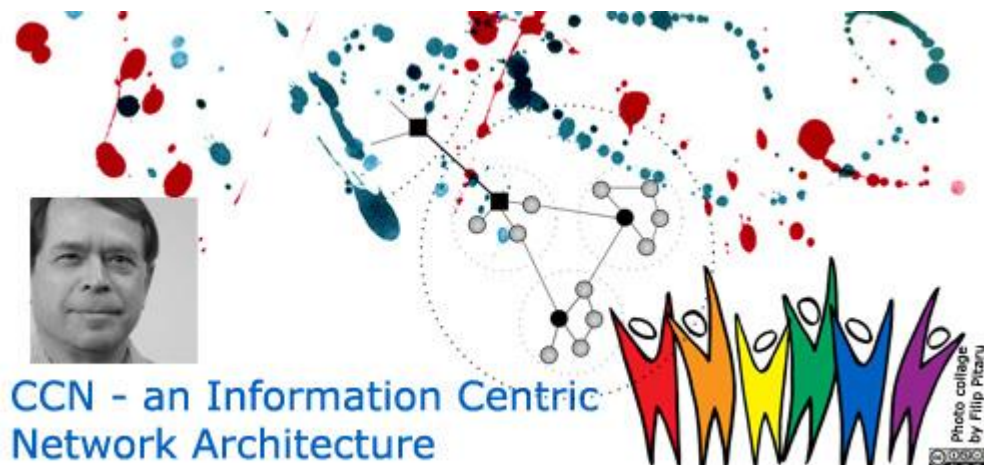


Content Centric Networks: ethical and legal implications, privacy and security concerns of new experimental technologies

Undergraduate project by Filip Pitaru



CCN - an Information Centric Network Architecture

Abstract

A Content Centric Network represents a new alternative to the existing architecture that supports the Internet. As the name itself suggests, the proposed solution focuses on content/information rather than referencing a physical location where the data would reside, as the current conversational model does.

This paper presents the concepts behind CCN and analyses some of the security and privacy concerns that need to be addressed before implementing it on a large scale.

Table of Contents

Abstract.....	2
1. Introduction	4
2. The Content Centric Networks Architectural Model	7
2.1 The theory behind CCN	7
2.2 CCN example	9
2.3 Hardware requirements and measured efficiency	11
3. Security in CCN.....	13
3.1 The credential system	13
3.2 Security concerns	17
3.2.1 Caching.....	17
3.2.1.1 The unencrypted privacy trail	18
3.2.1.2 Legal implications in the EU	22
3.2.2 In content we trust.....	23
3.2.2.1 The WikiLeaks example.....	25
3.2.3 The geolocation of anonymity	28
3.2.3.1 Understanding geolocation.....	28
3.2.3.4 Potential solutions to the privacy vs. reachability paradigm.....	34
4. Implementation	35
4.1 Advantages of CCN in a wireless mesh network environment.....	37
4.2 Device interoperability.....	39
5. Conclusion.....	42
References.....	45

1. Introduction

In his study "Beyond Moore's law: Internet growth trends", Roberts (2000) states that the trends concerning the expansion of the Internet are very similar to those discovered by Gordon Moore regarding the semiconductor industry and the exponential increase over time of processing power.

However, design issues have been noticed almost immediately after the boom of the World Wide Web and authors such as Scott Shanker in his "Design issues for the future Internet" (1995) advocated that the use of real time applications like video and voice will lead towards a very congested network, as the original design did not take into account for such heavy bandwidth usage.

In contrast with recent VoIP, teleconferencing and web video advancements, recent studies such as the one undertaken by the European deep packet inspection services provider Ipoque (2008, online) for the Portuguese ISP TVTEL and the Austrian University of Innsbruck, have shown that peer to peer applications (P2P) can consume up to 64% of the network traffic. At the time of these studies, the amount of traffic was done by only 10% of the users, their number slowly increasing.

While telecommunications companies have invested heavily in their infrastructure (e.g. most submarine communication cables have been installed since 1998 – source Wikipedia) and modern domestic routers have reached an average processing speed of 350 MHz¹, the

¹ Netgear, Buffalo Technologies and Linksys websites have been consulted.

TCP/IP protocols that are behind this gigantesque infrastructure still focus on location in a more and more content-centric based Internet.

According to statistics offered by Kantar Group Company (2011, online), it seems that we deal in digital information more than we would have ever

expected, in an increasingly networked world, where potentially everybody can reach anyone, absolutely anywhere.

For the UK alone, the Internet penetration rate has reached 82.5%, with people spending an average of 22 hours per week online, 29% of them using their time for entertainment. Please see Figure 1.




Figure 1: United Kingdom Internet usage statistics

Not only the information exchange flow and storage has increased tenfold for companies or governments and we are already dealing with multiple functional virtual economies and currencies but private individuals also seem to be making the transition toward the virtual space.

That would translate into quite a fraction used from our ontological space in dealing with information, which is now stored and transferred almost solely in digital format.

The ordinary person satisfies its daily basic communication needs through email, VoIP or instant messengers, creates and stores emotionally bound experiences in digitized formats (photos, videos, recordings, writings, artistic expressions) and also deals in and with shared digital information regarding its work or education related

experiences.

While an alternative solution to the exhaustion of IPv4 addressing space has been the IPv6 addressing scheme and for most online entertainment providers, Content Delivery Networks (CDN) have proved to be an efficient way of increasing the downtime in terms of user access to their services by replicating content closer to the end user, these approaches still do not complement organic behaviour, which CCN seems to closely mimic.

The “natural approach” in this case is the fact that humans cannot remember or associate the long numerical address which stand behind the IP protocol (even more so with 64 bit address) and they prefer using names, to which they can attach meaning and are also much easier to remember (e.g. www.google.com vs. 209.85.149.147).

However, the evolutive pressure behind changing the current communication paradigm does not rely on the fact that a conceptual metamorphosis took place and the actual physical location of information has become irrelevant to the end user in today’s digital era, as people are interested only in retrieving content immaterial to where it resides. This would be insufficient to support such a large scale infrastructural change if it would not also present economic advantages.

Initial findings by Uichin Lee et al. (April 2011) have shown that CCN provides a much higher efficiency in terms of energy use compared to the current architectural system and Lauinger (Sept. 2010) has shown that security seems more robust in most cases than the current model.

It is this author’s goal to further present in this paper the inner

workings of CCN in relation with current various security flaws of the TCP/IP architecture and to explore some of the practical, ethical and juridical implications that this new architecture could raise, a subject which has been so far avoided as focus has been only on theoretical principles of actual functionality.

2. The Content Centric Networks Architectural Model

2.1 The theory behind CCN

Compared to the TCP/IP model, a packet in CCN does not address a location (IP) but the actual content, having no notion of host on its lower level.

The communication principal is based on the notions of *Interest* and *Data*, very similar to the http "get" and "response" actions.

A user would broadcast an *Interest* in a chunk of content and any node (holding information regarding the request) that would "hear" it would send a reply. The information would be transmitted only when there is a request for it and the *Interest* would then be consumed. Just as the DNS provides information on the location of each domain name, a similar top down hierarchical database would provide routing information to the closest content that matches any particular *Interest*.

As content can be hashed, that makes it unique and thus traceable and prone to lookups just as IP addresses are.

In order for the whole process to work, Jacobson et al. (2009A) have identified that CCN routers must include three crucial components:

- a content store (CS)
- a pending interest table (PIT)
- a forwarding information base (FIB)

For each pending *Interest* the node will verify the Content Store and if the information is there, it will return it to the interface that requested it. Else ways, it will forward the interest to the next hop which is determined by the Forwarding Information Base and it will create an entry in the pending interest table so that the router will know where to send it when the specified chunk arrives.

Unlike the existing system though, CCN can query multiple sources for data (broadcast), leaving traces of where the *Interest* has originated from without the risk of creating loops. When *Data* flows toward the source by following the traces left in the routers by *Interests*, by consuming the *Interest* it lets the router know that any other *Data* that might come shortly after can be discarded. Also, a time to live (TTL) implementation will discard any leftover *Interests*, leaving to the source the responsibility to retransmit the request. This system assures that the fastest route is taken, which in most cases will also probably be the shortest path.

As Detti et al. (2010) best describe the routing functionality of this “*advertised-based*” architecture at the protocol level by mentioning that instead of having routers advertise their IP subnets, they will advertise the names of the content they can provide access to.

2.2 CCN example

We can assume that there are two hosts, A and B, each of them requesting the same content, e.g. a particular YouTube video, with host A having been the first one to send an *Interest* in the content, as seen in *Figure 2*.

In this case, the request reaches the first node and is looked up similar to current URI's through the Forwarding Information Base, with the request being forwarded to the next router and so forth until it reaches the signed domain to which the information belongs to, in this case *youtube.com*, that will know to provide the exact match to the requested information. The interest is also added to the Pending Interest Table along the way so that each router will know where to send back the reply and to avoid sending duplicate requests.

As the requested chunk of information travels back, each router will consume the *Interest* request and will store in the Content Store the particular chunk which was requested.

For this particular case, as we are dealing with a lot of content, the dynamic storage allocation algorithm presented by Carofiglio et al. (2011) for CCN seems quite appropriate, as this way the cache will be allocated dynamically and information stored until the space is needed by new data but taking also into account known cache algorithms like Least Frequently Used (LFU) or Least Recent Used (LRU).

As host B sends a request for the same information, as soon as the Interest reaches Router X, the router performs a lookup and

discovers that the content is already stored in its Content Store and therefore replies immediately with the data, without forwarding the request further so that it will reach the actual original source.

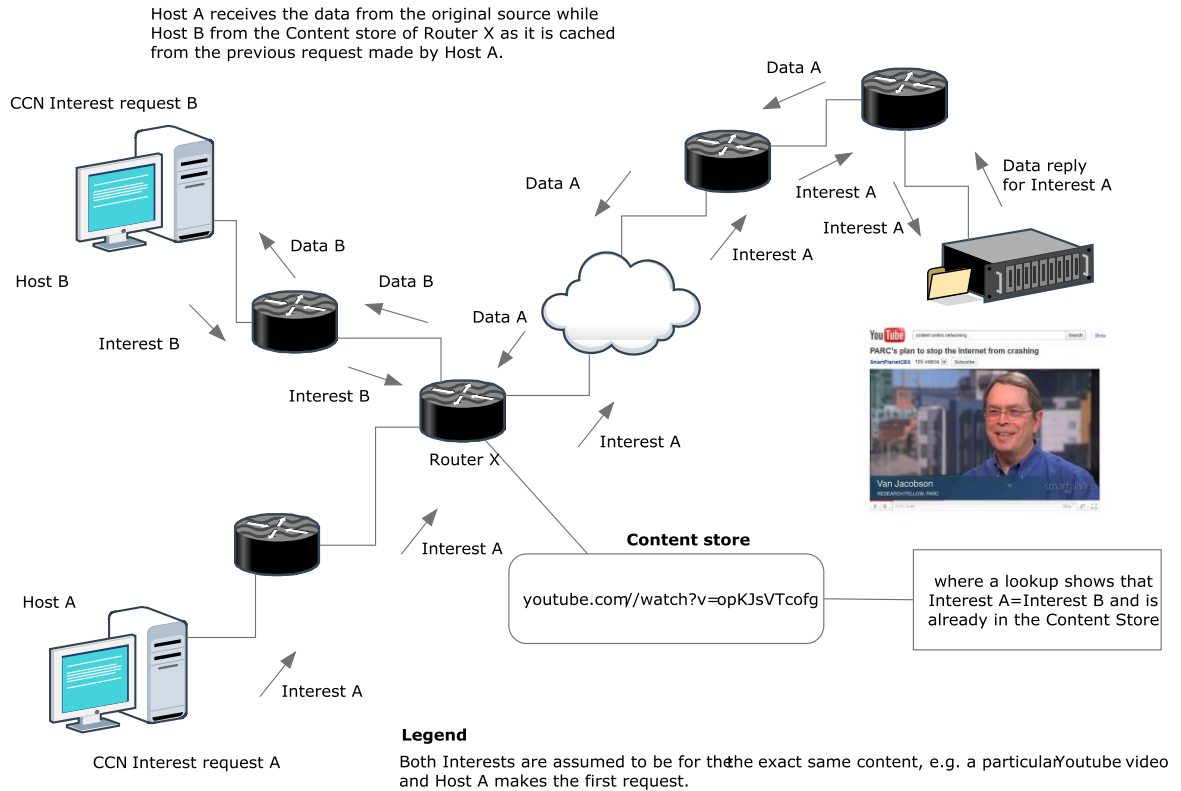


Figure 2: CCN functionality example excluding the digital signed authority identification request (e.g DNS like authority)

One can see from this example that Host B not only that it receives the same data as Host A but its requests never reach the original server where the data resides (e.g. in this case the YouTube server which hosts it), as somebody else close by in the topology has previously requested it.

Not only that Host B has a better quality of service than Host A, which was the first one to pull the data, but the YouTube server saves energy by responding to only one user therefore peering/bandwidth costs are reduced as the second user was able to retrieve the same data from a closer source.

While the principle is good, energy has been saved, efficiency achieved and the overall experience of the end user has improved,

privacy and security issues and the legal implications of caching copyrighted data will be explored separately by this author in further chapters.

2.3 Hardware requirements and measured efficiency

While the still experimental CCN daemons² can be run on any Unix/Linux or Android OS based systems, the caching system on which CCN is based and takes advantage of assumes future ISP routers with appropriate resources to handle the extra memory load. Arianfar et al. (2010) have proved that such a router is economically feasible and that in order for it to handle a 10 seconds long full cache at 10 Gb/s, it would need to have a 100 Gb DRAM content store and 324 Mb reserved for the index table, the last preferably made of SRAM chips in order to benefit from their energy saving and access speed features which this type of technology advertises. Even by using slower access DRAM chips but with an increased capacity, in this case 12 GB, it would use only up to 100kWh per year (excluding other router components).

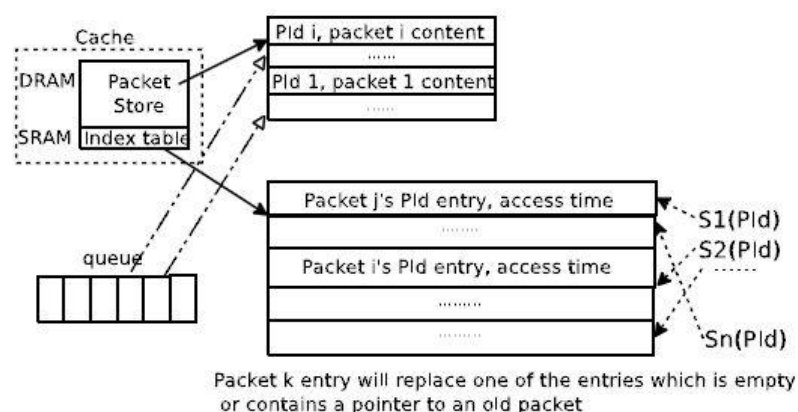


Figure 3: Structure of cache and queue system as detailed by Arianfar et al. (2010) for CCN routers

² A daemon is a computer program that runs in the background

Contrasting the statistics offered by Jacobson (2009A) regarding the efficiency of CCN over TCP (Figure 4) with those offered by Akamai about their bandwidth usage (Figure 5), we would find out that actually already at least 20%³ of the Internet's web traffic is using some form of caching, at the application level though and provided as a paid service to content creators and that CCN is optimized a-priori for this kind of traffic.

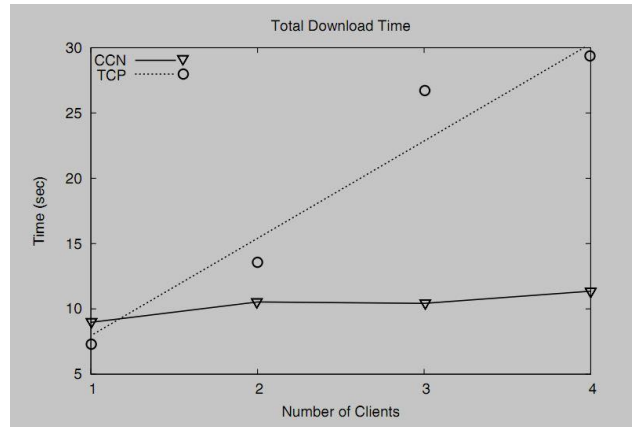


Figure 4: Measured efficiency of CCN compared to TCP/IP in a multi-user environment as experimented by Jacobson (2009A)

CCN is indeed outperformed by 20% by TCP when dealing with single requests for unique content. However, when multiple requests are addressed, the likeliness of a bottleneck decreases exponentially and the speed remains constant while for TCP it decreases as the number

of host increases, as it cannot handle simultaneously all of the connections. While the simulated traffic statistics backed up by results published also by Lauinger

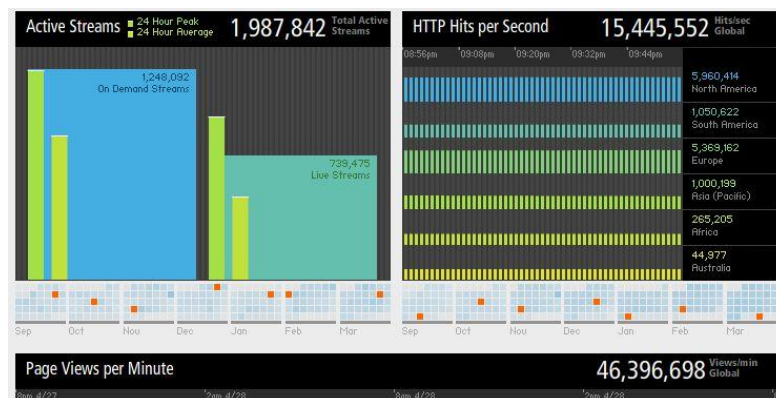


Figure 5: Real Traffic Statistics offered by Akamai

(2010) most certainly prove that CCN could indeed be a useful and economically feasible architecture in a real world deployment and that cached traffic is a happening reality through CDN's, a question

³ The statement belongs to Akamai and real traffic statistics are provided at www.akamai.com/html/technology/dataviz3.html [Accessed 20.04.2011]

which this author plans to explore in a further chapter is also raised: How supportive existing business models such as those which provide content delivery services will be toward implementing a free Internet architecture that makes their reason for existence obsolete?

3. Security in CCN

As described by Smetters et al. (2009), CCN was also designed to supplement the security flaws present in the current architecture. For the moment, security is implemented and bound to the containers where the content resides and at the network level through extensive mechanisms such as VPN's⁴ or TLS encryption⁵, which are still vulnerable to attacks such as "man in the middle".

By securing the actual content, some of these problems can be solved and it is irrelevant where the content is pulled from (be it from the original source, cache or some other server), as it can be authenticated as the original and desired content. Also, the theory states the content itself can be encrypted even though this would provide a processing overhead.

3.1 The credential system

A proposed method by Smetters (2009) for further securing the actual content is by authenticating the link between names and content, which would be generated into additional metadata and added to the chunk of information. By doing so, the content can be identified in other sources irrespective of its name, which is thus free to change.

⁴ Virtual Private Network

⁵ Transport Layer Security

As seen in *Figures 6 and 7*, the CCN naming structure is hierarchical and within each domain there can be multiple signed authorities, responsible for the authenticity of the content. Each authority can identify itself uniquely through a key, just as any piece of actual information made available by the respective authority, also with the possibility of having a distributed along the nodes credential system thus eliminating a single point of failure, which would also allow different attributes as described by Yu et al. (2011).

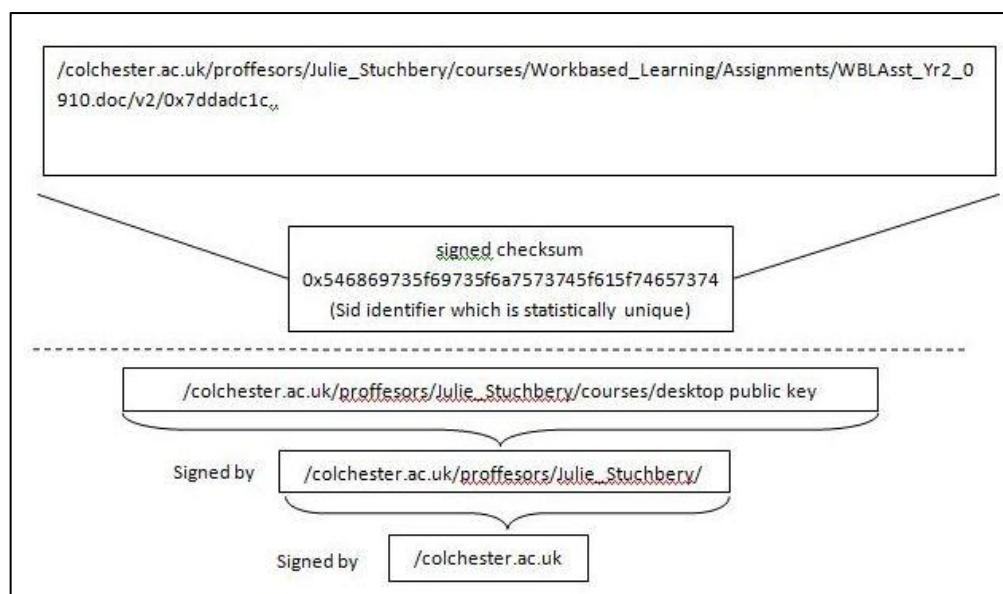


Figure 6: CCN can associate content namespaces with publisher keys. Example based on the Palo Alto white paper on Content Centric Networks. CoNEXT09, December.

Respectively, in *Figure 4*, a 32 bit credential table scheme is shown for K_{pc} , which in this case has the assigned value of A2E95CA6. The PC belongs to the Colchester Institute IT Department and has its public key stored at `pc.itd.colchester.ac.uk/pubkey` and the credentials at `pc.itd.colchester.ac.uk/creds`.

A credential query such as `pc.itd.colchester.ac.uk/creds-/query?subject= A2E95CA6` would show that an Android phone namely Kmobile, identified by the unique value - 87146546 has the credential to write to the Kprinter identified also by 39AF4668. By

taking into account the equation (K_{pc} , K_{mobile} , Write, $K_{printer}$) or ($A2E95CA6$, 87146546 , Write $39AF4668$) we discover that K_{mobile} has the permission to write or print to $K_{printer}$.

Credentials		
Issuer	subject	attributes
A2E95CA6	78C23435	Read pc.colchester.ac.uk/file2
A2E95CA6	87146546	Write 39AF4668
A2E95CA6	87146546	Read
A2E95CA6		Write
A2E95CA6		Read/write pc.colchester.ac.uk
B13CF329		Write printer.colchester.ac.uk
A..		Execute pc.colchester.ac.uk/file5
3..		Read
		Execute

Figure 7: Example of how a credential system would work under the Yu et al. (2011) concept.

A potential problem would be that the flooding queries which are made in order to discover the appropriate certified resource could impose an overhead on the network. However, a solution has also been proposed by the above mentioned researchers, where if the hash values are known, it is possible to locate the resource by using a DHT routing algorithm similar to that which is now used in P2P systems.

As the Internet is not only based on the principles of retrieving data but also on publishing it, a "on-demand" publishing model is incorporated in CCN, the theory behind it having been made available by Jacobson (2009B) in his research paper concerning VoIP over CCN, where he shows that it is possible to construct the name of a

desired piece of content beforehand.

Not all issues concerning the naming scheme and trust authorities have been fully addressed yet. If the architecture imposes a mandatory RCA (Root Certificate Authority) than this would require extra costs for every publisher as they will have to purchase them. As detailed by Kutscher et al. (2011), a possible solution in this case would be the use of a *Simple Distributed Security Infrastructure (SDSI)*⁶ or an alternative where publishers would provide their own resolution system. However, as there will be no central authority involved and SDSI has only been used in closed solutions or for academic purposes and demonstrations, there is no knowhow on how to implement it at such a scale.

There is also the question if existing CA providers will actually support an open and free system that threatens to make their business model obsolete, somewhat a similar situation with current Content Delivery business models, which may lose profits or even their reason for existence in a CCN architecture, as it was mentioned in the previous chapter.

Another issue raised at the Dagstuhl Seminar in February 2011, a conference where academics from all over the world working on various content-centric architecture projects have met, was if a credential system based on a legacy where only trusted authorities can generate keys and not the publishers themselves can become a tool for repressing free speech. The issue was raised for another content centric architecture, namely ICN, but the question addresses

⁶ SDSI is a prototype implementation for a public-key infrastructure developed at MIT with support from NASA and DARPA, which has the role of establishing identity without Certification Authorities. More information is available at <http://groups.csail.mit.edu/cis/sdsi.html>

just as well potential misuses of CCN and this author will extrapolate upon a hypothetical given case further along this paper (e.g. WikiLeaks).

3.2 Security concerns

As CCN is still a highly experimental technology and the only publicly made implementations (available at Project CCNx⁷) are related only to chat and VoIP daemons, many of the security concerns are purely theoretical and they deal mostly with architectural concerns, as none have been actually tested yet.

Based on the existing published data and through theoretical extrapolations and also some experimentation, Lauinger (2010) has identified some important security concerns, some of them which could overlap with issues of *privacy*.

3.2.1 Caching

As CCN is entirely based on the notion that content can be cached along the way, this also means that potentially documents, emails or even private conversations over chat or VoIP can leave traces along the routers which lay between the communicating hosts. Assuming that for the time being there is no mechanism in place to immediately discard such data from the cache (e.g. an “immediate delete” flag field for each chunk of content) and that there is no encryption service of the actual content, it can easily be retrieved. Even if the content is not complete, it can still be interpreted or approximated by

⁷ Project CCNx is a long term research program intended to develop and expand the CCN architecture. The project is sponsored by the Palo Alto Research Center (PARC) and the source code is currently licensed under the GPLv1 and LGPLv2.

using different cryptanalysis methods such as frequency analysis. Queries can be made through *Interests in a* very similar fashion to the MySQL injection attacks and by timing the responses to each node, one could potentially find the caches where the content was stored and to retrieve the whole content even after the exchange has already happened.

This can prove more dangerous than the current “man in the middle” attacks as one does not have to be present and record the entire transfer but can easily do it after the transfer already took place, as the information will still be stored along the way in the nodes.

3.2.1.1 *The unencrypted privacy trail*

Even though this author is not fully aware of the technical or ethical implications that a “delete data” field would have in the response Data packets, he sees no reasons for which this could not be applied/activated at the application level, thus mentioning that the data in question should not be stored in the cache along the nodes.

One may argue that such a field is not needed as the encryption overhead could handle the actual potential misuse of information (making it useless to an attacker which does not have the appropriate decryption key) and that by implementing it one renders the whole architecture vulnerable, as some copyright holders could decide to use it extensively for all their content, thus actually undermining the benefits of the entire architecture as efficiency was shown to be less than that of TCP/IP for single hand requests.

A “private” or “delete” field as seen in *Figure 8* could also be over-used for advertising/monitoring purposes by organizations, depending on the answer to the question if web pages will be treated and retrieved as a single object or not. While the end user definitely benefits from retrieving content faster, companies will lose the ability

to fully monitor user access to their content, especially by those for which advertising is a business model.

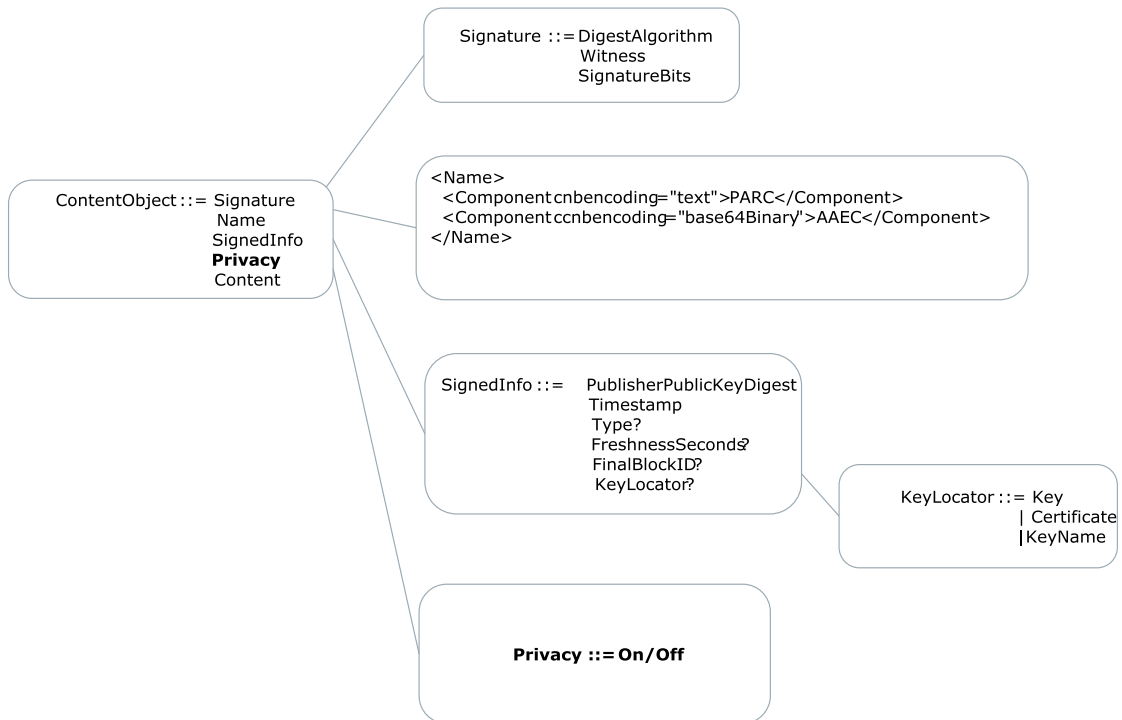


Figure 8: The structure of a data packet in CCN with the inclusion of a Privacy field

If, for example, only a fraction of the 5 million views are registered for the newest and most popular video on YouTube, as everybody is retrieving it from nearby caches in the nodes, than the company cannot estimate and promote its advertising potential even though the information is disseminated with less bandwidth costs. Privacy⁸ would be gained in this case though but in the detriment of the content provider.

There is also a potential for misuse by individuals, groups, organizations or governments if duplicate information would be created (hash wise) but with the sole difference between the original

⁸ Privacy is used here with the meaning of freedom from unwarranted monitoring and social profiling

and the “fake” consisting of this “privacy” field. One could use multiple requests (e.g. botnets) from various places to eradicate the original information from the caches after the original source has been taken down by requesting only the fake information.

Arguing from the privacy point of view⁹ however, there is absolutely no reason for having a CCN voice conversation which both parties desire for it to be private to be recorded, as basically this is what the architecture allows by storing it along the way in the cache. This statement is also true from the efficiency point of view as it will be a bi-directional conversation, therefore proving less efficient than TCP/IP as shown by Lauinger (2010).

The same functionality principle as above would apply to a private email which will be stored not merely just in the receiver’s inbox or the sender’s outbox but also in the network, along the way between the two. Even if deleted by both parties, the email would still exist in the network until the data from the caches would be overwritten. Such an outcome would transfer the security and privacy issues into the hands of users and the creators of content, which will be solely responsible for implementing stronger encryption algorithms to overcome the risk of having private data stolen from a cache. However, no encryption algorithm can fully guarantee that the content will not be decrypted in the near future with the appropriate knowledge and processing tools. Even 128 bit AES encryption has been proven that it can be cracked as of recently, as shown by Gilbert (2009).

⁹ Privacy is used in this context as viewed by Froomkin (2000), meaning “informational privacy” - “the ability to control the acquisition or release of information about oneself”

The risk of such a scenario happening is extremely low but not at all farfetched if one bears in mind that an attacker could use the processing power of a botnet or, in the case of governments, that of a supercomputer such as the new projected \$895 million NSA¹⁰ one (US National Defence Budget Estimates 2011), scheduled to be built until 2015 and expected to be faster than some of the world's supercomputers put altogether.

It is to be noted that the example above is only hypothetical, with the sole intent to blatantly show that the tools to perpetrate such an attack are real and that any new communication architecture must be constructed under the ethical assumption that free speech should also be a legacy for future generations, irrelevant of ideology or political discrepancies. This author does not know though the purpose of the NSA supercomputer which constitutes classified information and its power has been extrapolated based only on the specified wattage consumption of the project from the unclassified budget documents (*Figure 9*) which is supposed to be 60MW, compared to the 6MW required to function by the supercomputer Blue Gene or the 2.35MW which IBM Roadrunner uses.

¹⁰ The National Security Agency (NSA) is a cryptologic intelligence agency of the United States Department of Defence (Definition provided by Wikipedia) [Accessed 20 April 2011]

UNCLASSIFIED											
1. COMPONENT NSA/CSS DEFENSE		FY 2012 MILITARY CONSTRUCTION PROGRAM						2. DATE February 2011			
3. INSTALLATION AND LOCATION Fort Meade, Maryland				4. COMMAND NSA/CSS				5. AREA CONSTRUCTION COST INDEX 1.00			
6. PERSONNEL STRENGTH IC Community Installation		PERMANENT			STUDENTS			SUPPORTED			
a. AS OF		OFF	ENL	CIV	OFF	ENL	CIV	OFF	ENL	CIV	
b. END FY					x	IFIED					
7. INVENTORY DATA (\$000)											
A. TOTAL ACREAGE										TBD	
B. INVENTORY TOTAL AS OF DEC 2010										TBD	
C. AUTHORIZED NOT YET IN INVENTORY										0	
D. AUTHORIZATION REQUESTED IN THIS PROGRAM										860,579	
E. AUTHORIZATION INCLUDED IN FOLLOWING PROGRAM										399,939	
F. PLANNED IN NEXT THREE YEARS										431,000	
G. PLANNING AND DESIGN COST										35,000	
H. REMAINING DEFICIENCY										0	
G. GRAND TOTAL										895,579	
8. PROJECTS REQUESTED IN THIS PROGRAM:											
CATEGORY		PROJECT NUMBER		PROJECT TITLE			COST (\$000)	DESIGN START	COMPLETE		
141		TBD		HIGH PERFORMANCE COMPUTING CENTER (FY12)			\$29,640	Nov 2010	Sep 2011		
				PLANNING AND DESIGN (FY12)			\$35,000				
9. FUTURE PROJECTS:											
a. INCLUDED IN FOLLOWING PROGRAM											
CATEGORY		PROJECT NUMBER		PROJECT TITLE			COST (\$000)				
141		TBD		HIGH PERFORMANCE COMPUTING CENTER (FY13)			\$399,939				
b. PLANNED IN NEXT THREE YEARS											
CATEGORY		PROJECT NUMBER		PROJECT TITLE			COST (\$000)				
141				HIGH PERFORMANCE COMPUTING CENTER (FY14)			\$431,000				
10. MISSION OR MAJOR FUNCTION Agency activities are classified.											

Figure 9: Unclassified US Defence budget estimates detailing the plan for an NSA High Performance Computing Centre totalling 60 MW of technical load.

3.2.1.2 Legal implications in the EU

The legality in the European Union of storing cached information in the nodes should be protected under the same principles enounced by the resolution of the European Court of Justice in the case of Google vs. Luis Vutton (2010), which states that "a service provider cannot be held liable for the data which it has stored at the request of an advertiser, unless, having obtained knowledge of the unlawful nature of those data or of that advertiser's activities, it failed to act expeditiously to remove or to disable access to the data concerned", as detailed by Manolea (2011).

However, the interviewee mentioned above, which is specialised in Internet law and privacy and also heads the Association for Technology and Internet from Bucharest, has recognized the problem

of transforming ISP's, which have been so far forwarders of information into information holders. This move might force them to take appropriate measures under the pressure of copyright lobbyists in order to be able to enforce the policy of removing or disabling access to the data in question once having knowledge of the data having been stored in their cache.

The current legal climate regarding enforcing content filtering policies for ISP's viewed as mere forwarders of information is still uncertain in Europe, with the case of Sabam vs. Tiscali (2010) having been referred for resolution to the European Court of Justice (ECJ).

However, given the recent report of the European Commission (2011) on the topic of *net neutrality*, which discourages the common practice of offering unequal Internet traffic by ISP's, this author can assume that any resolution of the ECJ will take into consideration both this report, as well as the potential implications which could undermine the Articles 8 and 10 of the European Convention on Human Rights (freedom of expression and the right to privacy).

The juridical implications of a Content Centric Network architecture in Europe if we presume that the ISP Tiscali shall win the above mentioned case could lead though to a different outcome, as a new case could be constructed under the assumption that an ISP is now an information holder and must have the means to delete or restrict access to stored material.

3.2.2 In content we trust

As opposed to the TCP/IP model where we tend to trust only the source of the content by its certified locality (e.g. trusting that we are actually retrieving content from google.com irrelevant to what that content contains), CCN thanks to its caching ability means that we must trust the content itself as the locality becomes obsolete (e.g. the content that we think is being pulled from youtube.com actually is

retrieved from the cache of the nearest hop as somebody else happened to make the request for the same material beforehand). This issue raises the question of how can one guarantee through a formal mechanism that the content retrieved is actually the latest content and not one at its previous version.

While the technical feasibility has not been explored yet in any of the read published papers regarding CCN, this author has thought of the possibility of having a "time stamp" attribute applied to content, thus permitting a lookup for the most recent content (e.g. retrieve closest content to current date matching the specified identifier).

The persistence of data in the network (as the network **is** the data¹¹ in content centric architectures compared to the network only transporting the data in the current architecture) has equally the potential to encourage free speech or to be used maliciously for keeping unwanted or illegal data in the network even after the original source has disappeared.

By scheduling *Interests* in such a way that they will arrive just before the data in the cache will be discarded, someone could potentially keep the information in there or help propagate it to other nodes by using botnets to request that particular data to those nodes.

In the free speech scenario, this could prove beneficial as any mechanism of repression might be able to shut down the original source but will not be able to stop the information from propagating through the network. On the opposing side, lies the position of some law enforcement agencies or national security institutions, which would prefer to stop any illegal or confidential information from

¹¹ This affirmation is based on the context that some information will co-exist in the network as opposed to having it at the border of the network and addressed each time by locality or in transit.

potentially remaining in caches throughout the nodes.

3.2.2.1 The WikiLeaks example

A perfect scenario for CCN or other content centric based architectures is one which allows us to extrapolate on a real world example, that of the WikiLeaks¹² website (Figure 10).



Figure 10: WikiLeaks website

This particular case is relevant as no other website in the history of the Internet has generated so many debates on the freedom of speech issue in relationship with the actual Internet architecture, functionality and jurisdiction issues.

One view is that the website encourages transparency, responsibility and accountability from governments around the world and that even though the Internet is a global place, it should remain a free one and under no standardized jurisdiction, thus always allowing free speech and the possibility to propagate information.

However, cultural, political and ideological differences have led to

¹² WikiLeaks is an international non-profit organisation that publishes submissions of private, secret, and classified media from anonymous news sources, news leaks, and whistleblowers. (Definition retrieved from Wikipedia) [Accessed 27 April 2011]

opposing views, not just by governments and regimes worldwide but also among members of the same administration, country, groups or even family members. Their counter arguments support the view that the website acts as a spy, that it favours transparency over privacy, with some official US statements (BBC, 2010) even accusing the organization of terrorism as it represents "an attack on the international community".

In such a difficult context to comprehend, where the issue in question has not found yet a resolution and opinions are divided, we can clearly see that the current TCP/IP architecture has actually helped the website to remain available after multiple shutdown attempts and to continue to sustain its self-declared mission of making the information available to everyone.

In a similar CCN scenario, where information can be blacklisted through a list of censored names propagated through the nodes, as a potential countermeasure described by Lauinger (2010) to storing and propagating unwanted information in the caches, the website or its mirrors could have easily been shut down. This scenario implies that some degree of international regulation and standardization will exist between the nodes, permitting the enforcement of the blacklist everywhere. If the mechanism in question exists but is enforced at the local level though (country, region, ISP, etc), the result could be the same, limited or no access at all to the wiki leaks resources in that area.

Assuming that certification is mandatory in a content-centric network architecture and that there is one centralized certificate generating authority (CA), the logical inference resulting from these statements is that the overall values under which the Internet would function will be mandated, imposed and enforced by the cultural and judiciary

values of the country where the CA resides.

It is to be noted that certification attests the identity of the data through a public key in this case and not merely the location (e.g. website).

This would mean that if WikiLeaks is deemed inappropriate by the country where the CA resides, it could potentially be shut down by making their certificate obsolete.

If we assume no mandatory certification but a centralized certificate generating authority none the less under the same scenario as above, the implications are different as mirrors would be created but the information could not be certified as belonging to WikiLeaks, allowing for the creation of false duplicate content.

One particular scenario has been raised (Ethics, privacy and security in content centric architecture group 2011) about the potential use for censorship of content centric networks.

If there is a mechanism in place which is able of identifying each piece of information as unique (hash or certificates) and forwards data based on the actual information itself, how easy would it be to filter "unwanted" data?

Most likely, even by slightly modifying the data, it can still be identified through similitude analysis.

Even if the respective data could be moved to another location this would be irrelevant under a CCN type architecture as one would only have to ban the actual information and not the location, which is easily mirrored in today's TCP/IP system. Does this new type of architecture actually enforce at its lower levels the possibility for having the best content filtering systems yet developed?

3.2.3 The geolocation of anonymity

“**whois**¹³ tracking us” is another issue which CCN promises to solve, encouraging anonymity. The implications will be explored though in depth throughout this subchapter.

3.2.3.1 Understanding geolocation

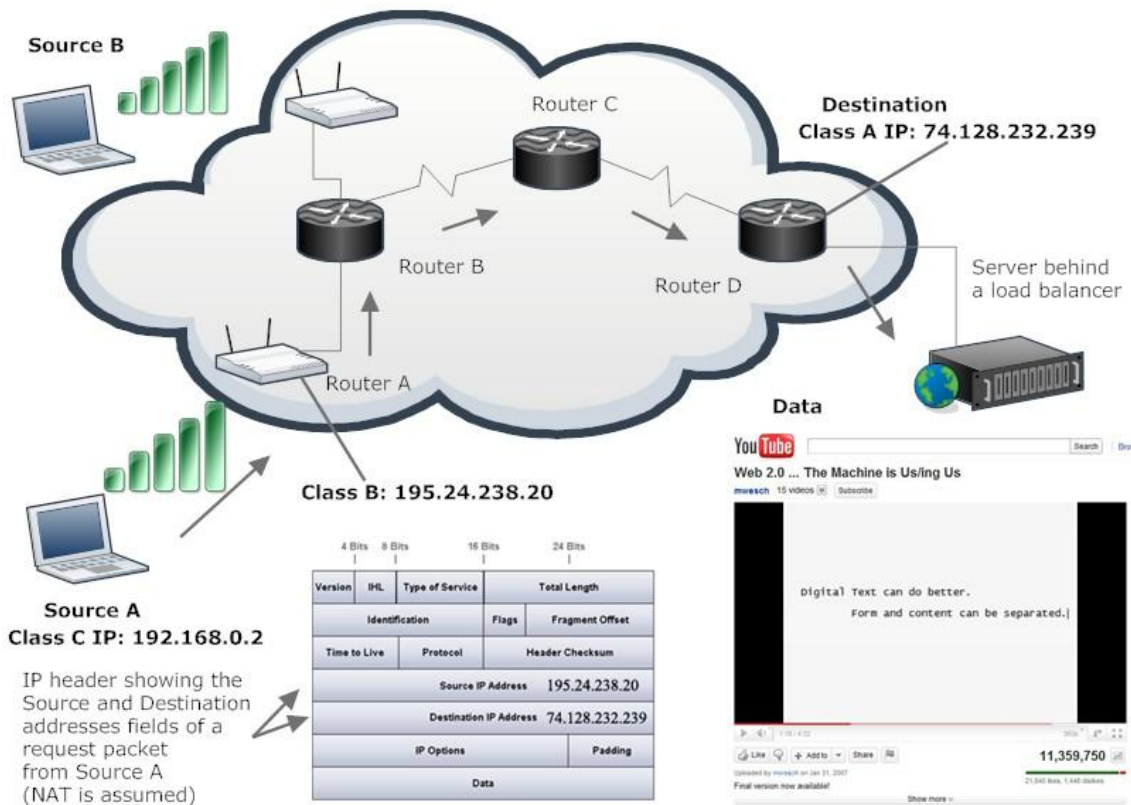


Figure 11: Geolocation as viewed from the TCP/IP model

In the case of the existing Internet infrastructure, the *Destination* web server knows of the *Source* which has requested the data as it will know its unique IP address. Even using NAT translation and Class C subnets in the local area network, the address of the border

¹³ WHOIS is a query and response protocol that is widely used for querying databases that store the registered users or assignees of an Internet resource, such as a domain name, an IP address block, or an autonomous system (Definition offered by Wikipedia) [Accessed 27 April 2011]

wireless router of **Source A** as shown in *Figure 11* (above) will still be known, as otherwise the YouTube server in this particular case would not know where to return the data.

As ICANN¹⁴ oversees the distribution of IPv4 and IPv6 addresses through IANA¹⁵, there is a database with the range of addresses given to the various ISP's around the world. Through various queries regarding the source IP address, the YouTube server can find the geographical location of the data seeker with a margin of error up to a couple of kilometres if we assume optimum conditions.

This *service* is used for many security or business reasons, from tracking unlawful activities to deploying localized marketing strategies or optimizing traffic through the use of various statistics and analysis gathering tools.

Many companies use it to restrict access to their content (e.g. BBC iPlayer which is only available in the UK but not abroad, last.fm which is available in certain countries or US available only entertainment provider Hulu). The reasons usually invoked for censoring access based on geolocation can be bandwidth related - as most of these companies are creating heavy traffic - or copyright related, as various content owners must consent to allow distribution for specific regions.

Such services which rely on geolocation can though be easily circumvented even by an inexperienced computer user by using

¹⁴ Internet Corporation for Assigned Names and Numbers (ICANN) is a non-profit corporation with the headquarters located in the US which oversees a number of Internet-related tasks. (Definition based on Wikipedia) [Accessed 27 April 2011]

¹⁵ The Internet Assigned Numbers Authority (IANA) is the entity that oversees global IP address allocation among others tasks and is operated by ICANN. (Definition based on Wikipedia) [Accessed 27 April 2011]

proxies¹⁶ or VPN's, thus masquerading the real IP address of the user and having the request originate from a different place (IP).

3.2.3.2 Geolocation in CCN

As it can be seen in the scenario from *Figure 12*, CCN is not based on an end-to-end protocol like TCP/IP, thus rendering geolocation useless, at least in the current development stage.

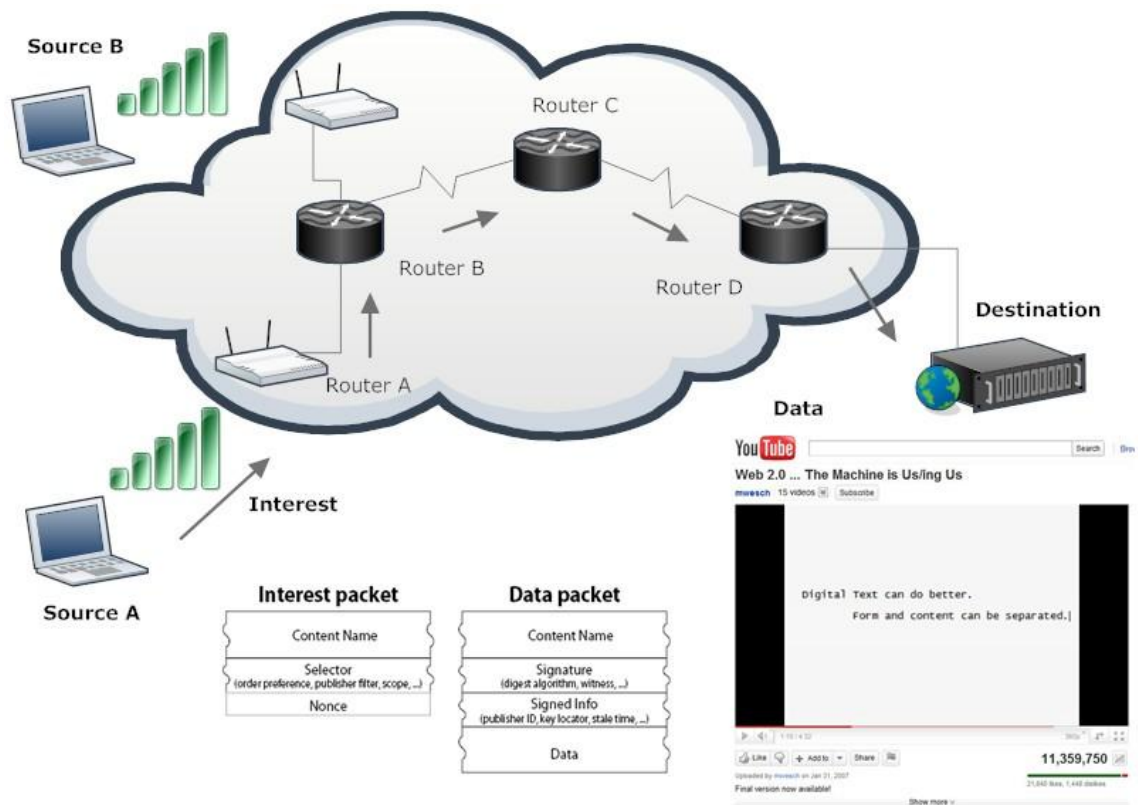


Figure 12: Geolocation as viewed from the CCN model

In this CCN scenario, Source A requests the information from Router A, who forwards the request to Router B, who forwards it further to Router C, which does the same thing until it reaches the actual place where the data lies. However, at the destination, the only thing that

¹⁶ A proxy server acts as intermediary between client requests and the end destination

Router D knows is that Router C has forwarded a request for a particular YouTube video, with no knowledge of where the request has originated from. Each router remembers only the interface on which the *Interest* came from, thus knowing to delete the *Interest* and return the data through that interface when it arrives.

We can assume that if host B asks for the exact same information within the timeframe that the chunks of information are still present throughout the caches of routers B and C, than the YouTube server not only that it will not know who was the source which asked for the information but the second request will go unnoticed completely as it will never even reach the server.

This can potentially affect current load balancing techniques used by organizations as well, since without knowing the end destination the data must return through the interfaces which have forwarded the Interests and who know the way back to the requesting source.

3.2.3.3 *The Nichomachean Ethical Implications of Geolocation*

While used extensively by companies around the world for various reasons as mentioned before, it is safe to assume that most websites from the Internet are actually using at least some form of geolocation services by simply implementing traffic counters. As an example, Figure 13 is showing services using geolocation on the BBC website which were identified and blocked by using the Ghostery extension for Google Chrome.



Figure 13: Advertising and tracking services using geolocation on the BBC website

Many questions have arisen though from the ethical and legal points of view, as the service is bound to overlap in certain circumstances

with the right of data privacy (Vica, 2011).

How will companies use this kind of information in correlation with other data and what is the potential for misuse? Is behavioural tracking for the purpose of targeted advertising an ethical method? Where will the company store the data provided by a user and how safe is it in their hands?

At least some of these last security concerns can be answered by the recent attack on the Sony network, which resulted in the stolen data for 77 million members (Yin, 2011).

In California, the recently proposed Senate Bill 761 (2011) is trying to enforce consumer protection by making companies which collect, use or store information from a consumer to be legally bound to provide an opt-out mechanism. The bill though somewhat contradicts with other current US legislative efforts of having ISP's log communication data.

In the EU, the Directive 2006/24/EC regarding data retention mentions that ISP's must retain from 6 months to 2 years the following:

- the source of a communication (IP address);
- the destination of a communication (IP address);
- the date, time and duration of a communication;
- the type of communication;
- the communication device;
- the location of mobile communication equipment.

While the directive has been adopted in certain EU countries, others have found that it violates their constitutional rights of privacy,

confidentiality in communications and freedom of speech (Romanian Constitutional Court Decision, 2009).

In the UK the retention of communications data is done under part 11 of the Anti-terrorism, Crime and Security Act (2001) for the purpose of national security.

Given a CCN scenario however, in EU countries where the above mention Directive is enforced through national laws, the element of geolocation becomes somewhat obsolete, as one would be identified by the actual content which is retrieved, as that would be interpreted as the destination instead of the IP address.

While discussions are still pending and views differ between member states if the data retention directive under its current form violates human rights, under a content centric architecture this would most likely be the case as it would contravene to both Articles 8 and 10 of the European Convention on Human Rights since the ISP would have to retain information of the actual content of a communication rather than its mere destination and source.

A point raised by Manolea (2011) regarding increased anonymity through the lack of a geolocation mechanism to identify the requesting source of information in CCN is that law enforcement agencies might oppose such an architecture which favours the privacy of the individual versus the controlling actions of the State.

If we assume no interference from any government agency or legal conflicts which can arise from the retention of data in countries where such a practice is enforced by law, there is still the question of how will content publishers react in adopting and supporting a content centric networks infrastructure if they cannot profit from tools they have been accustomed to and which are an integrated part of their

business model.

In California, the Senate bill 761 (2011) mentioned earlier is opposed on common grounds by powerful organizations among which we can find Google, Yahoo, Facebook or Time Warner.

3.2.3.4 Potential solutions to the privacy vs. reachability paradigm

While geolocation may not work in CCN, current methods for localizing the source of a communication regardless if the originating IP address is known or hidden behind VPN's or proxy's in TCP/IP do exist, based on analysis of the latency between packages correlated with known topology maps.

Huffman et al. (2005) have patented for the NSA such a method, which could be applied just as well to CCN, as long as there are at least some chunks of data retrieved from the original source.

While the option of having a privacy field has been discussed in previous chapters, thus forcing permanent retrieval from the original source, another method envisioned by this author would be by simply generating unique objects after each request at the application level, thus allowing some packets to arrive to the source of the content to retrieve them while the rest of the data will be pulled from nearby caches.

A query similar to "whois" can be made to a database with the values between packets and the location in the topology of the server which has received the *Interest* and in the case of web browsing a cookie would suffice to identify a unique user, thus establishing identifiable information.

This method would still allow geolocation services to function under CCN and it could also potentially solve the problem of having a record of where the request originated from for those providers whose

business model is dependent on user data.

Lauinger (2010) offers a different method of accountability which is enforced at ISP level for users but diminishes completely the privacy and anonymity which CCN seemed to offer. His scenario assumes that ISP's can provide users with unique signing keys which can be renewed just as dynamically IP addresses are. The metadata that they would contain can identify both the ISP and country or location of the user.

4. Implementation

For all current information-centric networks which are under development in various institutions around the world, be it CCN, PSIRP, 4WARD-NetInf or Dona, one of the major issues beside actual functionality is the capability of the architecture to reflect the possibility of a real world deployment.

In an ideal scenario, academic arguments about efficiency, privacy or even the natural tendencies of evolution in networks as explained by Dorogostev et al. (2002) could suffice to support an information centric architecture. However, in a complex business environment various incentives have to exist for all parties involved in the process, meaning end users, ISP's, content providers and developers.

Ahlgren et al. (2011) approach in theory some of the incentives required by each category mentioned above, establishing that end users and application developers are most likely to actually adopt the change, as there would be no extra costs associated with it, at least for new users purchasing or wanting to upgrade their equipment. However, Ahlgren et al. (2011) argue that having end users adopt the technology might not be enough, especially since without the

benefits of intermediary caches there would be no performance differences from the current architecture or potentially even lagging behind it. This author has identified though, as it will be explained in a further subchapter, that there is still an incentive to implement such a technology at the user level as it promotes device interoperability.

As explained by Carrea (2010), the optimal solution for a large scale implementation would be at first the adoption of a content centric architecture overlaid over TCP/IP at ISP level, starting with Tier 1 networks and ending with Tier 3 as seen in *Figure 14*.

Such an implementation would reduce initial investment costs while still providing adequate services by minimizing bandwidth usage between the interconnecting peers.

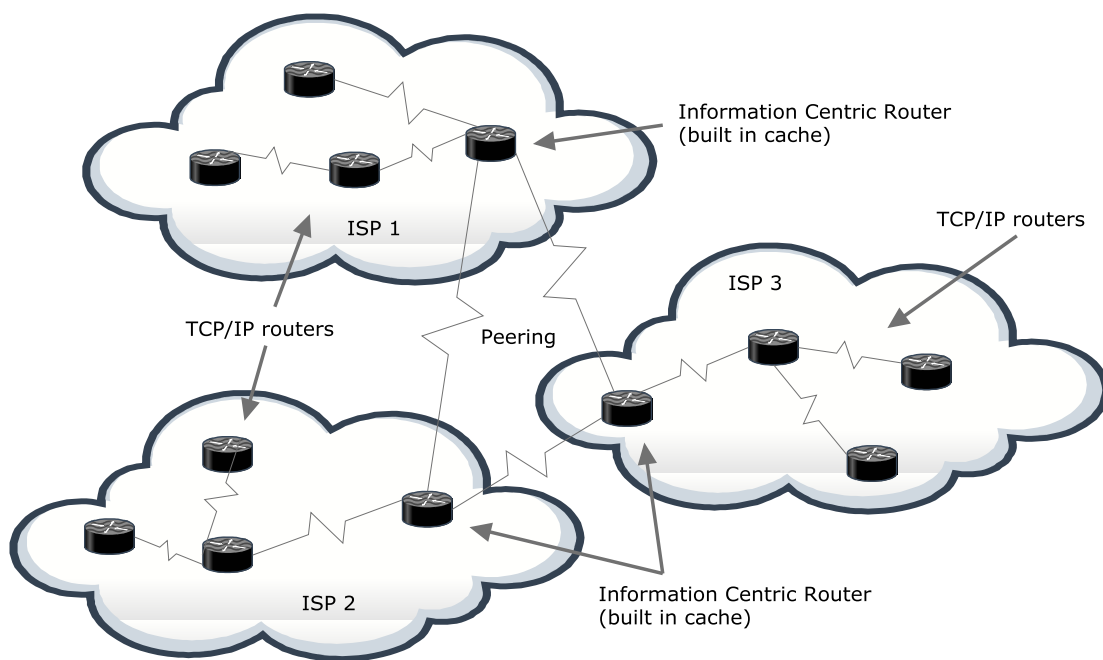


Figure 14: Localized ISP implementation of CCN

The likelihood of ISP's implementing an overall CCN like architecture will depend largely on the costs involved in creating the infrastructure and also on their vulnerability to the rigour of the law, as analyzed in

subchapter 5.2.1.2 - Legal implications in the EU.

Since many CDN organizations as Akamai have undertaken the investment costs involved in storing local content caches at ISP premises, thus enabling the ISP to provide a better service but without some of the costs that would have been associated with it, it is unclear if any will be willing to jeopardize such an economically fruitful relationship.

4.1 Advantages of CCN in a wireless mesh network environment

While for ISP's the deployment of CCN or any other content centric information architecture is yet uncertain and the adoption depends on multiple factors, this author has identified a potential future business model which could be supported by some equipment manufacturers and ISP's in a different context.

According to Internet World Stats (2011), the penetration rate for the Internet in the EU is at 58%, with many rural or even urban areas still having access to modem like bandwidth. A major downside toward increasing availability and speed rests with the major costs associated with deploying an adequate infrastructure.

However, given the recent EU proposal (2010) to free up the 800 MHz radio frequency band from January 2013 for WiMax or LTE, a new possibility arises under the form of wireless mesh networks, with the goal of helping to eradicate somewhat the notion of digital divide at least in Europe.

Since CCN seems to incorporate game theory elements within its conceptual structure, just as mesh networks do, it would be suited for such an environment, emphasizing access to local or duplicate content without consuming much external bandwidth.

While this CCN scenario (*Figure 15*) has not been proven yet in experimental conditions, the assumptions made by this author are based on the proved efficiency of CCN when asking for duplicate content, the principles of locality applied to sharing digital information in small communities as detailed by anthropologists Mesch et al. (2003) and the potential implications of releasing the 800 MHz spectrum which is currently reserved in EU countries.

Given the timeframe for achieving homogeneity at the national level within the EU in regards to implementing the EU proposal above mentioned and the decrease of the cost of storage (£/Mb), an information centric architecture could prove to be an incentive as well for ISP's willing to extend their range of customers and services without the heavy cost associated with building a wired infrastructure.

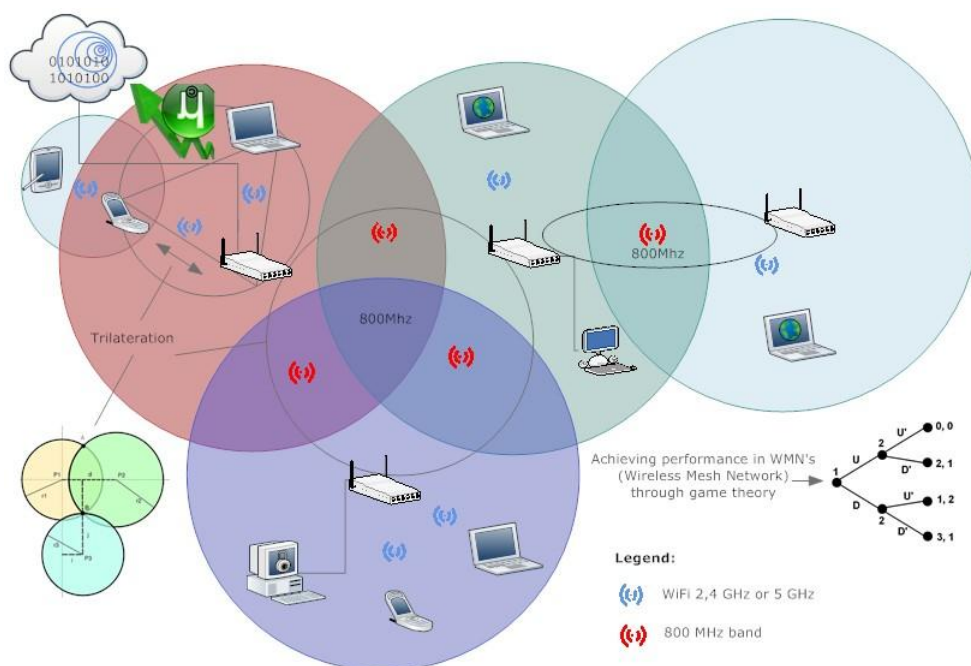


Figure 15: CCN in conjunction with wireless mesh networks as interpreted by the author

4.2 Device interoperability

A source of discontent for today's digital users is the lack of interoperability between devices in what should be an information ubiquity like environment.

While a user can have many devices, ranging from desktop PC's, laptops, PDA's, tablets, smart phones or other electronic devices which nowadays can store or transmit digital information (e.g. camera's, GPS systems, etc..), there are no easy means for retrieving the information from its closest point and between devices.

Because of the TCP/IP system, each device is forced to have a unique IP address in order to communicate with *likewise* capable devices.

This usually implies either a router which identifies or maps each device or having users manually set up the devices in order for them to become interoperable. An example is having two laptops, a business one and a home one, running a similar program with data which needs synchronizing. In order to retrieve the data and synchronize in a Wi-Fi environment, both laptops have to be either connected to a router, to an ad-hoc network or manually configured by providing IP addresses.

By using the same switching logic present for example in the Virtual WiFi utility¹⁷ and a broadcast, the communication between wireless devices in an information centric architecture could simplify the retrieval of data as a device could be able to ask for the information to all the devices in a certified network and retrieve it from the fastest source or straight from the device which actually holds it as

¹⁷ Virtual Wi-Fi is a project developed by Microsoft Research. It abstracts a single WLAN card to appear as multiple virtual WLAN cards thus allowing simultaneous connections to multiple networks. (Available at: <http://research.microsoft.com/en-us/um/redmond/projects/virtualwifi/default.htm>)

seen in *Figure 16*.

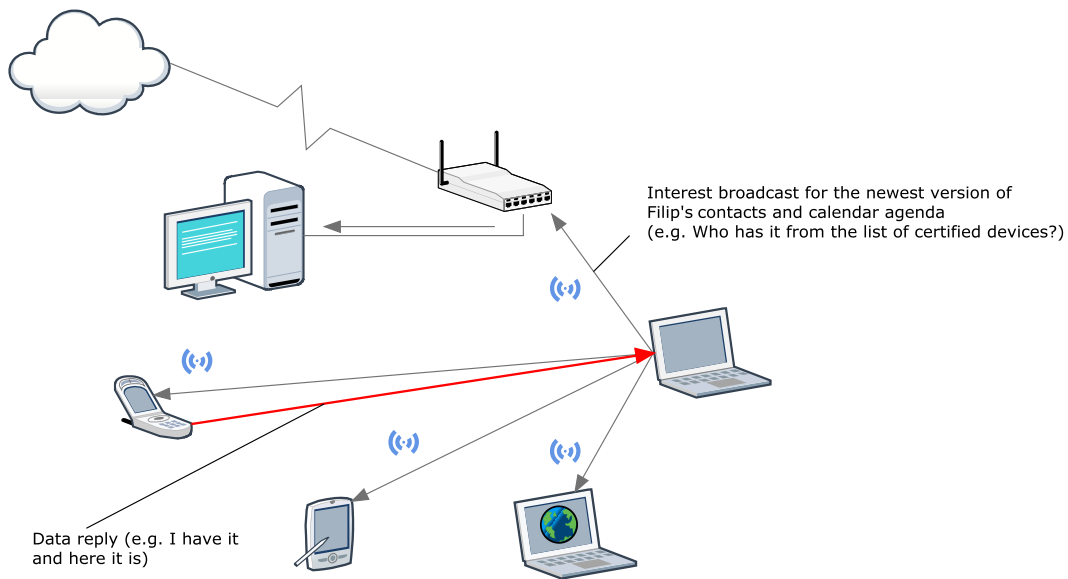


Figure 16: CCN like scenario showing interoperability between devices

If CCN protocols will be open source, then application developers will have a great incentive to implement them overlaid over TCP/IP, allowing users to easily interconnect devices and applications, a demand which increases exponentially as the number of devices powered by operating systems that one holds also increases. While not relevant¹⁸ from the statistical accuracy point of view - a reason for which this survey is not included in the methodology section - this author has found by surveying 50 people that 100% possessed at least two devices (computer and a mobile phone) and 25% of those more than two devices, with all of them desiring easier interoperability between the devices and their various applications (synchronizing contacts, addresses, bookmarks, etc). While the above survey can be easily repudiated from the scientific point of view, as

¹⁸ A survey for estimating the need for interoperability between devices for the UK population with a 5% margin of error would need at least 1035 respondents and for the Essex region at least 425 people assuming a margin of error of 15%. (Survey method obtained from *Sociological Surveying* by Pr. Dr. Septimiu Chelcea, ISBN: 973-9105-30-0)

its error margin would be more than 25% for a region like Essex (e.g. a population of 1,7 million) or potentially higher as the sample was chosen from people from urban areas using technology on a daily basis, it can be backed up by common sense in the case of users which are familiar with technology and whom also own or use at least two devices, thus having the need for such interoperability.

The above mentioned survey can also be sustained by a different argument, using as an example only the rate of iPad sales (20 million since launch, figures provided by Apple Inc. in the mass media), thus showing that users are already migrating toward a new range of mobile devices while still using personal computers. The iPad was chosen as an example in this case as it assumes that the user already has a PC/Mac in order to activate it and restore it to factory settings and it is also a device which cannot fully replace the functions of a traditional computer.

As we are dealing with a great range of system distributions (UNIX, Linux, Windows, Android, Symbian, IOS, etc), it is hard to conceive that manufacturers and developers will adopt a proprietary protocol that will increase their production costs and as an example we can use the case of intra domain routing protocols, in particular OSPF and the Cisco proprietary EIGRP. As indicated by Graziani (2008), OSPF is the preferred routing protocol in most cases. This is also due to compatibility issues, as it can be implemented between all professional routing devices and one does not need to have Cisco equipment or license it from the organization in order to benefit from the advantages of EIGRP.

The increasing number of gadgets can be in itself an incentive for developers and manufacturers to implement such a content centric

architecture, if only for the reason of allowing better local interconnectivity between *smart*¹⁹ devices.

5. Conclusion

As technology becomes an increasingly intricate part of our lives and an efficient gnosis tool, it is bound to be analyzed and subjected to the same ethical rigours as many other concepts with similar effects. Jonas (1979) has identified, even before the Internet or personal computers have materialized and become accessible to the general public that “modern technology touches on almost everything vital to man's existence -material, mental, and spiritual”.

While his statement was not specific in nature, the extent to which both technologies mentioned earlier have been adopted and the rate at which they developed have made them probably more prone to ethical debates than any other technology in the history of mankind, including the much controversial subject of nuclear power and its potential implications. This can be considered already a fact²⁰, supported by the sheer number of Internet related legal cases since the deployment of the World Wide Web in 1991 and the expansion of the Internet ever since.

Because of the global nature of the Internet and the interlaced relationship between power and information as interpreted in a chomskyan²¹ approach, the way the Internet tends to function seems to mimic free human interaction and communication, allowing at least in theory an equal access to its resources and the possibility of

¹⁹ Operating system and Wi-Fi capabilities are assumed

²⁰ Famous U.S. Internet Court Decisions can be accessed at www.internetlibrary.com

²¹ Noam Chomsky is an American linguist, philosopher and cognitive scientist whose views are that information helps shape public opinion and whoever controls the information flow can also control the process of manufacturing consent.

exchange at a cultural, intellectual and spiritual level immaterial to where a user resides. Information can be indeed manipulated, restricted, censored or altered and blocked maliciously but simultaneously it can also be disseminated and accessed as a whole through various mechanisms which the current network architecture allows, despite all the attempts to eradicate it.

As it was presented in the WikiLeaks case scenario, the morality of having information censored is not an easy task to process as it is bound by cultural differences in the interpretation of what constitutes freedom of speech.

The apparently abstract qualities accepted as universal (e.g. problem of universals) are in fact just as bound to locality as the current TCP/IP architecture, as opinions differ from country to country, government to government or at the individual level when one debates about what are the principles on which access to information should be achieved online and what constitutes appropriate information.

If one assumes that freedom of speech and the right to privacy are not merely conceptual and malleable principles stated by the Universal Declaration of Human Rights (or the similar European Convention of Human Rights) but their universality is actually unanimously accepted as something that should apply to all human beings, then any new architecture which promises to fundamentally change how people exchange information should be analyzed and based on the same shared values from the ethical point of view. Any major technological change which potentially could affect these values at a global scale, thus affecting millions of people or even more, must be rigorously checked before implementation and also

must be made sure that its purpose is to help enforce the values mentioned earlier for future generations and not to undermine them.

It is this author's opinion - supported by the various interviews and appropriate documentation on the subject used to uncover what and how ethical principles should apply in this case - that CCN is indeed a promising and viable future Internet architecture. However, there are still many questions left unanswered on what its final standardized form will be, so that it will not overlap with ethical and moral issues related to basic human rights.

It is also this author's opinion and conclusion that if standardization in this case shall rest on the premises of principle locality in regards with how privacy and freedom of speech are approached ethically and not by allowing a "reductio ad abstractum" approach - where the principles themselves are enforced at the architecture level as opposed to questioning where they should or they should not be - than CCN could potentially become a mechanism for information censorship and manipulation.

References

Ahlgren, B., Dannetwitz, C., Imbrenda, C., Kutscher, D., Ohlman, B. (2011), "A survey of Information-Centric Networking", Schloss Dagstuhl - Leibniz Center for Informatics, Dagstuhl Seminar Proceedings, ISSN 1862-4405

Anti-terrorism, Crime and Security Act 2001 (c. 24), Part 11, UK Parliament

Arianfar, S., Nikander, P., Ott, J. (2010), "On content-centric router design and implications", ReARCH '10: Proceedings of the Re-Architecting the Internet Workshop, ISBN: 978-1-4503-0469-6.

BBC, 2010. Wikileaks diplomatic cables release 'attack on world', [Online] November 2010. Available at: www.bbc.co.uk/news/world-us-canada-11868838 [Accessed on 02 May 2011]

Carofiglio, G., Gehlen, V., Perino, D., (2011), "Experimental Evaluation of Storage Management in Content-Centric Networking". [Online] To appear in Next Generation Networking and Internet symposium, IEEE ICC 2011, Kyoto, Japan. Available at <http://www.ccnx.org/content/content-centric-networking-resources> [Accessed 20 March 2011]

Carrea, L. (2010), "Current CCN development" [Interview], Essex University with Filip Pitaru. 04 February 2010.

CCNA exploration companion guide, Pearson Education, London,
ISBN-10: 1-58713-204-4

Ciuche, D. (2011). "Content centric network architecture and potential security issues". [Interview] Bullgurad Internet Security in Bucharest with Filip Pitaru. 12 April 2011.

Daras, P., Semertzidis, T., Makris, L., Strintzis, M.G., (2010), "Similarity content search in content centric networks", MM'10: Proceedings of the international conference on Multimedia, ISBN: 978-1-60558-933-6.

Deti, A., Blefari-Melazzi, N. (2010), "Network-layer solutions for a content-centric Internet". [Online] 21st International Tyrrhenian Workshop on Digital Communications (ITWDC) - Italy, European FP7 Project, Available at www.ict-convergence.eu/documents [Accessed 28 April 2011]

Directive 2006/24/EC of the European Parliament and of the Council of 15 March 2006 on the retention of data

Dorogovtsev, S.N., Mendes, J.F.F. (2002), "Evolution of networks", Advances in Physics, Vol. 51, No. 4, 1079-1187, DOI: 10.1080/00018730110112519

EU proposal 52010PC0471 from 20 September 2010, Proposal for a

Decision of the European Parliament And Of the Council establishing the first radio spectrum policy programme

European Commission (2011), "Commission Communication: The open internet and net neutrality in Europe (19.04.2011)". [Online] Available from http://ec.europa.eu/information_society/policy/ecomm/doc/library/communications_reports/netneutrality/comm-19042011.pdf [Accessed 05 April 2011]

European Convention for the Protection of Human Rights and Fundamental Freedoms as amended by Protocols No.11 and No. 14 signed on 4 November 1950.

Ethics, privacy and security in content centric architecture group 2011, Minutes from 08 May 2011, Bucharest, Romania

Froomkin, M. (2000), "The Death of Privacy", Stanford Law Review 52(5):1461-1543.

Gilbert, H., Peyrin, T. (2009), "Super-Sbox Cryptanalysis: Improved Attacks for AES-like permutations", Cryptology ePrint Archive, Report 2009/531

Google vs. Luis Vutton. Case 236/08 [2010], European Court of

Justice [Online], Available at <http://curia.europa.eu/jurisp/cgi-bin/form.pl?lang=EN&Submit=rechercher&numaff=C-236/08>
[Accessed 05 April 2011]

Graziani, R., Jonson, A. (2007), „Routing protocols and concepts”, Cisco Press, ISBN-13: 978-1587132063

Huffman, S.M., Reifer, M.H., 2005, „Method for geolocating logical network addresses”, U.S. Pat. 6,947,978 assigned to the National Security Agency.

Internet World Stats (2011), “Internet usage and population statistics in Europe” [Online], Available from:
<http://www.internetworldstats.com/stats4.htm> [Accessed 02 May 2011]

Ipoque GmbH (2008), University of Innsbruck Case Study. [Online] Available from: <http://www.ipoque.com/userfiles/file/ipoque-case-study-innsbruck-english-web.pdf> [Accessed 03 February 2011]

Ipoque GmbH (2008), TVTEL Case Study. [Online] Available from: www.ipoque.com/userfiles/file/ipoque-case-study-tvtel-english-web.pdf [Accessed 03 February 2011]

Jacobson, V. et al. (2009B), “VoCCN: Voice over Content-Centric Networks”. [Online] ACM ReArch'09, Available at www.parc.com/content/attachments/voccn-voice-over-ccn-

[preprint.pdf](#) [Accessed: 10 February 2010]

Jacobson, V.; Smetters, D. K.; Thornton, J. D.; Plass, M. F.; Briggs, N.; Braynard, R., (2009A) "Networking named content". Proceedings of the 5th ACM International Conference on Emerging Networking Experiments and Technologies (CoNEXT 2009); Rome, Italy. NY: ACM; 2009; 1-12.

Jonas, H., (1979), „Toward a Philosophy of Technology“ [Online], The Hastings Center Report, Vol. 9, No. 1, (Feb., 1979), pp. 34-43, Available at www.jstor.org/stable/3561700 [Accessed 07 July 2010]

Kantar Group Market Research (2011), "Digital life statistics". [Online] Available from: www.discoverdigitallife.com [Accessed 02 March 2011]

Kazi, A.W.(2010), "Prefetching Bloom filters to control flooding in content-centric networks", CoNEXT'10 Student Workshop: Proceedings of the ACM CoNEXT Student Workshop, ISBN: 978-1-4503-0468-9.

Kutscher, D., Ahlgren, B., Karl, H., Ohlman, O., Oueslati, S., Solis, I. (2011), "Abstracts Collection - Information-Centric Networking", Schloss Dagstuhl - Leibniz Center for Informatics, Dagstuhl Seminar Proceedings, ISSN 1862-4405

Lauinger, Tobias (Sept. 2010), "Security & Scalability of Content-Centric Networking". [Online] Master's Thesis, Technische Universitat Darmstadt, Germany and Eurécom, Sophia-Antipolis, France. Available at <http://tuprints.ulb.tu-darmstadt.de/2275/1/ccn-thesis.pdf> [Accessed 04 February 2011]

Manolea, B. (2011), "Juridical implications in content centric network architectures" [Interview], The Association for Technology and Internet (APTII) from Bucharest with Filip Pitaru. 27 March 2011.

Mesch, G.S., Levanon, Y. (2003), „Community Networking and Locally-Based Social Ties in Two Suburban Localities“, City & Community Volume 2, Issue 4, pages 335–351, December 2003, DOI: 10.1046/j.1535-6841.2003.00059.x

Project CCNx™, 2010. Open source project exploring the next step in networking. [Online] Available at www.ccnx.org [Accessed 05 February 2010]

Reardon, D. (2006), "Doing your Undergraduate Project", Sage Publications, ISBN: 978-0-7619-4206-1

Roberts, L.G. (Jan. 2000), "Beyond Moore's law: Internet growth trends", Computer, Volume 33 Issue 1, IEEE Computer Society, DOI 10.1109/2.963131

Romanian Constitutional Court Decision no.1258 of Oct. 8 regarding data retention by ISP's, 2009, Official Gazette no. 798 of Nov. 23, 2009

Sabam vs. Tiscali. Case C-360/10 [2010], European Court of Justice [Online], Available at <http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=OJ:C:2010:288:0018:019:EN:PDF> [Accessed 05 April 2011]

Senate bill No. 761 of the California State Senate introduced on 18 February 2011 on regulating entities that collect, use, or store online data containing covered information from a consumer in the respective state

Shenker, Scott (Sept. 1995), "Fundamental design issues for the future Internet", IEEE Journal on Selected Areas in Communications, Volume 13, Issue 7, DOI 10.1109/49.414637

Smetters, D. K., Jacobson, V., (2009), "Securing network content". [Online] PARC technical report. Available at www.parc.com/content/attachments/securing-network-content-tr.pdf [Accessed 10 February 2010]

Smetters, D. K., Jacobson, V. (2009), "Securing network content". PARC technical report TR-2009-01.

Uichin, L., Rimac, I. And Hilt, V. (2010) "Greening the internet with content-centric networking". e-Energy '10 - Proceedings of the 1st International Conference on Energy-Efficient Computing and Networking, DOI: 10.1145/1791314.1791342

US National Defence Agency (2011). National Security Agency Military Construction, Defense-WideFY 2012 Budget Estimates. [Online] Available at http://comptroller.defense.gov/defbudget/fy2012/budget_justification/pdfs/07_Military_Construction/12-National_Intelligence_Agency.pdf [Accessed 03 May 2011]

Vica, C. (2011). "Applied ethics in content centric network architectures" [Interview], Research center in applied ethics at University of Bucharest with Filip Pitaru. 10 March 2011.

Wong, W., Nikander, P. (2011), "Towards Secure Information-centric Naming". [Online] National Physical Laboratory Conference, Securing and Trusting Internet Names, SATIN 2011, Available at <http://conferences.npl.co.uk/satin/papers/satin2011-Wong.pdf> [Accessed 20 April 2011]

Yin, S., 2011. „Sony Accounts Hacked? Here's What You Should Do Now, Soon, and Later". PC Magazine [Online]. Available at: www.pcmag.com/article2/0,2817,2384910,00.asp [Accessed 05 May 2011]

Yu, I., Song, B., Son, J., Baik, D.K. (2011), "Discovering credentials in the content centric network", 2011 International Conference on Information Networking (ICOIN), ISBN: 978-1-61284-661-3.