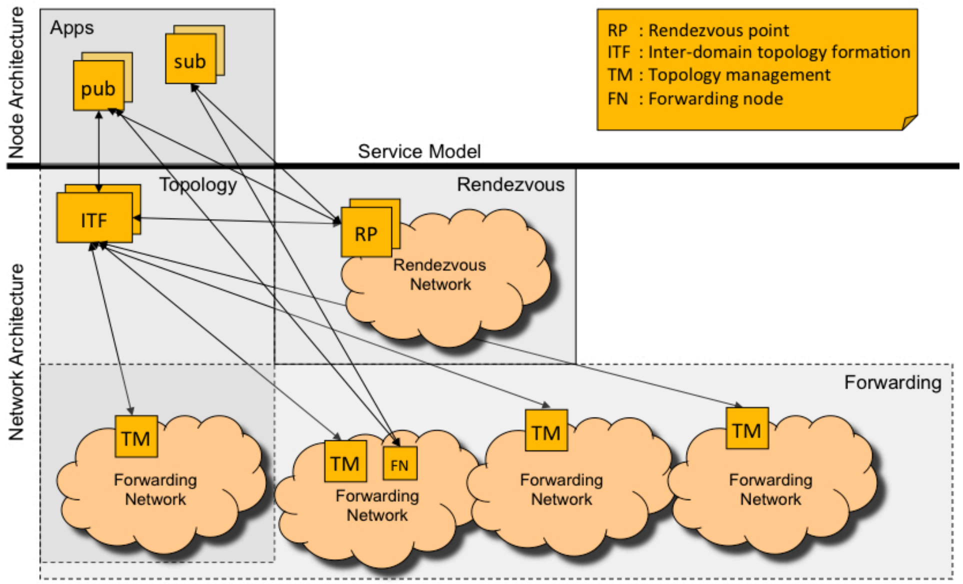

Figure 1: Conceptual Architecture

The network architecture itself consists of three main functions: *rendezvous*, *topology*, and *forwarding*. Generally, the *rendezvous* function implements the matching between publishers and subscribers of information. The matching is realized for a particular part of the overall information graph that is constructed by the application. The matching is performed by at least one rendezvous point which is directly associated with the identifier of the scope that it performs the matching over. In other words, rendezvous points match the semantic-free information items within the scope they are serving. With more than one rendezvous point possible for a scope, requests for information items within that scope can be routed either to all rendezvous points or to the "best" rendezvous point, using anycast-like functionality.[7] Furthermore, rendezvous points implement policies associated with the matching, such as access control.

Upon having matched a publication and one or more subscriptions, an inter-domain forwarding graph is created in negotiation with the inter-domain topology formation (ITF) function. This is based on some form of location for the publisher and subscriber on the level of autonomous systems (ASes). Furthermore, any applicable policies as well as peering and transit relationships among ASes are included into the operation. Hence, there exists a rich set of policies attached to potentially every information item. Unlike the Border Gateway Protocol (BGP), this approach allows for multiple ITF functions, each offering different sets of peering and transit opportunities that were exposed to them. This establishes the potential for peering markets in which the ITF providers serve as routing service providers. Choice can be achieved by ASes publishing peering and transit relations to various ITF functions, usually constrained by policies governing these relations,

[7] The term "anycast" refers to a technique in which datagrams from a single sender are routed to the topologically nearest node in a group of potential receivers that are all identified by the same destination address.

while particular (sets of) ITF functions are chosen for topology formation. The desire to separate the tussle of (policy-based) inter-domain path selection and inter-domain forwarding requires that transit ASes cannot make additional policy-based decisions on traversing packets, e.g. changing the next peering hop after the path selection decision.

After constructing inter-domain paths between the forwarding networks to which publishers and subscribers are attached, intra-domain paths need to be constructed. This is done in collaboration with the AS-internal topology management function, which instructs its local forwarding nodes (FNs) to establish paths to local publishers and/or subscribers or to serve as transfer links between ASes. As in the current Internet, we do not prescribe any particular intra-domain forwarding mechanism with the one constraint that the local mechanisms should support the traffic policies chosen by the ITF function.

Let us consider an example of content delivery in this proposed ICN architecture. The rendezvous network (RENE) will be responsible for matching requests for content between the subscriber and a publisher, realized via their local RENEs. Upon having matched a subscription and a publication of a specific information item (e.g. a video file), a forwarding graph will be created in negotiation with the topology managers (to construct the intra-domain paths), as well as with the ITF function (to establish paths between ASes). Finally, local ISPs forward the information item based on the constructed delivery relation.

A TUSSLE TAXONOMY

We now turn to the various conflicts that can occur in the architectural context we outlined. We start with examples of conflicts, some of which we will deepen in our later use cases. From these examples, we then formulate a tussle taxonomy that can guide our work on exploring the tussle space for information-centric networking.

Some Examples of Potential Tussles

Tussles about what content we want and what we get: A common problem in today's Internet is the delivery of content that users do not actually want.[8] In today's Internet, spamming has no sufficient cost (for the spammer) and still remains a common marketing tool for most advertisers. There is a tussle between end users and content providers that send spam e-mails, bulk messages or additional web pages that appear in users' browsers. Not only does this conflict with the users' interests, but it furthermore results in increasing congestion within networks and therefore increased costs for the delivery of the desired content. Although the information-centric architecture described

[8] Zoltán Gyöngyi and Hector Garcia-Molina, "Web Spam Taxonomy," paper presented at the first International Workshop on Adversarial Information Retrieval on the Web, Chiba, Japan, May 2005, accessed Feb. 27, 2012, airweb.cse.lehigh.edu/2005/gyongyi.pdf.

in the previous section addresses this conflict by introducing a publish/subscribe service notion,[9] new tussles may occur in an ICN architecture. For instance, malicious users could send fake requests to the rendezvous system of an ICN architecture, influencing the ranking system that is possibly implemented in the rendezvous point for the particular information. Such attacks are commonly known in today's Internet within ranking systems such as online shopping and the like. Hence, solutions to this problem need to be different than those in today's systems.

Tussles about what we need to expose in order to get what we want: Related to the issue of receiving wanted content, there is a recognized conflict that occurs when being required to reveal certain information in order to receive other information.[10] End users have become accustomed to gaining access to seemingly free content, albeit at the cost of revealing a plethora of information in the process of doing so. This is largely a conflict between end users and content providers, the latter gathering information about consumption on a large scale. Although data protection directives exist in various legislations, large-scale profiling is still considered as being in its infancy and the problems and impacts are still to be investigated.

Another conflict between end-users and content providers/owners is related to the so-called "piracy ecosystem." For example, the conflict around the SOPA (Stop Online Piracy Act) and PIPA (Protect Intellectual Property Act) bills in the United States Congress are representative examples of this technical and legal debate.

In both of the above tussles, we can recognize that an ICN architecture introduces a clear control point for such conflicts in the form of the rendezvous point for a particular information exchange.

Tussles about ownership of experience: Who "owns" the experience that is delivered to the end user is a conflict that already widely exists in today's Internet.[11] Users have their delivery contracts with their ISPs while often having additional agreements with content providers as well. Who "owns" the experience here? Recent research has already pointed to the problems arising in this constellation of relationships and the problems that result from the separation of opaque bit transfer at the IP level and the information exchange that is largely present in the World Wide Web today.[12] The ICN architecture described in the previous section introduces the rendezvous functionality as an intermediary between transport and end users. However, there is still a remaining tussle between the owners of actual delivery topology (represented by the topology manager function) and the

---

[9] It is important to understand that content is only delivered if a receiver indicated an interest in receiving it. Hence, it is the rendezvous point that becomes the place for mediation and therefore a crucial control point in the spamming conflict.

[10] Ian Brown, David Clark, and Dirk Trossen, "Should Specific Values Be Embedded in the Internet Architecture?" proceedings of the Re-Architecting the Internet workshop (ReARCH '10), Nov. 2010, accessed Feb. 27, 2012, http://dl.acm.org/ft_gateway.cfm?id=1921246&type=pdf&CFID=68093922&CFTOKEN=42506244.

[11] This conflict is exemplified by initiatives such as the YouView platform (accessed Feb. 16, 2012, http://www.youview.com) or the Content Delivery Networks Interconnection page at the website of the Internet Engineering Task Force (accessed Feb. 27, 2012, http://datatracker.ietf.org/wg/cdni/) where ISPs aim to reduce the impact of content providers dictating the user experience through "over the top" management.

[12] Dirk Trossen and Gergely Biczok, "Not Paying the Truck Driver: Differentiated Pricing for the Future Internet," Proceedings of the Re-Architecting the Internet workshop (ReARCH '10), Nov. 2010, accessed Feb. 27, 2012, http://dl.acm.org/ft_gateway.cfm?id=1921235&type=pdf&CFID=68093922&CFTOKEN=42506244.

owners of content (represented by the rendezvous function). Similar to today's bundled service offerings, ISPs might decide to offer rendezvous services, entering the game of brokering information in addition to delivering it. This could be countered, however, through regulatory enforcement of choice in selecting rendezvous services (similar to choosing your DNS service today). In conclusion, the conflicts are not much different but the modular boundaries, defined through the introduction of new architectural roles, could be different and therefore allow for different outcomes; something we elaborate on in our use case in the next section of this article.

Tussles about optimizing delivery networks: Related to the conflict over who owns the end user experience is that of optimizing the utilization of delivery networks.[13] One aspect of this conflict is the role of content delivery networks (CDNs). CDNs are widely used in the current Internet to optimize the delivery of content. In most deployments, large content providers pay CDNs to deliver their content more efficiently and with guaranteed latencies. ISPs collaborate with CDNs in order to perform such optimized delivery. But recent initiatives such as the UK-based YouView platform demonstrate the desire of ISPs to directly compete with CDN providers, such as Akamai, by replacing this overlay function with a natively supported function at the ISP level.[14] The prospect of offering lower prices for content distribution by directly exploiting the available infrastructure knowledge is what drives these efforts, albeit without a clear architectural basis for realization. Within the ICN architecture, the ability to directly offer a service equivalent to today's CDNs is given through the exposure of a dedicated topology manager function (see above). The boundary here lies in the interface between the rendezvous provider (representing experience requirements from the end user side and content provider side) and the topology function (representing operational requirements from the ISP side).

Tussles about interconnecting networks: As part of the aforementioned optimization tussle, there is a set of particular conflicts related to interconnecting individual transport networks.[15] One set of conflicts is that around the problem of optimizing across administrative boundaries, similar to proposals for providing inter-domain routing as a service.[16] Such optimization often requires revealing operational data, such as topology information, link and router loads, etc. – which is seen as highly confidential by the individual ISPs. Although collaborating ISPs have an incentive to be truthful about their topology in order to create win-win situations, there could be situations in which untruthful operation is seen as beneficial, resulting in an information exchange that unilaterally influences the choice of paths that are created (e.g. to shift load towards a particular ISP). We believe, however, that such an information-centric aspect enables the possibility to expose the exchanged information similar to end user content and apply similar ranking mechanisms that are used for the content itself.

---

13 Dave Clark, Bill Lehr, Steve Bauer, Peyman Faratin, Rahul Sami, and John Wroclawski, "Overlay Networks and the Future of the Internet," *Communications & Strategies* 63, 3rd q. (2006): 1-21.
14 "YouView," accessed Feb. 16, 2012, http://www.youview.com/.
15 Matthew Caesar and Jennifer Rexford, "BGP Routing Policies in ISP Networks," *IEEE Network* 9, no. 6 (2005): 5-11.
16 Karthik Lakshminarayanan, Ion Stoica, and Scott Shenker, "Routing as a Service," working paper, Report No. UCB/CSD-04-1327, Computer Science Division, University of California, Berkeley (2004).

Tussles about the incentive to interconnect transit and access ISPs: With the transit ISPs' business being based on the transport of content across its network, there is an obvious conflict between the desire to locally cache popular content at the access ISP level (not only for cost reduction towards the transit ISP but also to maximize the user's experience in terms of reduced latency). Hence, transit ISPs lack an incentive to participate in an architectural change that is driven by an information-centric viewpoint.[17] However, such conflict could be decisively different when moving towards a transaction-based cost model, which could be enabled by the information-centric nature of the architecture as proposed in the previous section of this article. Our use case in a later section addresses some of these conflicts.

Tussles about interconnecting individual rendezvous solutions within the ICN architecture: Brokering different namespaces (or parts of possibly different namespaces) on the Internet yields to recognized conflicts in terms of access to the namespace as well as the formation of federations for performing the brokerage tasks at hand.[18] The outcome of this tussle inherently influences aspects like reachability in the global information space and eventually the fragmentation of markets due to competing offerings. While interconnection incentives can be driven by economic as well as regulatory forces, desires to isolate counter these forces in areas where such isolation is required (e.g. for security reasons) or desired (e.g. for regionalization reasons).

Other examples of tussles, being left out for reasons of space, address issues of who defines identifiers, the structure for information, and who ensures a trustworthy execution of various functions. While our following taxonomy lists some of these particular conflicts, it is clear that only deeper elaboration and study in future research can shed more light on these important issues.

A First Estimation for a Tussle Taxonomy

Before elaborating on some of our examples in more detail throughout the following section, we first formulate a first estimation for a taxonomy of tussles that can be utilized for a systematic investigation of the larger tussle space.

Figure 2 below presents the various categories of tussles that we have identified. It can be seen that the categories are not mutually exclusive, e.g. security tussles related to information overlap with tussles in the information category with the latter being more concerned with the economic aspects of the information-centric perspective. We can see that the inner categories are all encompassed by the larger socio-economic tussle category that is concerned with the establishment as well as intervention of markets (the intervention driven by various socio-economic players).

---

[17] Jarno Rajahalme, Mikko Särela, Pekka Nikander, and Sasu Tarkoma, "Incentive-Compatible Caching and Peering in Data-Oriented Networks," *Proceedings of the Re-Architecting the Internet Workshop*, ACM New York (2008).
[18] Saikat Guha and Paul Francis, "An End-Middle-End Approach to Connection Establishment," *ACM SIGCOMM Computer Communication Review* 37, no. 4 (2007): 193-204.

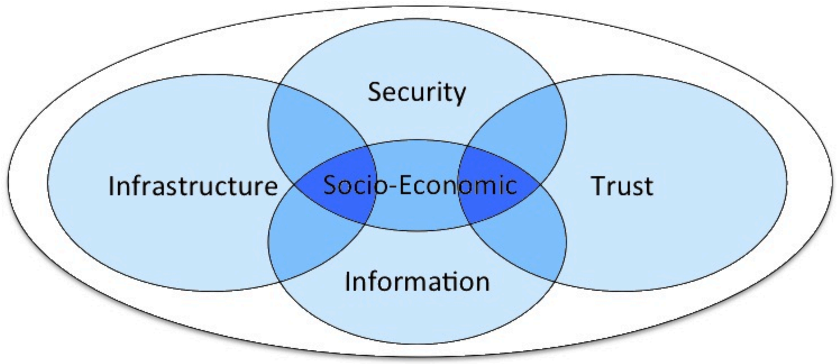

Figure 2: Tussle Categories

Table 1 below elaborates on our tussle taxonomy. It outlines the likely involved actors as well as the architectural functions affected in some form or another. What is missing from this taxonomy is the particular remedy that such an architectural approach provides in accommodating the tussles in each particular category. This is left for our tussle space exploration below, in which we furthermore return to architectural lessons learned from our work.

Table 1: A Tussle Taxonomy

Category	**Aspects of Conflicts**	**Actors Involved**	**Architectural Functions Affected**
Security	*Infrastructure security*: Who makes routing decisions? Who can define requirements affecting infrastructure security (such as path choices, load, etc.)? *Information security*: payload encryption and key management (e.g., self-certified vs. centralized), governance of identifier space (e.g. long-lived vs. short-lived identifiers), governance of information structures (e.g. changes in structure to avoid profiling) *Accountability*: conflict between accountability and privacy of actions as well as content	ISPs, content providers, rendezvous providers, key providers, end users, regulators	Rendezvous, topology management, key management for identifier space, network attachment at end nodes
Trust	*Trust in functions*: policy-compliant execution, isolation of functions possibly misbehaving *Trust in information*: provenance, confidentiality	ISPs, content providers, rendezvous providers, end users, regulators	All architectural functions

Category	Aspects of Conflicts	Actors Involved	Architectural Functions Affected
Information	*Information governance*: governance of identifier space as well as ownership of the defined information space *Brokering information*: policies for matching interests and availability, aspects of profiling usage and consumption for the benefit of, e.g. advertisements	ISPs, content providers, rendezvous providers, end users, regulators	Rendezvous, key management for identifier space
Infrastructure	*Brokering topological capabilities*: exposure of infrastructure information for optimized resource usage within and across networks *Delivering bits*: delivery of individual information items that is compliant to some agreed policy during route selection	Tier1 ISPs, access ISPs, end users, regulators	Topology management, forwarding
Socio-Economic	*Establishing flexible information asymmetries*: flexible exposure of stakeholder requirements, such as QoS or path selection, and association of pricing regimes with each with the ultimate goal of establishing an information asymmetry that results in a market structure *Defining functional boundaries*: Definition of modular boundaries along which to execute functions, including the enforcement of such boundaries through technological, market and regulatory means	All actors in the ecosystem	All architectural functions

USE CASE: ACCESS PROVISIONING

We now turn to use cases for conflicts, investigating how a new information-centric vision of Internet architecture could affect the business models of existing actors. Our first example is that of current ISPs and their core business of providing Internet access to end users and businesses alike.

In today's Internet, there is a business relationship between end users and Internet Service Providers (ISPs). End users usually pay a fixed price for connectivity to ISPs' access networks. Alternative Service Level Agreements (SLAs) have also been proposed, such as volume-based or congestion-based charging.[19] However, fixed pricing for connectivity seems to be the most common pricing scheme that ISPs employ. Often, ISPs offer end users a bundled service of connectivity as well as DNS services. While alternative DNS services are offered (e.g. Google Public DNS,[20] OpenDNS[21]), the majority of end users choose the default DNS service of their own ISP.

[19] Bob Briscoe, "Internet – Fairer is Faster," white paper TR-CXR9-2009-001, British Telecommunications Plc., Sept. 23, 2008, accessed Feb. 20, 2012, bobbriscoe.net/projects/refb/FairerFasterWP.pdf.
[20] Google, "Google Public DNS," accessed Feb. 20, 2012, http://code.google.com/speed/public-dns.
[21] "OpenDNS," accessed Feb. 20, 2012, http://www.opendns.com.

In an information-centric architecture, we can identify two new stakeholders: the Rendezvous Networks (RENE) and their individual rendezvous points, as well as the Topology Management/Internet Topology Formation functions (ITFs). Moreover, there are two different types of end users: publishers[22] and subscribers. The RENE serves as a platform on which end users subscribe for as well as publish an information item. The RENE is a federation of brokers who own the necessary information about the demand and the supply of the information items.[23]

This market structure seems very similar to a two-sided market. An economic network is a two-sided market if (a) there are two distinct groups of customers; (b) the value obtained by one kind of customer increases with the number of the other kind of customer; and (c) an intermediary is necessary for internalizing the externalities created by one group for the other group.[24]

In two-sided markets, there are traditionally three types of stakeholders: two end user groups and an intermediate platform. In the current Internet's market (see Figure 3 below), we identify the individual end users and content providers (the two end user groups), as well as the ISP (the platform).

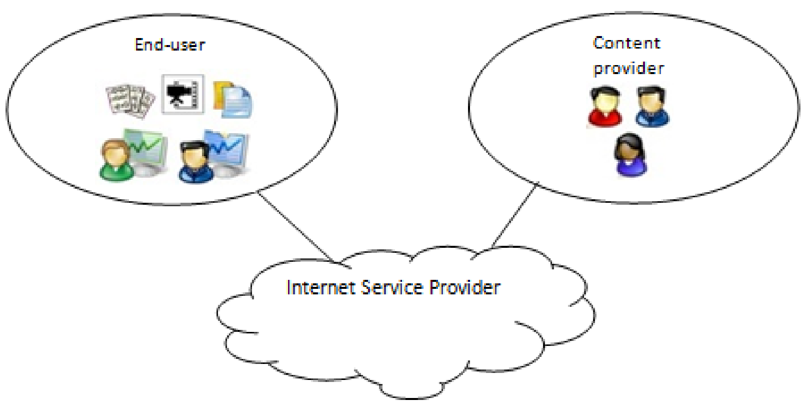

Figure 3: A Two-Sided Market in the Current Internet

It should be noted here that the role of network externalities is also of high importance. In particular, we can distinguish between two main sets of externalities in a two-sided market: the usage externality results from the interaction between two different user groups, whereas the membership externality refers to the installed base.[25]

The presence of externalities and the existence of two different prices raise the issue of price allocation. Since there are two different user groups, ISPs face two distinct types of demand.

---

[22] Caches can be seen as alternative publishers for the same content.
[23] We observe that today's DNS is similar to the rendezvous functionality within our architecture.
[24] Jean-Charles Rochet and Jean Tirole, "Platform Competition in Two-Sided Markets," *Journal of the European Economic Association* 1, no. 4 (2003): 990–1029.
[25] Jean-Charles Rochet and Jean Tirole, "Two-Sided Markets: A Progress Report," *RAND Journal of Economics* 37, no. 3 (2006): 645-667.

Therefore, the price structure will reflect the demand elasticity and externalities (in order to get both sides on board), as well as platform competition. Thus, the final end price is composed of a price paid by the web site and a price paid by Internet users.

However, we observe that the above phenomenon within today's Internet is quite different from an information-centric approach. In the previous example, the platform (namely the ISP) owns the network and charges based on the operating cost. The main difference lies in the new stakeholders that have appeared: the topology manager who owns the network, as well as the rendezvous network that could be seen as a broker between end users and topology managers. Thus, we can see the above case as a three-sided market (see Figure 4 below).

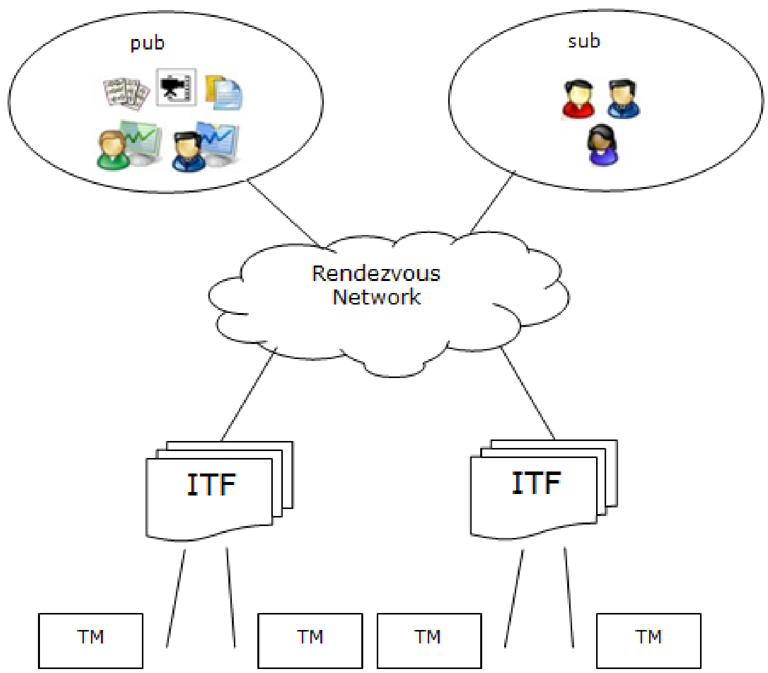

Figure 4: A Three-Sided Market in ICN

This might result in evolution of the current business models. For example, the rendezvous network could offer a fixed-fee SLA to publishers and subscribers for the connectivity service. Additionally, end users could pay extra usage-based fees and these fees (or a proportion of them) could "pass through" the ITFs via the RENEs.

We argue that the information-centric approach might possibly enable new market mechanisms. The current Internet does not provide sufficient economic mechanisms for stakeholders to express preferences. This is largely due to the fundamental separation within the IP architecture between the opaque bit transfer and the information exchange at upper layers. It is this separation that makes the establishment of pricing regimes an expensive solution, requiring out-of-band signaling solutions, which are largely limited to few services only. Instead, it has been proposed in the literature[26] to flexibly establish information asymmetries through utilizing implicit knowledge about the transferred information structures (without a need for expensive and error-prone deep packet inspection solutions). This could lead to an accountability framework for resource utilization that spans applications and networks.

In such a framework, subscribers could express their preference about the QoS (quality of service) for receiving a specific information item. Potential preferences could be giving higher or lower priority for the subscribed information item, receiving a file from a specific publisher, paying the lower price for an information item, etc. Respectively, publishers could also express their preferences about the QoS of their delivered information item, such as publishing the same content to different rendezvous points and have different SLAs with them – classes of services, etc.

With this approach, the various functions in the network could have all the necessary information about the demand and supply of the available information items without the need for an explicit signaling framework as in today's IP networks. Such demand/supply information could play a role in the establishment of final delivery graphs throughout the network. For instance, the final matching decision within the rendezvous point could be based on the feedback that inter-domain topology formation functions will provide to the rendezvous point about potential paths, their utilization, their metric of resilience, etc. Auction mechanisms could also be applied, in which bids will determine which ITF will be chosen by the rendezvous point for a specific data transmission. This could lead to different types of competition games between providers for inter-domain connectivity.

USE CASE: CONTENT DELIVERY

In this section, we take a closer look at the most common business models in Content Distribution Network (CDN) markets. In particular, we provide a brief analysis of the status in the current Internet. Furthermore, we investigate how such a business model could be affected or changed in an information-centric architecture.

Generally, CDNs perform content replication in order to efficiently distribute it close to end-users. To achieve this, CDNs globally locate surrogate servers with cached content to improve accessibility and lessen the load of the origin servers.[27] Additionally, CDN solutions result in reducing congestion within a network as well as the need for capacity expansion investments. From a technical standpoint, CDN overlays usually perform operations such as DNS dispatching, URL dynamic re-

[26] Trossen and Gergely; Jochen Wulf, Rüdiger Zarnekow, Thorsten Hau, and Walter Brenner, "Carrier Activities in the CDN Market – An Exploratory Analysis and Strategic Implications," paper presented at the 14th International Conference on Intelligence in Next Generation Networks, Berlin, Germany, Oct. 2010, accessed Feb. 20, 2012, http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=5640943.

[27] George Pallis and Athena Vakali, "Insight and Perspectives for Content Delivery Networks," *Communications of the ACM* 49, no. 1 (2006): 101-106.

writing, and HTTP redirection. Such request routing mechanisms are used to forward the requests to cached servers that are close to the end users.[28]

Let us investigate the most common classes of CDNs. The first category is the content-centric model. Here, the CDNs' operation is driven by the needs of the content owners. In particular, content providers pay CDNs to host and serve their content under the promise of lower latencies and generally increased QoS (see Figure 5 below).

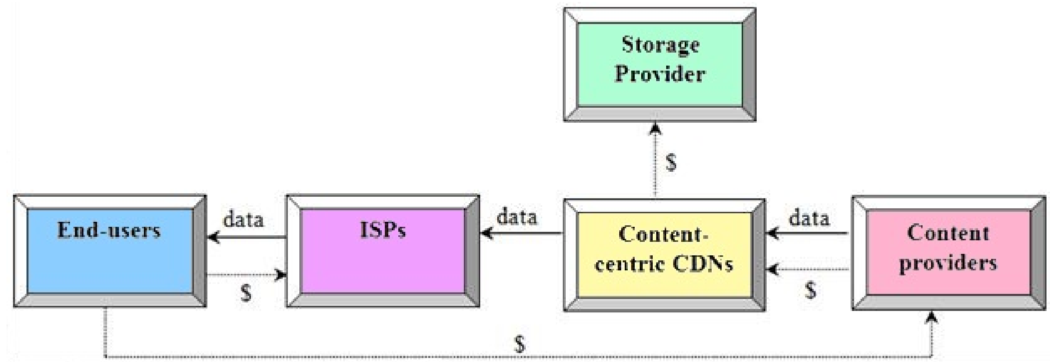

Figure 5: Money Flow in Content-Centric CDNs

Another category of CDNs is the access-centric model. This type of CDN is driven by the needs of Internet access service providers. Here, the money flow is different from the content-centric model, since an ISP pays the CDN to serve popular content from caches close to its customers (see Figure 6 below).

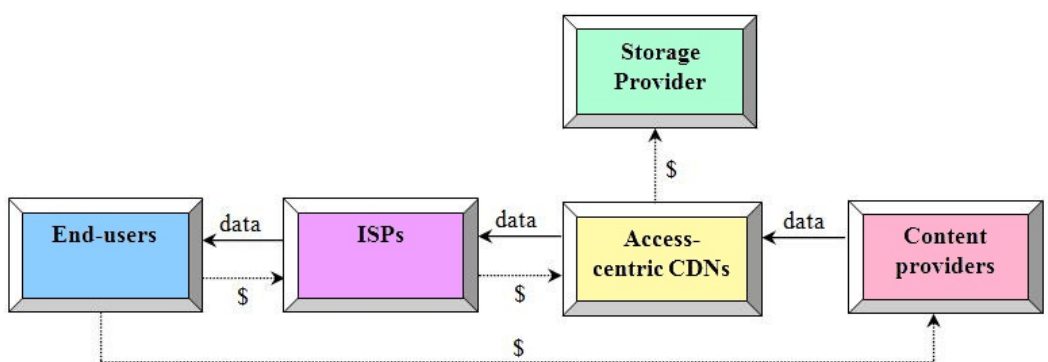

Figure 6: Money Flow in Access-Centric CDNs

[28] Clark, Lehr, Bauer, Faratin, Sami, and Wroclawski.

Access centric CDNs can be distinguished into two separate subgroups. Internal Service Provider CDNs aim at reducing the internal bandwidth load within an ISP's network by performing caching. These are mainly used to enhance service differentiation and premium content distribution (e.g. for Internet Protocol TV, Voice over Internet Protocol, etc.). On the other hand, External Service Provider CDNs perform caching in order to reduce the external bandwidth cost for ISPs.

An obvious problem in both models is the usual network-agnostic nature of CDNs since they do not own any network infrastructure. They have limited control over the network in which they operate, and the only available information they have is the IP address of a specific request. Hence, based on the proximity of the source address, CDNs redirect the request to the "best" server. Such criteria could be different in each CDN, e.g. latency, locality, etc.

However, utilizing network proximity is not always the right approach. First, this could result in a higher DNS request rate. Moreover, DNS requests by a CDN do not always provide the correct information about a user's network location,[29] e.g. in the case of a cache miss. Consequently, new trends appear in CDN markets.[30] One such trend is the establishment of CDNs directly through network owners, offering bundled services that take advantage of combining access provisioning and content hosting. Such an advantage arises from having the necessary knowledge about traffic bottlenecks and not being dependent upon bandwidth from third-party carriers. A similar approach for peer-assisted content distribution is the insertion of entities equipped with high resources (in terms of bandwidth and storage) that are controlled and managed by the ISP; these are called ISP-owned Peers (IoPs).[31]

Another trend in CDN markets is the cooperation of different stakeholders, such as hosting providers, access service providers, and backbone carriers (e.g. Inktomi's Content Bridge Alliance).[32] The role of access network providers is of crucial importance in such federation solutions, since the access network is a critical quality bottleneck for end users. On the other hand, the Content Delivery Network Interconnection (CDNI) efforts[33] propose to interconnect separate CDNs, supporting end-to-end content delivery while dynamically expanding the delivery footprint. However, this approach has similar locality disadvantages to current CDNs.[34]

Let us now investigate how these business models could be affected or changed in an information-centric architecture. First, the notion of content delivery is central to an information-centric

[29] Paul Vixie, "What DNS is Not," *Communications of the ACM* 52, no. 12 (2009): 43-47.
[30] Wulf, Zarnekow, Hau, and Brenner.
[31] Ioanna Papafili, Sergios Soursos, and George D. Stamoulis, "Improvement of BitTorrent Performance and Inter-Domain Traffic by Inserting ISP-Owned Peers," paper presented at the 6th International Workshop on Internet Charging and QoS Technologies, Athens, Greece, 2009.
[32] Rebecca Wetzel, "CDN Business Models – Not All Cast from the Same Mold," working paper, Wetzel Consulting LLC, accessed Feb. 21, 2012, http://wetzelconsultingllc.com/CDNArticle.pdf.
[33] Gilles Bertrand, Francois Le Faucheur, and Larry Peterson. Content Delivery Network Interconnection (CDNI) Experiments, Version 00, Feb. 16, 2011, accessed Feb. 21, 2012, http://tools.ietf.org/html/draft-bertrand-cdni-experiments-00.
[34] Locality issues arise when HTTP re-directs end up in parts of the CDN that are more associated with the business association of the end user (e.g. his home ISP) than his actual physical location. This could occur, for instance, when travelling or utilizing corporate VPN functions.

architecture so that the role of a CDN can be assumed by anybody in the network who might have the requested information available. Hence, caching becomes a natural part of the network's operation. This can be supported in ISP-driven models in which the topology information within the ISP's topology manager is utilized to determine "nearby" caches as potential publishers for information requested by subscribers. This equates to the access-centric model in today's Internet (albeit with the possibility of utilizing more of the ubiquitous resources that are available today on standard, user-managed equipment). The role of today's CDN providers would be differentiated towards a mere (storage) resource provider; a new role that could easily be assumed by a massive storage (cloud) provider.

But there is also the equivalent to the aforementioned content-centric model, in which the content provider aims at increasing the end user experience. This approach is enabled by the rendezvous function taking the main role in decision-making over publisher selection, albeit in collaboration with the transport network providers. It is clear that the selection considerations here are mainly driven by QoE (quality of experience) criteria, while the access-centric model is driven by QoS criteria. But similarly to the access-centric model, the role of traditional CDN providers would be that of mere storage providers. Both cases are represented in Figure 7 below. The only difference between the access-centric and content-centric models is that in the latter there is an extra money flow from content providers to ISPs.

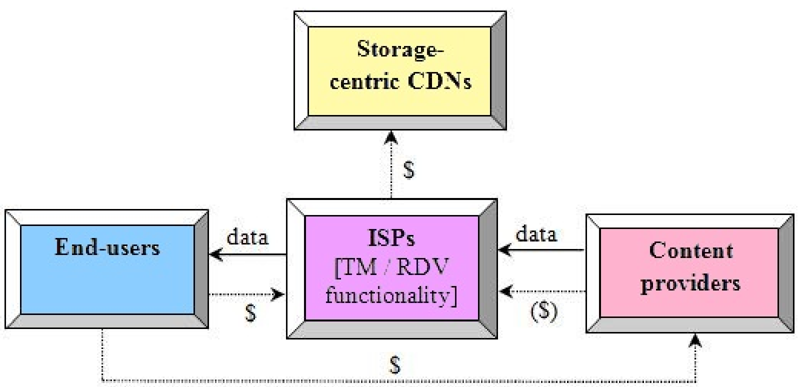

Figure 7: Money Flow in Storage-Centric CDNs

From this analysis, we can derive a very important potential impact on today's CDN market. The non-storage-related functionality of traditional CDNs will be inherently implemented by the functions of topology management and rendezvous in ICN architecture,[35] while the remaining storage provisioning can be realized either by managed storage providers (such as Amazon or other

[35] In other words, the complexity of CDNs in implementing the plethora of re-directions necessary to end up in the appropriate cache is largely implemented by core functions of the network architecture – a fundamental difference from today's IP world.

cloud storage providers) or through user-managed resources. We can conclude that ICN might evolve current business models by differentiating current players' roles, and by giving opportunities to new players entering the Internet market.

Scenario 1: Minimizing Congestion Cost

In our first scenario, we assume that there are two ISPs that have a peering Service Level Agreement (SLA). Both ISPs are connected to a higher-tier ISP. Moreover, we assume that both ISP's topology managers (TMs) interact.

In the current Internet, ISPs have no information about what "kind" of traffic traverses their network.[36] Techniques like DPI (deep packet inspection) are utilized to gain insight about the characteristics of their traffic (e.g. peer-to-peer or real-time traffic). However, DPI boxes cannot capture the entirety of the traffic due to, for instance, encrypted packets or tunneling over well-known ports (such as port 80).

On the other hand, in an information-centric architecture such information is available within the network and this could be a useful input for the topology managers. Let us illustrate this issue. In Figure 8 below, we can see a publisher (P_1) in ISP_1, who has available a very popular information item. In ISP_2, there are three subscribers (S_1, S_2, and S_3 respectively) that have subscribed to this information item.

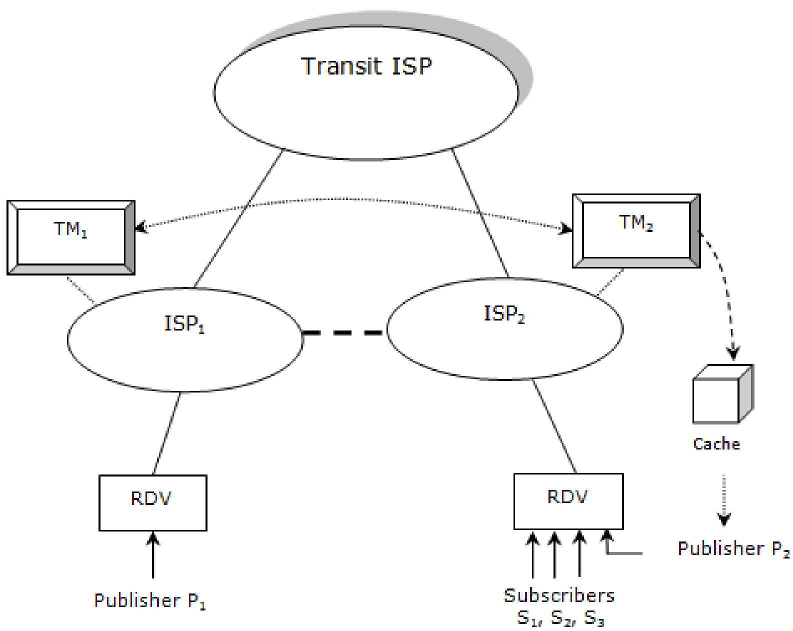

Figure 8: Using ICN to Minimize Congestion Cost

36 David Hausheer, Pekka Nikander, Vincenzo Fogliati, Klaus Wünstel, María Ángeles Callejo, Santiago Ristol Jorba, Spiros Spirou, Latif Ladid, Wolfgang Kleinwächter, Burkhard Stiller, Malte Behrmann, Mike Boniface, Costas Courcoubetis, and Man-Sze Li, "Future Internet Socio-Economics – Challenges and Perspectives," in *Towards the Future Internet – A European Research Perspective*, ed. Georgios Tselentis, John Domingue, Alex Galis, Anastasius Gavras, David Hausheer, Srdjan Krco, Volkmar Lotz, and Theodore Zahariadis (Amsterdam: IOS Press, 2009), 1-11.

Since both ISPs have the available information about the publishers and subscribers for a specific information item, the topology managers can interact for the purpose of caching and hence reduce the exchanged traffic between their networks. In particular, we observe that there are many subscribers in ISP_2 for an information item located in ISP_1. Since the ISPs have a peering relationship (so ISP_2 will not be charged by ISP_1 for the received traffic), TM_1 could inform TM_2 to cache the specific information item. Thus the cache server of ISP_2 could now act as a new publisher (P_2) in ISP_2. As a result, the new subscribers for this information item in ISP_2 will receive it with increased QoS. Moreover, both ISPs will reduce their traffic within their network (no traffic for ISP_1 and less traffic for ISP_2). Hence, both ISPs have an incentive to cooperate.

Scenario 2: Minimizing Transit Cost

In this scenario, we again assume two ISPs with a peering relationship. There is also another ISP (ISP_3), which has no peering agreement with ISP_1 and ISP_2. All three are connected to a higher-tier ISP. In Figure 9 below, we assume publisher P_1 has a very popular information item. In ISP_1 and ISP_2 there are a number of subscribers that are interested in this publication. However, neither ISP_1 nor ISP_2 has a peering relationship with ISP_3. Consequently, the traffic for the specific information item is passed through the transit ISP, resulting in increased transit costs for both ISPs.

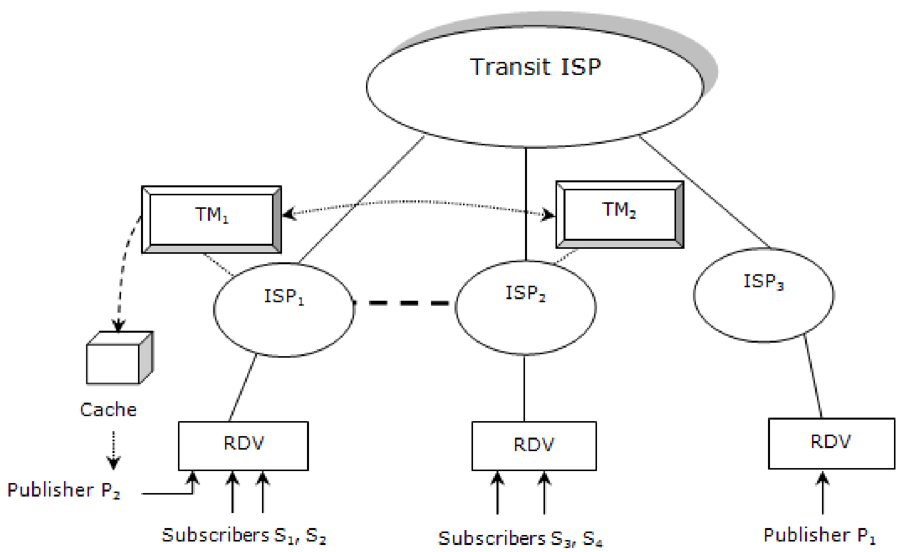

Figure 9: Minimizing an ISP's Transit Cost

In the information-centric architecture case, ISP_1 and ISP_2 can implement their incentive to cooperate, in order to avoid additional charging by the transit ISP, by having either ISP_1 or ISP_2 cache this information item in its server. This cache could act again as a new publisher (P_2) for the same information item. In this case, both ISPs will avoid additional charges by the transit ISPs. Moreover, ISP_3 will lower potential congestion within its network.

LESSONS LEARNED

In this article we have presented potential use cases in ICN architecture – with a focus on economics, business, and policy aspects. In this section we highlight several key points that should be taken from the aforementioned analysis.

Design for Choice

Various incentives on the regulatory and user side, as well as on the ISP side, emphasize the need for choice in revealing information and assembling functions acting on this information. This is driven by the stronger emphasis in the ICN architecture on the exchange of information, which inherently carries value for actors in the ecosystem. This is somewhat different than interconnection in the Internet today, which focuses on resource pooling and therefore cost minimization. Differentiation on the service or content level is hardly provided. Hence, any solution should consider mechanisms for choice. The ICN approach provides such incentives for choice in several functions, including the ITF function, the topology manager, and the rendezvous network. For example, in the use access provisioning case analyzed above, end users may express their preferences about QoS via the RENE provider. Such expression of choice can be implemented by accompanying the information structures (on which the network directly operates) with metadata information that defines the preferences and policies being taken into account for the particular structure to be delivered.

Design for Isolation

One expression of choice is a desired isolation of information spaces, each of which might be interconnected by its own provider or to each of which end users might connect through different ISPs (e.g. using specific financial network providers compared to regular ISPs). But also the enforcement of digital rights influences the incentive to widely interconnect within an isolated island of policy enforcement. Such regional power struggles already exist today and are likely to exist in the future. Any design must accommodate these influences.

In an ICN architecture, such aspects of isolation can be accommodated by the clear identification of the various roles that are responsible for functions of finding information, building a delivery graph, and eventually delivering the bits (of information). In particular, the second function today only exists in the many technological extensions to the Internet, realized in various middlebox solutions, all of which are only badly (if at all) exposed to the various actors in the system.[37]

In both scenarios presented in the content delivery use case above (concerning the reduction of congestion and transit costs) the ISPs already have the necessary information to perform caching, without the need for middlebox solutions (i.e. using deep packet inspection boxes to recognize the "kind" of traffic within their network).

[37] We recognize that various efforts, e.g. in the IETF, exist to expose such middleboxes. These efforts are driven by similar motivations to ours when it comes to designing for choice.

Design for Flexible Deployment

The need for evolvability of solutions has long been recognized.[38] Hence, any design should consider various deployment scenarios. One such scenario is being driven by vertical industries for initial adoption, significantly contrasting a full adoption model in which every player will need to adopt a new technology (like ICN). Consideration of various deployment options should be accompanied by a proper understanding of their market impact.

Decoupling Business Models

Another aspect is that of decoupling business models, such as interconnection models on the bit and information level. Such decoupling would occur by routing discovery requests along the same upgraph[39] connections that are being established through bit-level interconnection.[40] With that, one creates a strong alignment of the business models underlying both interconnections, namely that at the bit transport level and that at discovery level. However, such alignment is not necessarily upheld in reality, such as in the search space today.

Taking into consideration the aforementioned points, there are two major policy findings that stand out:

Clearly Defined Modular Boundaries Are Crucial. Recent research focuses on the role that modular boundaries play when needing to accommodate tussles at system runtime and minimize the impact of a function on other functions, reducing the overall dependence of players on other stakeholders.[41]

But while these works focus on the general importance of this architectural design principle, our work within an information-centric architecture context provides examples for such modular boundaries, namely those of the three crucial functions of finding information, building appropriate delivery graphs, and delivering bits of information. These main functions are well exposed and defined through the architecture. This aspect is important, for instance, in separating the information brokering from the transport of the final bits, in applying different types of policies, etc.

---

[38] David D. Clark, Karen Sollins, John Wroclawski, and Ted Faber, "Addressing Reality: An Architectural Response to Real-World Demands on the Evolving Internet," *ACM SIGCOMM Computer Communication Review* 33, no. 4 (2003): 247-257; Mark Handley, "Why the Internet Only Just Works," *BT Technology Journal* 24, no. 3 (2006): 119-129; Sylvia Ratnasamy, Scott Shenker, and Steven McCanne, "Towards an Evolvable Internet Architecture," *ACM SIGCOMM Computer Communication Review* 35, no. 4 (2005): 313-324.

[39] Within the so-called "valley-free routing" of the Internet, an "upgraph" connection is one to a transit provider. See Sophie Y. Qiu, Patrick D. McDaniel, and Fabian Monrose, "Toward Valley-Free Inter-Domain Routing," paper presented at IEEE International Conference on Communications (ICC '07), Glasgow, Scotland, June 2007, accessed Feb. 27, 2012, http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=4289005.

[40] Jarno Rajahalme, Mikko Särelä, Kari Visala, and Janne Riihijarvi, "Inter-Domain Rendezvous Scalability in Content Networking Architectures," Technical Report #TR09-003, Publish-Subscribe Internet Routing Paradigm, accessed Feb. 21, 2012, psirp.org/files/Deliverables/plugin-TR09-003_Rendezvous.pdf.

[41] Costas Kalogiros, Alexandros Kostopoulos, and Alan Ford, "On Designing for Tussle: Future Internet in Retrospect," *Proceedings of the 15th Open European Summer School and IFIP TC6.6 Workshop on The Internet of the Future* (2009): 98–107.

This separation effectively creates a tussle space, and therefore a market boundary, between bit- and information-oriented markets.

However, ancillary functions such as identifier governance, key management, and network attachment also have their clear role in the architecture. Hence, we assert that the ICN context provides a positive example for flexible "good" design along well-defined modular boundaries. Significant future work, however, is required to shed more light on the aspects of what defines "good" and "flexible" in terms of metrics – metrics that could potentially be generalized for other architectural use cases.

Flexibly Creating Information Asymmetries Is Key. Any formation of markets is based on creating as well as re-shaping appropriate information asymmetries between the various actors within the markets.[42] The access provisioning use case investigated above illustrates how such information asymmetries may affect actors' behavior in a three-sided market. Any architecture needs to accommodate such fundamental mechanisms in order to enable (and sustain) a flourishing ecosystem, i.e. to enable not just today's business models but also any future business models that the actors within the system strive to establish.

Within a communication system, this economic observation boils down to enabling a well-defined information exposure between various parties in the system (i.e. for congestion notification in today's Internet).[43] Generally, such exposure of information (such as policies, preferences, end user interests, and topological information being used for routing or observed congestion in the network) needs to be inherently supported by a communication system. We assert that a communication system that itself operates on information (within well-defined and exposed structures for this information) provides an improved ability for supporting such information exposure. We recognize, however, that such an assertion needs a larger pool of anecdotal evidence through an extended investigation of use cases; an effort being left for future work in the field.

CONCLUSIONS

Information-centric networking has attracted increasing attention in the network architecture community, both on the academic stage and within corporate research organizations. Given the attention that this particular range of architectural proposals for the Internet' s evolution has been receiving, it seems only natural to study the socio-economic playground and the tussles that such architectural context would bring about. This is not only important for understanding the socio-economic impact of such possible technological change. It is also a crucial exercise in understanding the viability of technological solution proposals within the wider socio-economic environment that is our society.

[42] George Hoffer and Michael Pratt, "Used Vehicles, Lemons Markets, and Used Car Rules: Some Empirical Evidence," *Journal of Consumer Policy* 10, no. 4 (1987): 409-414.
[43] Alan Ford, Philip Eardley, and Barbara van Schewick, "New Design Principles for the Internet," working paper, The Center for Internet and Society, accessed Feb. 21, 2012, cyberlaw.stanford.edu/files/publication/files/05207995.pdf.

We clearly recognize that this article can only be the starting point for exploring the tussle space for a global-scale communication system such as the Internet. But we believe that the initial tussle taxonomy in this article as well as our dedicated use cases provide a useful first insight. Our future work will focus on developing a methodology for exploring and better understanding the tussle spaces of evolving Internet architectures.

BIBLIOGRAPHY

Bertrand, Gilles, Francois Le Faucheur, and Larry Peterson. Content Delivery Network Interconnection (CDNI) Experiments, Version 00, Feb. 16, 2011, accessed Feb. 21, 2012, http://tools.ietf.org/html/draft-bertrand-cdni-experiments-00.

Briscoe, Bob. "Internet – Fairer is Faster." White paper TR-CXR9-2009-001, British Telecommunications Plc., Sept. 23, 2008. Accessed Feb. 20, 2012, bobbriscoe.net/projects/refb/FairerFasterWP.pdf.

Brown, Ian, David Clark, and Dirk Trossen. "Should Specific Values Be Embedded in the Internet Architecture?" Proceedings of the Re-Architecting the Internet workshop (ReARCH '10), Nov. 2010. Accessed Feb. 27, 2012, http://dl.acm.org/ft_gateway.cfm?id=1921246&type=pdf&CFID=68093922&CFTOKEN=42506244.

Caesar, Matthew and Jennifer Rexford. "BGP Routing Policies in ISP Networks." *IEEE Network* 9, no. 6 (2005): 5-11.

Clark, Dave, Bill Lehr, Steve Bauer, Peyman Faratin, Rahul Sami, and John Wroclawski. "Overlay Networks and the Future of the Internet." *Communications & Strategies* 63, 3rd q. (2006): 1-21.

Clark, David D. Karen Sollins, John Wroclawski, and Ted Faber. "Addressing Reality: An Architectural Response to Real-World Demands on the Evolving Internet." *ACM SIGCOMM Computer Communication Review* 33, no. 4 (2003): 247-257.

Clark, David D., John Wroclawski, Karen R. Sollins, and Robert Braden. "Tussle in Cyberspace: Defining Tomorrow's Internet." *IEEE/ACM Transactions on Networking* 13, no. 3 (2005): 347-356.

Ford, Alan, Philip Eardley, and Barbara van Schewick. "New Design Principles for the Internet." Working paper, The Center for Internet and Society. Accessed Feb. 21, 2012, cyberlaw.stanford.edu/files/publication/files/05207995.pdf.

Google. "Google Public DNS." Accessed Feb. 20, 2012, http://code.google.com/speed/public-dns.

Guha, Saikat and Paul Francis. "An End-Middle-End Approach to Connection Establishment." *ACM SIGCOMM Computer Communication Review* 37, no. 4 (2007): 193-204.

Gyöngyi, Zoltán and Hector Garcia-Molina. "Web Spam Taxonomy." Paper presented at the first International Workshop on Adversarial Information Retrieval on the Web, Chiba, Japan, May 2005. Accessed Feb. 27, 2012, airweb.cse.lehigh.edu/2005/gyongyi.pdf.

Handley, Mark. "Why the Internet Only Just Works." *BT Technology Journal* 24, no. 3 (2006): 119-129.

Hoffer, George and Michael Pratt. "Used Vehicles, Lemons Markets, and Used Car Rules: Some Empirical Evidence." *Journal of Consumer Policy* 10, no. 4 (1987): 409-414.

Hausheer, David, Pekka Nikander, Vincenzo Fogliati, Klaus Wünstel, María Ángeles Callejo, Santiago Ristol Jorba, Spiros Spirou, Latif Ladid, Wolfgang Kleinwächter, Burkhard Stiller, Malte Behrmann, Mike Boniface, Costas Courcoubetis, and Man-Sze Li. "Future Internet Socio-Economics – Challenges and Perspectives." In *Towards the Future Internet – A European Research Perspective*, edited by Georgios Tselentis, John Domingue, Alex Galis, Anastasius Gavras, David Hausheer, Srdjan Krco, Volkmar Lotz, and Theodore Zahariadis, 1-11. Amsterdam: IOS Press, 2009.

Internet Engineering Task Force, Content Delivery Networks Interconnection (CDNI.) Accessed Feb. 27, 2012, http://datatracker.ietf.org/wg/cdni/.

Kalogiros, Costas, Alexandros Kostopoulos, and Alan Ford. "On Designing for Tussle: Future Internet in Retrospect." *Proceedings of the 15th Open European Summer School and IFIP TC6.6 Workshop on The Internet of the Future* (2009): 98–107.

Lakshminarayanan, Karthik, Ion Stoica, and Scott Shenker. "Routing as a Service." Working paper, Report No. UCB/CSD-04-1327, Computer Science Division, University of California, Berkeley (2004).
"Named Data Networking." Dec. 31, 2011, accessed Feb. 16, 2011, http://www.named-data.net/.
"OpenDNS." Accessed Feb. 20, 2012, http://www.opendns.com.
Pallis, George and Athena Vakali. "Insight and Perspectives for Content Delivery Networks." *Communications of the ACM* 49, no. 1 (2006): 101-106.
Papafili, Ioanna, Sergios Soursos, and George D. Stamoulis. "Improvement of BitTorrent Performance and Inter-Domain Traffic by Inserting ISP-Owned Peers." Paper presented at the 6th International Workshop on Internet Charging and QoS Technologies, Athens, Greece, 2009.
"PURSUIT." Accessed Feb. 16, 2011, http://www.fp7-pursuit.eu/PursuitWeb/.
Qiu, Sophie Y., Patrick D. McDaniel, and Fabian Monrose. "Toward Valley-Free Inter-Domain Routing." Paper presented at IEEE International Conference on Communications (ICC '07), Glasgow, Scotland, June 2007. Accessed Feb. 27, 2012, http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=4289005.
Rajahalme, Jarno, Mikko Särela, Pekka Nikander, and Sasu Tarkoma. "Incentive-Compatible Caching and Peering in Data-Oriented Networks." Proceedings of the Re-Architecting the Internet Workshop, ACM New York (2008).
Rajahalme, Jarno, Mikko Särelä, Kari Visala, and Janne Riihijarvi. "Inter-Domain Rendezvous Scalability in Content Networking Architectures." Technical Report #TR09-003, Publish-Subscribe Internet Routing Paradigm. Accessed Feb. 21, 2012, psirp.org/files/Deliverables/plugin-TR09-003_Rendezvous.pdf.
Ratnasamy, Sylvia, Scott Shenker, and Steven McCanne. "Towards an Evolvable Internet Architecture." *ACM SIGCOMM Computer Communication Review* 35, no. 4 (2005): 313-324.
Rochet, Jean-Charles and Jean Tirole. "Platform Competition in Two-Sided Markets." *Journal of the European Economic Association* 1, no. 4 (2003): 990–1029.
———. "Two-Sided Markets: A Progress Report." *RAND Journal of Economics* 37, no. 3 (2006): 645-667.
Trossen, Dirk and Gergely Biczok. "Not Paying the Truck Driver: Differentiated Pricing for the Future Internet." Proceedings of the Re-Architecting the Internet workshop (ReARCH '10), Nov. 2010. Accessed Feb. 27, 2012, http://dl.acm.org/ft_gateway.cfm?id=1921235&type=pdf&CFID=68093922&CFTOKEN=42506244.
Trossen, Dirk, Mikko Sarela, and Karen Sollins. "Arguments for an Information-Centric Internetworking Architecture." *ACM SIGCOMM Computer Communication Review* 40, no. 2 (2010): 27-33.
Vixie, Paul. "What DNS is Not." *Communications of the ACM* 52, no. 12 (2009): 43-47.
Wetzel, Rebecca. "CDN Business Models – Not All Cast from the Same Mold." Working paper, Wetzel Consulting LLC. Accessed Feb. 21, 2012, http://wetzelconsultingllc.com/CDNArticle.pdf.
Wulf, Jochen, Rüdiger Zarnekow, Thorsten Hau, and Walter Brenner. "Carrier Activities in the CDN Market – An Exploratory Analysis and Strategic Implications." Paper presented at the 14th International Conference on Intelligence in Next Generation Networks, Berlin, Germany, Oct. 2010. Accessed Feb. 20, 2012, http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=5640943.
"YouView." Accessed Feb. 16, 2012, http://www.youview.com/.