A Tussle Analysis for Information-Centric Networking Architectures

Alexandros Kostopoulos[1], Ioanna Papafili[1], Costas Kalogiros[1], Tapio Levä[2],
Nan Zhang[2], and Dirk Trossen[3]

[1] Athens University of Economics and Business, Department of Informatics, Athens, Greece
[2] Aalto University, Department of Communications and Networking, Espoo, Finland
[3] University of Cambridge, Computer Laboratory, Cambridge, UK
{alexkosto,iopapafi,ckalog}@aueb.gr,
{tapio.leva,nan.zhang}@aalto.fi, dirk.trossen@cl.cam.ac.uk

Abstract. Current Future Internet (FI) research brings out the trend of designing information-oriented networks, in contrast to the current host-centric Internet. Information-centric Networking (ICN) focuses on finding and transmitting information to end-users, instead of connecting end hosts that exchange data. The key concepts of ICN are expected to have significant impact on the FI, and to create new challenges for all associated stakeholders. In order to investigate the motives as well as the arising conflicts between the stakeholders, we apply a tussle analysis methodology in a content delivery scenario incorporating socio-economic principles. Our analysis highlights the interests of the various stakeholders and the issues that should be taken into account by designers when deploying new content delivery schemes under the ICN paradigm.

Keywords: information-centric networking, content delivery, future internet architecture, tussles, incentives, socio-economics, value network.

1 Introduction

Over the recent years, an increasing number of users gain access to the Internet via numerous devices equipped with multiple interfaces, capable of running different types of applications, and generating huge data traffic volumes, mostly for content. Traffic stemming out of these activities implies increased cost for the Internet Service Providers (ISPs) due to the congestion in their networks and the generated transit costs, as well as unsatisfactory Quality of Service (QoS) for some end-users.

This exponential growth of content traffic has been initially addressed by peer-to-peer applications, or Content Distribution Networks (CDNs). CDNs consist of distributed data centers where replicas of content are cached in order to improve users' access to the content (i.e., by increasing access bandwidth and redundancy, and reducing access latency). These CDNs practically formulate overlay networks [1] performing their own traffic optimization and making content routing decisions using incomplete information about customer's location and demand for content, as well as utilization of networks and available content sources. Similarly ISPs perform individual traffic optimization using proprietary, non-native and usually non-scalable solutions for

F. Álvarez et al. (Eds.): FIA 2012, LNCS 7281, pp. 6–17, 2012.

traffic monitoring and shaping (e.g. Deep Packet Inspection (DPI) boxes for peer-to-peer traffic) and have no incentive to reveal information about their network to CDNs. This *information asymmetry* often leads to a suboptimal system operation.

Information-centric Networking (ICN) postulates a fundamental paradigm shift away from a host-centric model towards an information-centric one. ICN focuses on information item discovery and transmission and not on the connection of end-points that exchange data. Thus, ICN has the potential to address efficiently the aforementioned information asymmetry problem by including traffic management, content replication and name resolution as inherent capabilities of the network.

What remains the same is that the Internet is a platform composed of multiple technologies and an environment where multiple stakeholders interact; thus, the Internet is interesting from both the technological and the socio-economic viewpoint. Socio-economic analysis comprises a necessary tool for understanding system requirements and designing a flexible and successful FI architecture.

A first attempt to investigate socio-economic aspects of FI in a systematic manner was performed by Clark *et al.* [2]. They introduced the '*Design for Tussle'* principle, where the term '*tussle'* is described as an *'ongoing contention among parties with conflicting interests'*. It is obvious that the need for designing a *tussle*-aware FI has emerged to enhance deployment, stability and interoperability of new solutions. Although there are plenty of counter-examples of adopted protocols/architectures that do not follow the Design for Tussle principle, tussle-aware protocols and architectures are expected to have better chances for adoption/success in the long-term [3].

The need for *understanding the socio-economic environment, the control exerted on the design, as well as the tussles arising therein* has been also highlighted in [4]. The purpose of this work is to explore and analyze the tussles that may arise in ICN, as well as to consider the roles of different stakeholders; below, we present a tussle analysis methodology which extends the methodology originally developed within the SESERV project [5], and apply it in the content delivery scenario. We focus on the *tussle spaces* of name resolution, content delivery and caching.

This paper is organized as follows: In Section 2, we present our methodology for identifying tussles among different stakeholders. Then, Section 3 provides an overview of representative information-centric networking architectures developed in the PURSUIT [6] and SAIL [7] research projects. In Section 4, we focus on a use case for content delivery; we identify the involved stakeholders and major functionalities and roles that they can take, and then investigate the potential tussles among the stakeholders. Finally, in Section 5, we conclude our remarks.

2 A Methodology for Tussle Analysis

This section provides a generic guide for better understanding the impact of a technology on the stakeholders' strategies, as well as on how other technologies might be used and deployed. Below, we extend the methodology presented in [8] and combine it with the Value Network Configuration (VNC) method introduced by Casey *et al.* [9]. The tussle analysis methodology consists of the following steps:

1. Identify all primary stakeholder roles and their characteristics for the functionality under investigation.

2. Identify tussles among identified stakeholders.
3. For each tussle:
   (a) Translate knowledge into models by assessing the mid-term and long-term impact to each stakeholder;
   (b) Identify potential ways for stakeholders to circumvent negative impacts, and the resulting spill-overs.
4. For each circumventing technique, apply steps 1-4 again.

The involved stakeholders usually express their interests by making choices that will affect the technology by deciding which technologies will be introduced, how these will be dimensioned, configured, and finally, used. All these collective decisions will eventually determine how technology components will operate and produce outputs that are valuable for these stakeholders. Technology outputs are assessed by each stakeholder individually and can affect real-world interactions (e.g. payments, price competition, price regulation and collaboration) or trigger new technology decisions. Such interactions allow the Internet to evolve and act as a living organism (Fig. 1).

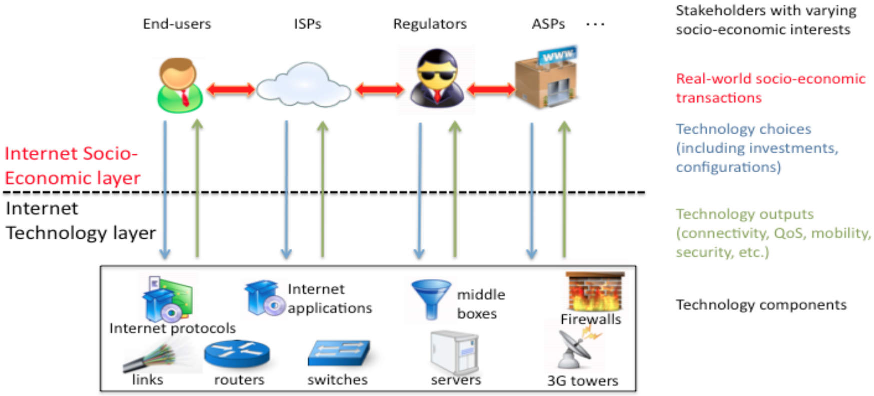

Fig. 1. The Socio-Economic layer and Technology layer of the Internet ecosystem

Several techniques or methods can be used to perform each of the aforementioned steps. In this paper, we show how the VNC method [9] can be incorporated in the tussle analysis. What makes the VNC method a particularly useful tool for tussle analysis is the separation of the stakeholders (or *actors* as Casey *et al.* call them) from the functional roles the actors can take, thus allowing us to analyze multiple role combinations instead of limiting to a single value network.

Identifying functional roles - defined in [9] as a set of activities and technical components, the responsibility of which is not divided between separate actors in a particular scenario- is central to the VNC method. Because roles hold economic and strategic value, the actors fight for their control. The tussles emerge when there is a conflict of interest between the actor controlling the role and the other actors affected by it. Depending on which actor controls a role, the tussle outcomes and the circumventing techniques vary, which further motivates the usage of the VNC method.

The VNC method emphasizes technologies' role in defining the possible value networks by identifying also the technical components and technical interfaces between them. By doing this, the method improves our understanding of the relationship between the technical architecture (a set of technical components linked to each other with technical interfaces, such as protocols) and the value network configuration (role division and related business interfaces among actors). This is important in analyzing whether the technology is designed for tussle [2], i.e., if the technical design allows variation in value networks. Fig. 2 presents the notation presented in [9] that can be used to visualize the roles and VNC.

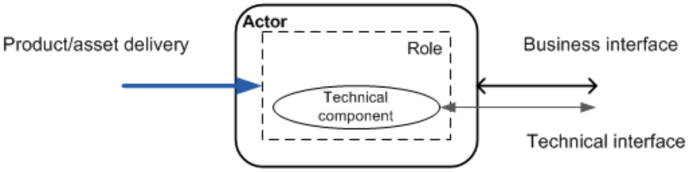

Fig. 2. Notation of VNC methodology

After identifying the involved stakeholders as well as the tussles among them, the next step is to translate knowledge into models and provide quantitative analysis. In [10] a toolkit is suggested that uses mind-mapping techniques and system dynamics to model the tussles. System Dynamics (SD) [11] is a useful tool to evaluate dynamic interactions between multiple stakeholders, by simulating the possible outcomes (e.g., how technology diffuses) when multiple stakeholders interact. The main focus is on the assessment of outcomes and their evolution over time, since possible reactions can be modeled. After having captured the causality models, relevant socio-economic scenarios may be formulated to investigate the potential consequences in the Internet market. We do not conduct SD analysis in this paper due to space constraints.

3 Overview of ICN Architectures

Diverse research projects, such as PURSUIT [6], SAIL [7] and NDN [12] are emphasizing the need to move towards an ICN architecture. In this section we briefly present an architecture overview of ICN in order to provide the necessary background. We focus on the Publish/Subscribe (pub/sub) model adopted by PURSUIT and the Network of Information (NetInf) introduced by SAIL.

3.1 Publish/Subscribe

In the PURSUIT pub/sub paradigm, information is organized in *scopes*. A scope is a way of grouping related information items together. A dedicated matching process ensures that data exchange occurs only when a match in information item (e.g., a video file) and scope (e.g., a YouTube channel) has been made. Each packet contains the necessary meta-data for travelling within the network. Fig. 3 presents a high level picture of the main architectural components of the pub/sub architecture.

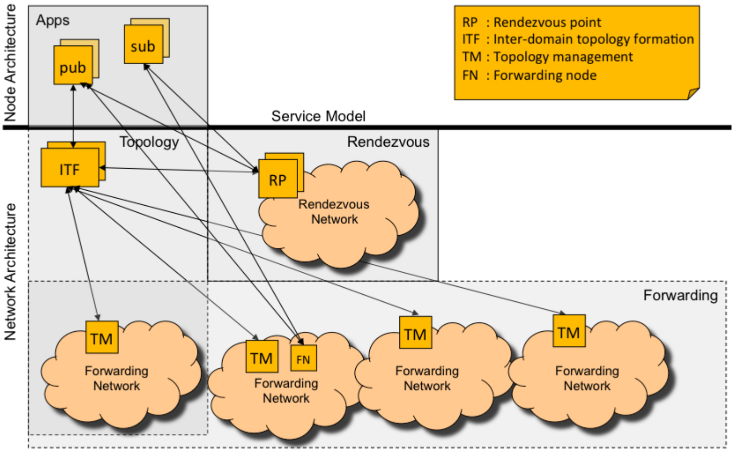

Fig. 3. A Publish/Subscribe architecture for ICN [13]

At the *application* level, the pub/sub components implement applications based on basic ICN services, enabling publications and subscriptions towards information items within particular scopes.

At the *network* level, the architecture itself consists of three main functions: *rendezvous*, *topology* and *forwarding*. The rendezvous function implements the matching between publishers and subscribers of information based on several criteria. Moreover, the rendezvous service provides additional functionalities to implement policies associated with the matching, such as access control. When a publication is matched with one or more subscriptions, an inter-domain forwarding graph is created in negotiation with the inter-domain topology formation (ITF) function. After constructing inter-domain paths between the forwarding networks to which publisher(s) and subscriber(s) are attached, intra-domain paths need to be constructed. This is done in collaboration with the AS-internal topology management (TM) function, which instructs its local forwarding nodes (FN) to establish paths to local publishers / subscribers or to serve as transfer links between ASes.

3.2 Network of Information

The SAIL Network of Information (NetInf) aims at three architectural objectives: i) unique naming regardless of the Named Data Object's (NDO's) location and without a hierarchical naming structure; ii) receiver-oriented NDO delivery; and iii) a multi-technology and multi-domain approach, where any underlying technology and network can be leveraged [14]. The NetInf network consists of Name Resolution System (NRS) nodes and NetInf router (NR) nodes, which are illustrated in Fig. 4.

NetInf supports both name-based routing and name resolution. Name resolution is enabling scalable and global communication: NDOs are published into the network and registered by the NRS. Specifically, the NRS is used to register the network locators of NDO copies in the underlying network, which can potentially provide packet-level routing and forwarding functionalities. The NDO request can be resolved by the NRS into a set of network locators, which are used to retrieve a copy of the

NDO from the optimum source based on a pre-defined criterion. At least one global NRS must exist in the NetInf network, but also intra-domain NRS' are possible.

The NetInf router node accepts NetInf names as input and decides how to route the request so that eventually a NDO is returned to the previous-hop NetInf node. This routing decision could be either towards a NRS or directly towards the NDO source, the latter of which represents the name-based routing scenario. In addition, NetInf cache servers for content replication can be placed both in the NR nodes and the NRS nodes.

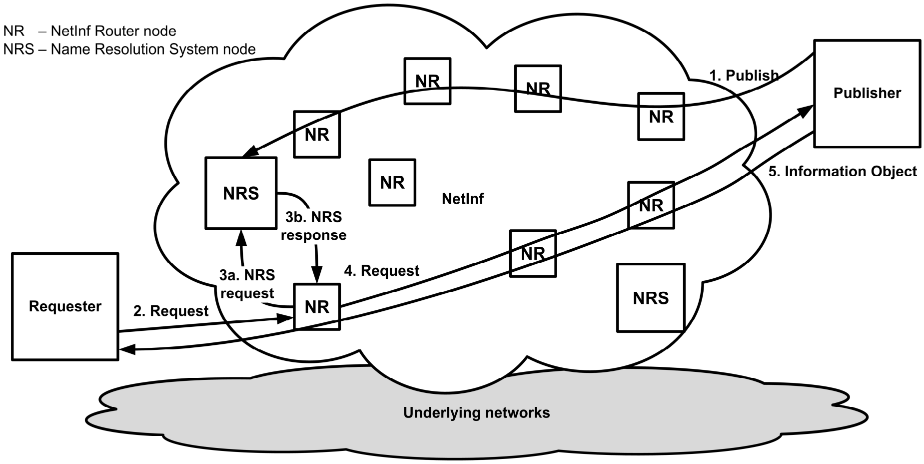

Fig. 4. NetInf high level architecture

Fig. 4 also shows the high level content retrieval process in NetInf. First, (1) a NDO owner publishes the NDO into the network by adding it to the NRS registry. When a (2) request for a NDO occurs, the NetInf router can either (3a) forward the request to a NRS for (3b) the set of locators or it can (4) directly forward the request to the NDO source, depending on whether the NetInf router knows where the NDO is. Finally, (5) the NDO is returned to the requester via the same route as the request and the NDO can be cached on every node that it passes.

4 Tussles in Information-Centric Networking

In this section, we focus on the content delivery use-case in a *generic* ICN architecture and apply our combined tussle analysis and VNC methodologies to it. We first look into the intra-domain scenario and then build incrementally on the inter-domain scenario. As the first step of our methodology, we identify here major functionalities, group them into roles and list the stakeholders that can take up these roles. Then, in the second step, we perform tussle analysis on a per functionality view.

4.1 The Content Delivery Use-Case

As illustrated in Fig. 5, we consider two Access Network Providers (ANPs) that employ ICN to offer content delivery services to their customers. The two ANPs are connected through transit links to an Inter-Connectivity Provider (ICP). Both ANPs employing ICN have deployed their own networks of Caches. Within the ANPs premises, local NRSs are also provided, which are connected to a global NRS service. The NRSs could be controlled by either the respective network infrastructure provider (ANP or interconnectivity provider) itself, or by a third-party. Potential subscribers of an information item exist in both ANPs; however, only a single publisher (P_1) of that specific content exists initially, in ANP_1.

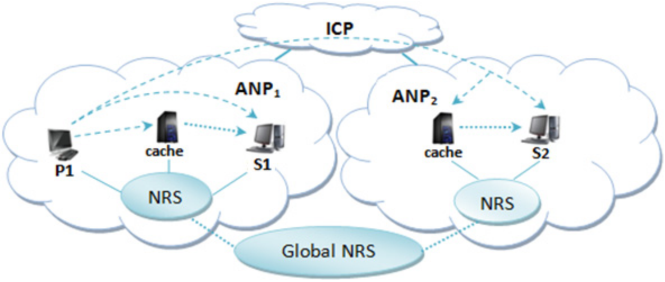

Fig. 5. Content delivery in ICN architecture

Intra-domain Scenario. We assume that P_1 in ANP_1 publishes an information item to his local NRS, and the local NRS advertises the publication to the global NRS. Then, S_1 in ANP_1 sends a subscription for an information item to the local NRS of its ANP. The local NRS identifies that the requested information item is published within the ANP and matches P_1 with S_1. If more subscriptions for the same information item occur, the ANP may also decide to cache the content to another location in order to achieve load balancing and to provide higher QoS to its customers (subscribers).

Inter-domain Scenario. Let us now assume that S_2 in ANP_2 also subscribes to his local NRS for the same information item. Since, the information item is not published within ANP_1, the local NRS informs the global NRS about this subscription. The global NRS, who is aware of P_1, matches P_1 with S_2. ANP_2 may cache the information item in his caching elements, in order to serve potential new subscribers.

4.2 Functionalities, Roles, Stakeholders

Based on the aforementioned use-case, we identify the key functionalities and map them to five key roles (Table 1). There are multiple stakeholders in position to control these roles, which would lead to different outcomes. Here, we focus on the role allocation visualized in Fig. 6, since it is a representative case to take place in ICN. In our setup, the *content access management* (i.e. AAA) role can be taken by either the Content Provider (CP) or the ANP, the *name resolution* is taken by either the ANP or a third-party provider (i.e. a Rendezvous Network (RENE) provider in [6]), whereas the other four roles are assigned to the ANP. The chosen role allocation differs from the typical situation in the market today where other stakeholders, such as CDN providers or CPs, control the name resolution, caches and content network.

Table 1. Key roles and functionalities in ICN content delivery

Role	***Functionalities***
Name Resolution	Content directory control, names to locations resolution, rendezvous, matching, applying policies
Content access management	AAA (Authentication, Authorization, Accounting)
Cache management	Cache servers control, content selection for being cached, cache updating
Cache location ownership	Cache locations control
Content network management	Content network resources selection, path selection, QoS

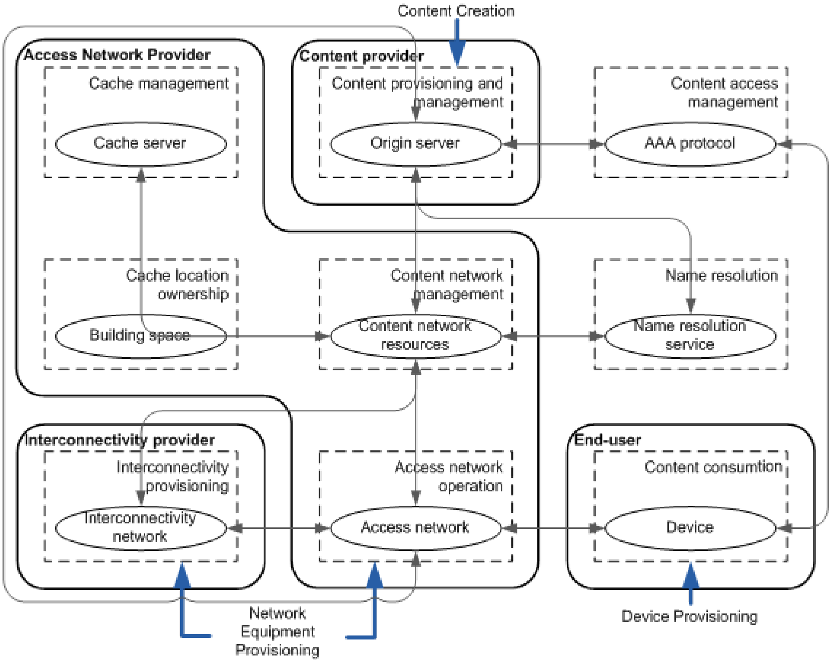

Fig. 6. Generic Value Network Configuration (VNC) for content delivery in ICN

The major stakeholders that can take up the aforementioned roles in our scenario are presented in Table 2. We also use parentheses to include the additional roles that could be potentially taken up by stakeholders in other scenarios. Additionally, we include the CDN providers, as well as the regulators that exist in current Internet, although their interests and actions are not subject of this analysis.

4.3 Tussle Analysis

In this section we identify tussles related to key roles listed in Table 1. Each tussle is described with references both to the use case (Fig. 5) and the VNC (Fig. 6).

Table 2. Stakeholder - basic role mapping

Stakeholder	***Basic role***
End-user	Content consumption, (content creation)
Content Provider (CP)	Content creation, (content access management)
Internet Service Provider (ISP)	
- Access Network Provider (ANP)	Access network operation, cache management, cache location ownership, content network management, (name resolution , content access management)
- Inter-Connectivity Provider (ICP)	Interconnectivity provisioning to ANPs, (name resolution)
NRS provider	Name resolution
Content Distribution Network Provider (CDN), e.g. Akamai	*Cache management, cache location ownership, content network management, name resolution*
Regulator	*Competition regulation*

Tussles related to name resolution

Spam Requests Tussle: The local NRS may decide to replicate the requested information to his own cache like the rendezvous in the pub/sub model. In this case, the local NRS (or RENE) adds a subscription in his message towards the publisher asking the information to be forwarded also to the ANP's cache. Thus, an NRS could issue a request for another stakeholder (e.g. the end-user) for an information item that the latter is not interested in (spam). This combined service contradicts the *functionality separation* as dictated in [2], since the rendezvous also performs content management besides its main function, i.e., name resolution.

Net Neutrality Tussle: The global NRS is potentially in a position to favor specific CPs by promoting their content over the content of other CPs, or by filtering the information items provided by the latter ones. Additionally, if the local NRS is provided by the ANP (similar to today's ISPs' DNS service bundled with access provisioning), there is an incentive for the NRS to forward the subscription to the local publisher. If the content is *not* locally published, then the ANP-owned local NRS (NRS_2) may refuse to further handle the request to avoid fetching the information object from a remote publisher or the cache of a competing CDN to avoid increasing ANP_2's interconnection costs. The latter case is also known as a "walled garden". Ideally this situation is avoided by having architectures that allow competition in the resolution service; otherwise a regulator would have to ensure that end-users are allowed to send their subscriptions to the NRS of their choice.

Conflicting Optimization Criteria Tussle: When multiple sources can serve a request, a tussle occurs due to actors' different preferences for the one to be used (e.g., cost concerns, performance attributes, regulatory constraints, or other local policies). For example, localization of traffic due to caching and content replication affects the volume exchanged between ANPs, as well as ANPs and ICPs. If the local NRS forwards the content requests to local caches, both the interconnection costs of ANPs and revenues of ICP decrease. This is naturally positive to ANPs but negative to ICPs.

Similarly, an ICP-owned global NRS may forward a subscription originated from a local NRS to publishers that are located behind a transit link, even if the information item was also available through a peering link (a different scenario than the one in Fig 5). The same situation could appear if the local NRS is provided by a third-party, similar to, e.g., Google's DNS, which may have different incentives. Such conflicting optimization criteria might imply a straightforward increase of interconnection cost for the ANP, and possibly degraded end-users' Quality of Experience (QoE).

As it is obvious, the actor who controls the name resolution is able to restrict or even determine the available options to others. However, such an actor (like an ANP when the end-user has used a different NRS provider) may still be able to use a different source than the proposed one. For example in [6], after the final matching of a publisher and a subscriber by the Rendezvous Network, the Topology Manager may create a path between the subscriber and a different publisher (i.e., an ANP's own cache server)[1]. This could be the case when the end-user or the NRS provider cannot verify which publisher has been actually used.

Furthermore, other stakeholders could enter the name resolution market. In an extreme case, even a CP may react by providing also his own NRS. For example, YouTube could serve its information space by redirecting end-users to servers according to its own criteria). Such an NRS may also be provided as a premium service to other CPs. However, in both cases, client configuration by the end-users is required.

Finally, traditional CDN providers (like Akamai) could also react by announcing all the content items (publishers and caches) they are aware of to multiple NRS providers, or even deploy their own name resolution servers.

Nevertheless, the name resolution role is central to ICN and of high interests to the most stakeholders in this setup.

Tussles related to content access management

Access Control Tussle: If the ICN architecture does not clearly specify how to limit access to certain end-users, the ANP may serve the subscriptions from its local cache without consulting CP's AAA system. This would destroy CP's business, especially if it is based on transactional payments from end-users, but also if he sells advertising or information about content usage. A proposed solution is presented in [10], where the RENE could act as an accountability broker between the end-users and CPs.

Content Usage Statistics Tussle: When the content is provided from local caches controlled by multiple stakeholders, the CP may lose visibility on how its content is used. This information has value, because payments from advertisers to CP and from CP to content makers are often based on the popularity of content.

Privacy Tussle: Finally, a control tussle may rise between the stakeholder managing content access and the end-users, since the former can use personal and transactional data for purposes not approved by the end-user to make a profit, e.g. to sell data to marketing companies.

---

[1] Here, we assume that the Topology Manager is aware of the information item ID.

Tussle related to cache management

Content Freshness Tussle: The content cached in the ANP's caches may be outdated, because the ANP may be reluctant to update the content in order to reduce his interconnection (i.e., transit) costs. Then, the end-user's quality of experience degrades, since he does not receive the most recent information.

Tussles related to cache location ownership

Cache Placement for Revisiting Interconnection Agreements Tussle: Tussles here mostly involve ISPs since existing interconnection agreements may not be justifiable if a new cache was added. Hence, ISPs may try to affect peering ratios in advantageous ways (e.g. create such an imbalance that violates their peering agreement). For example, an ANP deploying his own cache content network and having a peering arrangement with another ANP (which does not own a content network) may break this agreement in hopes of providing transit service to the latter one. Similarly, an ICP who sees its revenues being reduced may decide to adjust transit prices or enter the content delivery market by providing global NRS services.

Tussles related to content network management

Network Information Tussle: An ANP may provide inaccurate information (or no information at all) about its network topology, dimensioning, current utilization, etc., fearing that this sensitive information could be revealed to its competitors. However, this may have a negative impact on the effectiveness of selecting publishers and consequently paths between publishers and end-users that meet the QoE constraints posed by the latter. For example, in case there are two publishers for a particular request, one of them may seem more appropriate (although it may not be), if its own ISP is untruthful by providing biased network information (e.g. lower delay in a path).

5 Discussion

ICN brings new challenges in the Internet market, since name resolution services may be offered by different stakeholders in order to meet their own optimizing criteria; either by the ANP, or by a third-party (such as a search engine or a significant CP). Such major stakeholders of today's Internet are highly expected to extend their activities to offer NRS' in ICN.

Additionally, there is a crystal clear incentive for an ANP to deploy ICN, in order to enter the content delivery market. Due to the information-oriented nature of the network, an ANP could deploy his own caches, which implies that the ANP will gain more control of the content delivery. Therefore, under suitable business agreements, this will imply increase of his revenue, while simultaneously reducing his operational costs due to more efficient content routing and reduction of the inter-domain traffic.

Moreover, CPs and end-users will also be affected; i.e. CPs will be able to provide their content through more communication channels to their customers, while end-users will enjoy increased Quality-of-Experience (QoE).

On the other hand, the emergence of ANP-owned CDNs will cause traditional CDNs to lose revenues and control over the content delivery market. Thus, legacy CDNs will probably react in order to maintain their large market share, or at least not exit the market. CDNs may deploy their own backbone networks to interconnect their own caches, but still they will probably not in position to deploy access networks to reach the end-users; this is ANPs' last frontier. Nevertheless, no matter how legacy

CDNs will react, such local CDNs owned by ANPs will (and already) be deployed (e.g. At&T's CDN). The evolution of this competition and the way that the system will be lead to an equilibrium is the subject of future investigation and analysis.

Our contribution in this paper resides in the identification and analysis of tussles in a generic ICN architecture, which should be considered by designers and engineers that aim at deploying new content delivery schemes for the FI.

Acknowledgement. The authors would like to thank G. Xylomenos, G. D. Stamoulis, G. Parisis, D. Kutscher C. Tsilopoulos and X. Vasilakos. The research of A. Kostopoulos and I. Papafili 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. C. Kalogiros is supported by the EU-FP7 *SESERV* project. The research of T. Levä and N. Zhang is supported by the EU-FP7 *SAIL* project. The research of D. Trossen is supported by the EU-FP7 *PURSUIT* project.

References

1. Clark, D.D., Lehr, W., Bauer, S., Faratin, P., Sami, R., Wroclawski, J.: Overlay Networks and the Future of the Internet. Communications and Strategies 63, 109–129 (2006)
2. Clark, D.D., Wroclawski, J., Sollins, K.R., Braden, R.: Tussle in Cyberspace: Defining Tomorrow's Internet. IEEE/ ACM Trans. Networking 13(3), 462–475 (2005)
3. Kalogiros, C., Kostopoulos, A., Ford, A.: On Designing for Tussle: Future Internet in Retrospect. In: Oliver, M., Sallent, S. (eds.) EUNICE 2009. LNCS, vol. 5733, pp. 98–107. Springer, Heidelberg (2009)
4. Brown, I., Clark, D., Trossen, D.: Should Specific Values Be Embedde. The Internet Architecture? In: ACM Context ReArch 2010 (December 2010)
5. EU FP7 SESERV project, `http://www.seserv.org/`
6. EU FP7 PURSUIT project, `http://www.fp7-pursuit.eu/`
7. EU FP7 SAIL project, `http://www.sail-project.eu/`
8. Kalogiros, C., Courcoubetis, C., Stamoulis, G.D., Boniface, M., Meyer, E.T., Waldburger, M., Field, D., Stiller, B.: An Approach to Investigating Socio-economic Tussles Arising from Building the Future Internet. In: Domingue, J., et al. (eds.) FIA 2011. LNCS, vol. 6656, pp. 145–159. Springer, Heidelberg (2011)
9. Casey, T., Smura, T., Sorri, A.: Value Network Configurations in wireless local area access. In: 9th Conference of Telecommunication, Media and Internet (2010)
10. Trossen, D., Kostopoulos, A.: Exploring the Tussle Space for Information-Centric Networking. In: 39th TPRC, Arlington, VA (September 2011)
11. Sterman, J.: Business dynamics: Systems thinking and modeling for a complex world. McGraw-Hill, Irwin (2000)
12. Named Data Networking (NDN) project, `http://www.named-data.net/`
13. Trossen, D., Sarela, M., Sollins, K.: Arguments for an Information-Centric Internetworking Architecture. ACM Computer Communication Review (April 2010)
14. Pöyhönen, P., Stranberg, O. (eds.): (D-B.1) The Network of Information: Architecture and Applications. SAIL project deliverable (2011)