

Scenarios and Research Issues for a Network of Information

Christian Dannewitz
University of Paderborn
Pohlweg 47-49
Paderborn, Germany
cdannewitz@upb.de

Kostas Pentikousis
VTT Technical Research
Centre of Finland
Kaitoväylä 1
FI-90571 Oulu, Finland
kostas.pentikousis@vtt.fi

René Rembarz
Ericsson Research
Ericsson Allee 1
Herzogenrath, Germany
rene.rembarz@ericsson.com

Éric Renault
TELECOM SudParis
Samovar UMR INT-CNRS
5157
91011 Évry, France
eric.renault@it-
sudparis.eu

Ove Strandberg
Nokia Siemens Networks
Linnoitustie 6
Espoo, Finland
ove.strandberg@nsn.com

Javier Ubillos
Swedish Institute of Computer
Science
Box 1263
SE-164 29 Kista, Sweden
jav@sics.se

ABSTRACT

This paper describes ideas and items of work within the framework of the EU-funded 4WARD project. We present scenarios where the current host-centric approach to information storage and retrieval is ill-suited for and explain how a new networking paradigm emerges, by adopting the information-centric network architecture approach, which we call Network of Information (NetInf). NetInf capitalizes on a proposed identifier/locator split and allows users to create, distribute, and retrieve information using a common infrastructure without tying data to particular hosts. NetInf introduces the concepts of information and data objects. Data objects correspond to the particular bits and bytes of a digital object, such as text file, a specific encoding of a song or a video. Information objects can be used to identify other objects irrespective of their particular digital representation. After discussing the benefits of such an indirection, we consider the impact of NetInf with respect to naming and governance in the Future Internet. Finally, we provide an outlook on the research scope of NetInf along with items for future work.

Categories and Subject Descriptors

C.2.1 [Computer-Communications Networks]: Network Architecture and Design; E.2 [Data Storage Representations]: Object representation

General Terms

Design

Keywords

Network of Information, content-centric, information-centric, data-centric

1. INTRODUCTION

Every part is disposed to unite with the whole,
that it may thereby escape from its own incom-
pleteness

—Leonardo da Vinci

Leonardo Da Vinci's notebooks are renowned. Beyond their artistic merit, according to Tufte [11], they are in fact excellent pieces of evidence-based work, were notes, measurements, dimensions, and figures merge together naturally, making up a consistent whole.

Tufte also writes that producing a piece of evidence of similar quality as Leonardo's notebooks is astonishingly difficult today, in particular if one uses the widely available state of the art computer programs. One has to use a word-processing application for writing text, a spreadsheet application for collecting and analyzing data, and a graphics application to produce the figures. Then all these "building" blocks will have to be put together in a typesetting/layout program. To seasoned computer users, this seems quite natural today. Yet, as Tufte reminds us, the original GUI designed by Xerox 30 years ago did not make a distinction between different applications. There were only the "desktop" and "documents". The latter could deal with all kinds of information (text, figures, tables, and so on) just like modern office automation suites can, but back then users did not perceive each document as an assortment of different media content. Instead, each document was seen as a coherent whole, a distinct information object.

In a similar sense, while the WWW started off as a collection of interconnected documents, today's Internet is dominated by a huge collection of incoherent data objects, spread among P2P networks, Wikis, communities, video sites, and other web pages. Many data objects are semantically related, i.e., they refer to the same topic, but this relationship is not represented in any way in the Internet. Likewise, the relation between different encodings (e.g. MP3, WMA) of the same information elements as well as exact copies of the same file are not represented in today's Internet.

The inability of today's Internet to represent the relationship between those closely related elements has at least two main disadvantages. First, information retrieval is unnecessarily complicated. While seasoned users are accustomed to using search engines, Wikis, communities, P2P networks, and so on, to locate information on the Internet, this represents a significant hurdle for less experienced users. Second, information dissemination is inefficient in today's Internet. For example, existing copies of data objects can not be used efficiently to support the data dissemination process as the relationship between copies is often unknown. Existing overlay solutions like Bittorrent try to reduce those disadvantages that stem from the fact that the Internet was originally designed for point-to-point communication and not data dissemination. Today however, the vast majority of Internet traffic is actually dissemination of information, not point-to-point communication or information exchange, as pointed out by Van Jacobson [2, 3]. In other words, the Internet is used as it was designed to be used only as an exception to the rule. Moreover, today's Internet is not very well-suited for reliable and trustworthy information dissemination. What we are after is a new network architecture that centers on information. An architecture that connects information producers and distributors with information consumers, not simply nodes with links and processes with their remote peers. We want to define an architecture fitting to the Information Age.

In this paper we attempt to motivate the work undertaken within the EU-funded project 4WARD (see www.4ward-project.eu) to define a new paradigm for computer networking. As we explain in the remainder, we have adopted an information-centric approach as we proceed in defining an alternative architecture for the future Internet. The work in this line of research has started only recently and the task is imposing. As such, in this position paper we can only present preliminary ideas and propose promising solutions. As one would expect, several issues will need to be addressed, and the proposed solutions will have to pass the tests of rigorous performance evaluation studies. But we are not there yet. At this stage of development we are looking forward to comments from the research community and our hope is that this paper can serve as a starting point for a greater discussion about what we call the Network of Information (NetInf).

This paper is organized as follows. Section 2 motivates the need for NetInf by introducing three scenarios. The first scenario, content distribution, reviews the conglomeration of methods to distribute information in the contemporary Internet and presents the concept of object hierarchies in NetInf. The second scenario presents the Augmented Internet, a concept that integrates the Internet with our real world, and the third scenario revisits mobile personal communications from a NetInf perspective. Section 3 describes the main

NetInf concepts, namely the identifier/locator split, the introduction of information and data objects, and reveals the first ideas about the NetInf API. Section 4 defines the scope of anticipated work in this research area and outlines items of future work. Finally, Section 5 summarizes this position paper.

2. SCENARIOS

The following sections will describe three scenarios with a strong information-centric background that will serve to illustrate research issues of the future Network of Information. We have chosen three very different and independent scenarios that pose many diverse requirements on the Network of Information to be able to evaluate and prove the benefits of our information-centric approach from many different angles. At the same time, those three scenarios represent use cases and aspects that we think will be very significant in the future Internet: content distribution, real-world/Internet integration, and mobility.

2.1 Content Distribution

Contemporary content distribution has effectively become an assortment of individual, often tailored and specialized solutions. Depending on the field of application, these solutions range from carefully engineered and managed Content Distribution Networks (CDNs) [7] owned by individual corporations all the way to more ad-hoc and unmanaged peer-to-peer networks such as BitTorrent-based file exchange overlays. In addition, a number of interesting technologies, such as content personalization and adaptation have emerged. Although on a first look these technologies seem quite different in the way they are approaching the problem, effectively they are in fact all dealing with getting content from information producers to information consumers. It is, therefore, natural to investigate how we can employ Networks of Information in the field of content distribution.

We envisage a solution where NetInf provides a uniform mechanism for accessing different types of content, seamlessly incorporating different types of content and distribution methods, including novel schemes that either blend existing ones to create a new form of distribution mechanisms or define completely new schemes. A key enabler to this approach is the overarching notion of information elements which is a NetInf architecture cornerstone and further generalizes previous proposals on content-centric networking [3]. In this architecture, the particular representation of information (say a song encoded in MPEG-1 Audio Layer 3), or how it is retrieved (via HTTP, mobile peer-to-peer, or BitTorrent) are subordinate to the information object itself.

In order to enable information objects at different semantic levels, we foresee also that a hierarchy of information objects can be modeled. A user might only be interested in listening to Beethoven's 9th symphony, without knowing about the different orchestras which have performed it. This can easily be modeled by an abstract representation of Beethoven's 9th which then refers to several recordings of concerts where it was performed. Fig. 1 provides three other examples of how we envision object hierarchies in the context of content distribution. At the top of the figure, a fairly straight-forward case is shown where an entry page of a web site is composed of a set of sub-components like text blocks and figures. The same concept can also easily be applied to a completely different information format, such

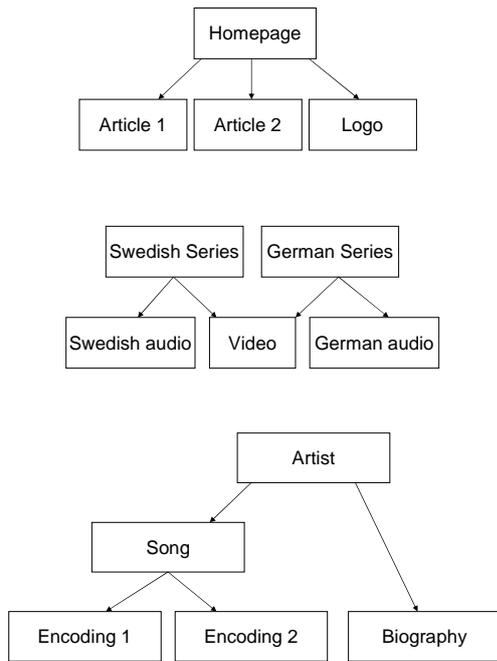


Figure 1: Three examples of object hierarchies in information objects

as a TV series (middle of Fig. 1), where information objects might correspond to the localized versions, sharing a common video and then linking to different audio files or subtitles. Finally, a wide range of elements can be understood as information objects, such as both a certain song and a biography being linked to an artist. The scenario on real-world objects presented in the next subsection takes this idea even further. As already explained, the more abstract notion of a song might then again link to specific recordings and their (digital) encodings of the song.

With this universal notion of content, we then have an overarching method at hand that can accommodate all different kinds of distribution schemes and network scenarios. From a user perspective, this implies that all that is needed is the identifier of the information object to be retrieved. Given this identifier, the NetInf infrastructure will then decide what the optimal source or sources are, and deliver the content to the user. “Source” in this context can be as simple as the original content owner’s web server, for example, but it may as well include retrieval from dedicated caches or peers that have already downloaded the content (or parts of the content). Possible distribution schemes can then range from simple point-to-point transfers all the way to BitTorrent-like methods. The algorithms used to determine the optimal delivery method can vary from network to network and therefore also apply to scenarios like a managed operator network.

We are also planning to explore the benefits of different types of content adaptation within this new network architecture. Transcoding, for example, has already been used in the contemporary Internet where properties like the resolution or bit rate of an information element are changed on the

fly. Scalable media, such as H.264/SVC, will gain in popularity in the foreseeable future, and are also included in our research agenda. Content personalization is also increasing in importance, along with blending of different content. In all cases, the actual treatment of the media happens outside the NetInf realm, but both the content elements that are fed into this process and the outputs of the process have to be handled by and integrated into the NetInf machinery. Last but not least, with the intrinsic cryptographic properties of information elements, also content protection is possible. Protection in this context not only refers to classical digital rights management (e.g. pay-TV and video on demand), but also to other forms of access control, such as users controlling who is allowed to see their holiday photos, possibly also adding an expiry date to the file.

2.2 Augmented Internet

As mobile Internet-enabled devices become more ubiquitous, users will want to use them to access more and more information and services in the real world surrounding them, for example about everyday objects, people they meet, or places they visit. When accessing information on the move, it is essential that the accessing information does not take the user’s attention away from the real-world activities. Unfortunately, mobile Internet access as we experience it today requires a lot of attention by the user and is therefore not suitable for many use case scenarios where users want to or have to focus on real-world activities. To support such scenarios, a smooth integration with the real world that enables service access without interrupting the user’s real-world work flow would have to be supported by Internet applications. But such applications are very hard to build on a large scale as the current Internet architecture does not provide a notion for real-world integration. Hence, providing conceptual network support for real-world-integrated Internet services will be a key challenge for the future Internet architecture. We call the concept of network support for real-world-integrated applications the *Augmented Internet* paradigm.

When users attempt to access information related to real-world entities, they are typically interested in the information itself and not in the nodes that deliver the information. For example, a tourist near the Eiffel Tower cares about opening hours, ticket cost, the history of the monument, and so on. If this information located in a server in Paris or elsewhere is irrelevant to the users. URIs such as www.tour-eiffel.fr provide an abstraction layer to the location of information, but nevertheless tie it to specific network nodes. We claim that real-world-integrated applications are inherently information-centric and that NetInf is much better suited to provide an architecturally sound infrastructure to enable and support real-world-integrated applications directly by the semantics of the network. Such applications pose two main requirements on an underlying infrastructure. First, the Augmented Internet needs a notion of virtual representations for physical entities that can cumulate and provide access to physical-entity-related services. Second, the Augmented Internet has to build and maintain a binding between the physical entity and its virtual representation on the future Internet. Our goal is to ensure that both requirements can be inherently met by NetInf.

The Augmented Internet paradigm combines multiple families of use cases that will be described in more detail now.

One main family combines use cases that enable the interaction with physical entities via entity-related Internet services directly from the real world. This facilitates a completely new real-world interface to access physical-entity-related services without the need for a conventional Internet search. A tourist could, for example, obtain information about buildings and places by simply pointing with his cellular phone at a building instead of performing a cumbersome full-text search. Any output interface like an audio headset or the cell phone display could be used to unobtrusively present the resulting information as illustrated in Figure 2. In this use case, identifying the monument can technically be realized by using physical attributes of the building such as, for example, its GPS position, estimated via the user's GPS coordinates and an electronic compass. These attributes will be used by the Augmented Internet infrastructure to build and maintain a connection between entity-related information and the physical entity itself.



Figure 2: Tourist application: Accessing information about buildings

In a similar manner, an Augmented Internet service could identify objects based on, for example, Radio Frequency Identification (RFID) tags and could execute Internet services related with the referenced object. For example, a user could “click on” a library book, thereby executing the online service to renew the book. Furthermore, a selection menu could be displayed to the user, providing additional options like accessing the bibliographic information of the book. All book-related information are cumulated in the book's virtual representation, provided and maintained by the information-centric network infrastructure. A simple Augmented Internet service could also enable a user to store personal notes in this virtual representation, hence directly binding them to the book instead of storing the notes in some separated text document that is difficult to retrieve later on.

Furthermore, an Augmented Internet could enable users to use web methods like bookmarking in the real world: a user could bookmark interesting products for later reference during a shopping tour by “clicking on” the physical object with the cell phone. Likewise, a user could add a virtual post-it to a specific object to leave a message for other users that could retrieve this message later on (e.g., by “clicking on” the object again).

Related scenarios have been discussed in some other research projects [1, 4, 6, 8–10, 12]. Our aim is to extract the common requirements of those scenarios and realize support for such scenarios based on the Network of Information.

2.3 Personal Mobile Communications Redux

The personal mobile scenario is intended to illustrate the effectiveness of using the Network of Information approach in supporting mobile users. This is done by keeping copies of the mobile user's data “close” to possible attachment points of the Internet. Opportunistic methods for pushing and pulling data lead to cheaper, faster, and more stable retrieval of data. We look at a mobile user's mailbox as the example data. This scenario is a non-dissemination scenario, where the user does not want the data to reach a broad audience. The intention is to explore the benefits of such a scenario in an information-centric network.

The scenario (Figure 3) is instantiated as a train passenger that wishes to retrieve some larger amount of data. The passenger's computer is equipped with a lower capacity cellular link (GPRS/UMTS/HSPA) as well as a higher speed WLAN connection (802.11a/b/g/n). When the train stops at stations, the passenger may use WLAN but is otherwise restricted to use the slower cellular network.

The data to be retrieved is at A (on the Internet) and the train is about to pass station B.

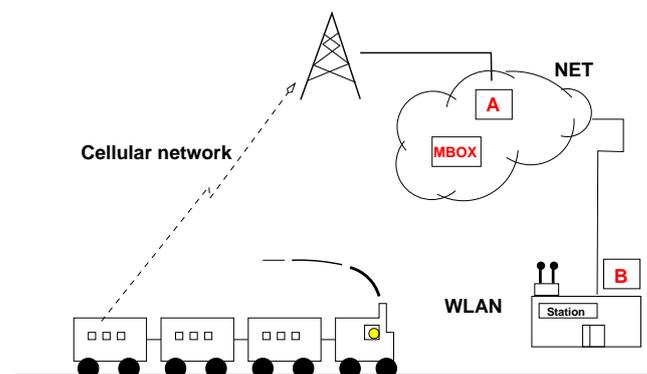


Figure 3: Train scenario: Preparing data for access

By replicating the desired data to the fixed stations (or other pre-distributed caches), the user may gain throughput and battery life time (by turning off interfaces until arrival at the next cache).

The scenario works as follows:

- The user initiates retrieval over the low capacity cellular radio link (GPRS/UTMS/HSDPA).
- The system concludes that retrieval will fail because of capacity constraints and interrupts the transfer.
- The system determines that (station) B is the next available WLAN access point.
- A is informed that communication will resume at B and pushes the remaining data to B.
- When the train arrives at B the communication resumes.

An alternative procedure is that the user communicates directly with B through the cellular link and sends a “pre-fetch” message to B. B would then follow by directly pulling the data from A.

To achieve this the network must have the ability to locate the closest copy of an object and be able to “move” the copy closer to where the user will be next or is likely to be next. This replication may either be triggered by some (low-bandwidth) control message or perhaps automatically by the network itself. Also, either the network or the client needs a way to predict which stations would be suitable to replicate data to. In a train scenario geographical data is not unlikely to exist, and in other daily situations it is not unlikely that the mobile device recognizes frequently seen open access points (e.g. FONtm nodes, see www.fon.com). It would also be desirable for the network to be able to estimate the amount of available network resources between two communication points for replication decisions.

These methods would provide benefits as compliments to the current synchronous Internet model, by providing an asynchronous model of communication to novel mobility aware applications. This would also allow decoupling of the wanted information from the originating nodes. It would provide bandwidth adaptation and improve usage enabling the use of transient access (which is often the case with short-ranged high-capacity radios).

3. CONCEPTS AND ARCHITECTURE

Several base concepts have been identified to enable the development of the Network of Information architecture in 4WARD. A fundamental concept that influences the overall architecture is the identifier/locator split. In NetInf, the functional overloading of IP addresses as both identifiers and locators will be eliminated via a conceptionally clean split of those two functions. Each object will have a locator pointing at the location of the data object in the network and a separate identifier, enabling the persistent identification of the object regardless of possible location changes or replication in the network.

The Network of Information differentiates between two kinds of objects as alluded already in Section 2.1: *data objects* and *information objects*. The central idea is that data objects contain the payload (e.g., an mp3 file with a certain encoding, like 128 kbps), while information objects represent a semantically meaningful entity like a certain song (e.g., Beethoven’s 9th symphony). Information objects enable users to find content independent of its specific representation (e.g. encoding) and independent of certain characteristics that *might* not be relevant to the user (e.g. the performing orchestra). Information objects themselves may be hierarchical as explained in Section 2.1, i.e., an information object does not necessarily point to a data object, it may also point to one or more other information objects. Furthermore, information objects can represent and describe physical entities of the real world (e.g. a building or a book) as described in Section 2.2, which constitutes a further generalization compared to existing ideas. In order to make this system work, the bindings and indirections between the different objects need to be handled efficiently, and accordingly this area constitutes an important topic for our research. Another important question of the ongoing research is to examine which parts of the metadata are stored in the information objects and which parts are stored in the data

objects.

Another concept is to reduce the need for governance in the Future Internet. The Future Internet will play an even more important worldwide role than today’s Internet already does. Hence, its availability and evolution should not be hindered by any single government or organization. As a result, the Network of Information will be self-managing as far as possible. For example, it will provide a flat naming mechanism (e.g., based on cryptographic identities) that enables any entity to name and access objects without having to rely on a naming authority.

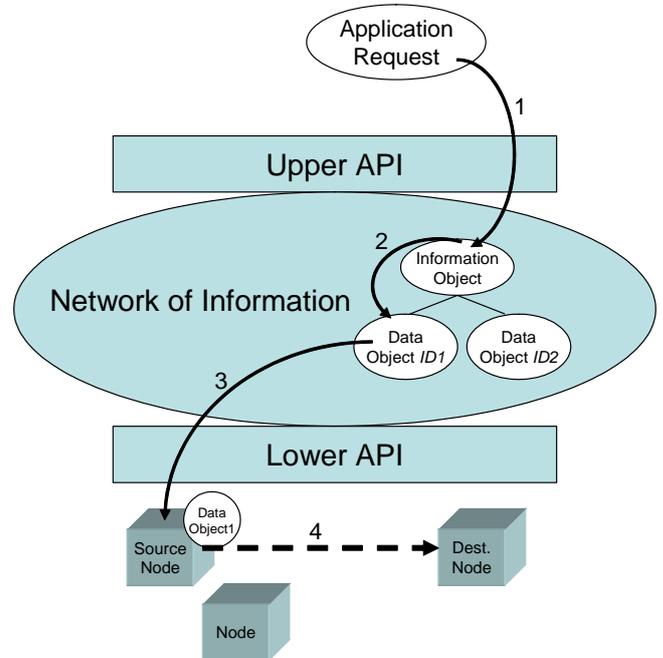


Figure 4: The Network of Information architecture overview

The Network of Information architecture has two main APIs (Fig. 4):

- The *upper API* provides access to the Network of Information for NetInf users (i.e., users or applications). NetInf users can access information objects and related data objects based on semantic information and metadata, independent of the storage location in the network.
- The *lower API* provides access to the underlying network to transport and distribute the data itself.

Above the upper API, several interfaces can coexist. Typically, a publish/subscribe mechanism, a socket-like interface, or any other interface could be used to update and/or search for data/information objects in the Network of Information. Likewise, the lower API does not rely on any particular network protocol. Therefore, the concept can support, for example, the current IP protocol (v4/v6) as well as post-IP solutions.

When a NetInf user wants to access an object via the upper API, the network first determines the list of information objects that may be of interest based on metadata provided by the NetInf user (1). Second, the desired data object(s)

has/have to be selected from the list of attached data objects (2). Finally, each data object identifier has to be resolved into an appropriate locator/set of locators (3) and has to be transferred to the requesting NetInf user via the lower API (4).

As alluded in the previous example, a single identifier can be bound to multiple locators. This can happen for two different reasons. First, the NetInf concept inherently supports replication of data objects to ensure data availability and to improve performance. Sections 2.1 and 2.3 have provided application examples for active management of data storage in the network. A certain replica can then either be specified by the NetInf user directly or selected by the Network of Information based on certain metrics such as server load (the access remains however completely transparent to the NetInf user). The second case involves very large data objects. In fact, files stored in the Internet are getting larger and larger each year. A good example is the development of websites like YouTube that provide millions of videos. To ensure an efficient data transfer in this case, large objects can be cut into smaller pieces (i.e. *chunks*) that can be downloaded from different locations simultaneously. This way, a BitTorrent-like mechanism can be used for downloading.

4. RESEARCH SCOPE AND OUTLOOK

The main objective of this work (NetInf) is to create a new framework based on the information-centric paradigm, creating a Network of Information where information objects have their own identity and are no longer bound to specific hosts or locations. Information objects are not only useful for traditional content distribution but should also be used for conversational services like telephony, and store-and-forward services like email. This objective is well in line with new Internet design discussions, e.g., by Van Jacobson [2, 3].

Of central concern to this framework is the anticipated explosion in the number of things to name and address, and the systemic changes in the trust reference relationships, something also recognized by the DONA architecture [5]. In NetInf, trust will need to be disassociated from where information resides. In the current Internet, authentication of the nodes where the information resides seems to be sufficient, although in practice it has often proved to be a fallacy. In NetInf it is authentication of actual information objects that matters, not authentication of storage locations. Also authorization needs to be performed on a per-information object level. Further, in today's networks, only the source node where the information resides is addressed. Clearly, there are many more information objects than nodes, and replication of the same object needs to be taken into consideration in order to enable more flexible and scalable data dissemination mechanisms than today's. Much of the challenge lies in the retention of the information object book-keeping, for example, the desire to resolve the location of data for retrieval purposes. The requirement on the resolution services in the network demands that it has to be scalable, efficient, and suit a more dynamic environment than existing resolution services like DNS. Mobility and multi-access are also fundamental to the NetInf architecture: we anticipate that the vast majority of information will be generated and consumed by nodes that have access to multiple access networks simultaneously and are typically mobile (or at least do not use the same, fixed attachment point all the time). The

content distribution scenario (Section 2.1) and the personal mobile scenarios (Section 2.3) highlight the dynamicity and efficiency demands, while the Augmented Internet scenario (Section 2.2) adds another dimension on how well the new network needs to scale.

The research area that is opening ahead of us is quite broad and one practical approach in addressing the challenges is based on scenario use cases, which assist in having reality checks for the proposed solutions. In other words, scenarios can be used to measure how well different candidate solutions can address these research problems/challenges. In a first approach we need to look at how the numerous information objects should be "registered" by information producers and located by information (re)distributors and consumers. The ability to find the right information object broadens the traditional network scope and the border between information search and traditional look-up resolution may start to blur. For instance, the addition of metadata to information objects is one way to handle the larger network addressing needs. Another approach is to enhance content delivery using network and end-user device capabilities, for example, network features like multicast and caching can help to address the efficiency and scalability challenges. One clear research direction comes from the caching capability, it indicates that data objects will be replicated in the network and multiple instances of the information can offer several retrieval alternatives for the end user. In this case, metadata can also be a feasible research direction to help in addressing these several copies of the same object. These cached copies can be available throughout the network and the research scope should include how to efficiently use caching. Furthermore, we are convinced that many more research challenges will arise along the way of developing a new architecture from scratch that is based on a completely different (i.e. information-centric) paradigm.

It is clear that the direction of research will mostly center around how to get this information centric framework to work reliably. The focus, like with the scenarios, concentrates on providing a clear interface for applications, traditional or brand new. However, providing a generic interface for all applications, many with differing scope and usage situation, requires that the NetInf works well for many purposes. As a consequence, we are also investigating the existence of several coexisting naming and resolution solutions in our research, as long as multiple, potentially complementing solutions support the overall naming and resolution framework.

5. CONCLUSION

The shortcomings of the current Internet have been well documented in recent years and a significant effort has started in order to redefine the core architecture of the future Internet. Paraphrasing Edward Tufte, we argue that the main handicap of the contemporary Internet is that it segregates information by its means of production and consumption. Tufte reminds us that the "user shouldn't have to know what an operating system is, or even an application". Similarly, we argue that users should not have to concern themselves with issues such as connectivity in a multiaccess world, content replication and retrieval when storage is abundant, or rely on any particular provider to retrieve or distribute their content.

We have presented scenarios for a future network archi-

ture that is information-centric and introduces research issues that have to be considered as we develop this architecture. We have introduced the main concepts underlying NetInf, our proposition for the future Network of Information and attempted to define a bit better the problem scope, which admittedly is quite broad. Our expectation is that the research agenda can only be drafted at this stage of development as many topics will emerge in the process.

Developing testbeds to showcase our ideas is high on our research agenda. So is modeling and simulating large NetInf networks and we expect that in the coming months we will be able to quantify the performance of NetInf in a concrete manner.

6. ACKNOWLEDGMENTS

We would like to thank our colleagues in the 4WARD Work Package 6, aptly titled *Network of Information*, for fruitful discussions. We would also like to thank Petteri Pöyhönen, Sabine Randriamasy, and Daniel Horne for review comments. We also express our gratitude to the General Chairs of MobiMedia 2008 for the opportunity to present our ideas at such an early stage of development.

This work has been carried out in the framework of the IST 7th Framework Programme Integrated Project 4WARD, which is partially funded by the Commission of the European Union. The views expressed in this paper are solely those of the authors and do not necessarily represent the views of their employers, the 4WARD project, or the Commission of the European Union.

7. REFERENCES

- [1] R. Barrett and P. P. Maglio. Informative things: How to attach information to the real world. In *UIST '98: Proceedings of the 11th Annual ACM Symposium on User Interface Software and Technology*, pages 81–88, New York, NY, USA, 1998. ACM Press.
- [2] V. Jacobson. A new way to look at networking. online: <http://video.google.com/videoplay?docid=-6972678839686672840>, August 2006.
- [3] V. Jacobson, M. Mosko, D. Smetters, and J. Garcia-Luna-Aceves. Content-centric networking. Whitepaper, Palo Alto Research Center, January 2007.
- [4] T. Kindberg, J. Barton, J. Morgan, G. Becker, D. Caswell, P. Debaty, G. Gopal, M. Frid, V. Krishnan, H. Morris, J. Schettino, B. Serra, and M. Spasojevic. People, places, things: Web presence for the real world. *Mobile Networks and Applications*, 7(5):365–376, October 2002.
- [5] T. Koponen, M. Chawla, B.-G. Chun, A. Ermolinskiy, K. Kim, S. Scenker, and I. Stoica. A data-oriented (and beyond) network architecture. In *SIGCOMM*. ACM, 2007.
- [6] P. Ljungstrand, J. Redström, and L. E. Holmquist. Webstickers: Using physical tokens to access, manage and share bookmarks to the web. In *DARE '00: Proceedings of DARE 2000 on Designing Augmented Reality Environments*, pages 23–31, New York, NY, USA, 2000. ACM Press.
- [7] T. Plagemann, V. Goebel, A. Mauthe, L. Mathy, T. Turetli, and G. Urvoy-Keller. From content distribution networks to content networks - issues and challenges. *Computer Communications Journal, Elsevier*, 29(5), 2006.
- [8] S. Pradhan, C. Brignone, J.-H. Cui, A. McReynolds, and M. T. Smith. Websigns: Hyperlinking physical locations to the web. *Computer*, 34(8):42–48, 2001.
- [9] J. Rekimoto and K. Nagao. The world through the computer: Computer augmented interaction with real world environments. In *UIST '95: Proceedings of the 8th Annual ACM Symposium on User Interface and Software Technology*, pages 29–36, New York, NY, USA, 1995. ACM Press.
- [10] H. Tarumi, K. Morishita, M. Nakao, and Y. Kambayashi. Spacetag: An overlaid virtual system and its applications. In *Proceedings of the International Conference on Multimedia Computing and Systems (ICMCS'99)*, volume 1, pages 207–212. IEEE, 1999.
- [11] E. R. Tufte. *Beautiful Evidence*. Graphics Press, 2006.
- [12] R. Want, K. P. Fishkin, A. Gujar, and B. L. Harrison. Bridging physical and virtual worlds with electronic tags. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 370–377, New York, NY, USA, 1999. ACM Press.