Computer Networks 56 (2012) 3705–3722

Contents lists available at SciVerse ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

Supporting the Web with an information centric network that routes by name

A. Detti *, M. Pomposini, N. Blefari-Melazzi, S. Salsano

CNIT, Department of Electronic Engineering, University of Rome "Tor Vergata", Via del Politecnico 1, Rome, Italy

ARTICLE INFO

Article history:
Received 26 February 2012
Received in revised form 9 July 2012
Accepted 9 August 2012
Available online 20 August 2012

Keywords:
Internet architecture
Future internet
Information-centric networking
Routing
Caching
Scalability

ABSTRACT

Information Centric Networking (ICN) is a new paradigm in which the network layer provides users with content, instead of providing communication channels between hosts, and is aware of the name (or identifiers) of the contents. A fundamental ICN operation is the routing of content requests towards a node that is able to provide the requested content. To meet this goal, different routing architectures have been proposed so far.

In this paper, we consider a network that uses a *routing-by-name* architecture, i.e. content requests are routed on the base of the content name by using a *name-based routing table*. We focus on the scenario of fetching Web contents, assuming to use ICN in place of traditional TCP/IP means. In this scenario we need to handle tens of billions of name-based routes, due to the high numbers of Web contents and to the limited aggregability of their names. Consequently, re-using the existing architecture of an IP router would result in two severe problems. First, the current Forwarding Information Base (FIB) technology is unable to contain all name-based routes. Second, implementing a so large Routing Information Base (RIB) requires a very costly hardware. In order to overcome these problems, we propose a routing-by-name architecture, named Lookup-and-Cache, where the FIB is used as a *cache of routes*, while the RIB is stored in a remote and centralized routing engine. By analyzing real Internet traces, we prove the effectiveness of the proposed architecture, which we also show to be feasible with current technology. In fact, our ICN nodes require to have only a limited set of routes in their FIB, even when supporting a high number of traffic flows.

We have implemented our proposed Lookup-and-Cache solution within the CCNx software framework and we used this implementation to assess system performance, such as download delay, lookup rate and fairness.

The paper is completed with a discussion on how ICN can be used not only to fetch Web contents but also for other scenarios.

1. Introduction

Information Centric Networking (ICN) is a concept proposed some time ago under different names [1,2], which is attracting more and more interest, recently (see e.g. the papers [3–6], the workshop [36] and the projects [6–11]). ICN proposes a shift from the traditional host-to-host communication to a content-to-user paradigm, which focuses on the delivery of the desired content to the intended users. The basic functions of an ICN infrastructure are to: (i) address contents, adopting an addressing scheme based on names (identifiers), which do not include references to their location; (ii) route a user request, which includes a "destination" content-name, toward the "closest" copy of

* Corresponding author. Tel.: +39 06 72597285.
E-mail addresses: andrea.detti@uniroma2.it (A. Detti), matteo.pomposini@uniroma2.it (M. Pomposini), blefari@uniroma2.it (N. Blefari-Melazzi), stefano.salsano@uniroma2.it (S. Salsano).

1389-1286/$ - see front matter
http://dx.doi.org/10.1016/j.comnet.2012.08.006

the content with such a name; and (iii) deliver the content back to the requesting host.

ICN basic concepts were first proposed in [1], with the so-called TRIAD architecture. TRIAD is an overlay network of content-routers. Content routers *route-by-name* content requests, i.e. they route content requests on the basis of the content-name towards the "best" suitable server. Once the server is reached, the content is delivered to the user by means of a plain TCP/IP session.

A second seminal paper is [2], which proposed the so-called DONA architecture. DONA introduced important issues, such as security and naming, and proposed to use self-certifying-names. Similarly to TRIAD, DONA uses an overlay network to route-by-name content requests towards the best server, and then delivers contents via TCP/IP. Therefore, the route-by-name procedures are involved only at the start of a download session and the network path followed by a content request could be different from the one used to deliver the data using TCP/IP.

The third seminal paper on ICN is [3], which proposed a so-called content-centric network (CCN). CCN does not rely on an underlying IP infrastructure. Thus, CCN could completely replace the IP layer. CCN procedures route-by-name requests for *chunks* of contents. Furthermore, CCN mechanisms also handle the delivery of chunks to the user, rather than delegate this operation to TCP/IP. To download the whole required content, a user sequentially downloads all its chunks. Hence, route-by-name procedures are continuously involved during a content download and require line-rate processing speed. The path selected to route-by-name a chunk request is also used in the reverse direction to deliver the corresponding chunk data. The matching between the request path and the data path facilitates the exploitation of *en-route* (or *in-network*) caches within network nodes [14]. CCN attracted a significant number of researchers to the ICN field, also thanks to the development and release of an open source implementation of the main CCN concepts, called CCNx [17]. In the following sub-section we report our view on pros and cons of ICN.

1.1. Pros and cons of ICN

In our view, an ICN would offer the following advantages.

Efficient routing by name: even though Content Delivery Networks (CDNs) offer efficient mechanisms to route contents, they cannot use network resources in an optimal way because they operate over-the-top, i.e. without knowledge of the underlying network topology. ICN would let ISPs perform native content routing with improved reliability and scalability of content access. This would be a built-in facility of the network, unlike today's CDNs.

In-network caching: caching enabled today by off-the-shelf HTTP transparent proxies requires performing stateful operations, because the delivery of cached content is based on the connection-oriented TCP. The burden of a stateful processing makes it very expensive to deploy caches in nodes that handle a large number of user sessions. ICN would significantly improve efficiency, reliability and scalability of caching, with interesting applications especially for video.

Simplified support for peer-to-peer like communication: without the need of overlay dedicated systems. Users could obtain desired contents from other users (or from caching nodes) thanks to content-routing functionality, as it is done today with specialized applications, which, once again, do not have a full knowledge of the network and involve only a subset of possible users.

Per-content quality of service differentiation: providing different performance in terms of both transmission and caching. Network operators (especially mobile ones) are already trying to differentiate the quality and the priority of content, but they are forced to resort to complex and hardly scalable deep packet inspection technologies. ICN would let network operators differentiate the quality perceived by different services without complex, high-layer procedures, and off-load their networks via caching, a very handy functionality, particularly for mobile operators who can differentiate the quality and the priority of content transferred over the precious radio real estate.

Handling of mobile and multicast communications: simplifying handovers and stateful nodes. As regards handovers, when a user changes point of attachment to the network, she will simply ask the next chunk of the content she is interested in, without the need of storing states; the next chunk could be provided by a different node than the one that it would have been used before the handover. Similar considerations apply for multicasting. Several users can request the same content and the network will provide the service, exploiting caches, without the need of overlay mechanisms.

Content-oriented security model: securing the content itself, instead of securing the communications channels allows for a stronger, more flexible and customizable protection of content and of user privacy. In today's network contents are protected by securing the channel (connection-based security) or the applications (application-based security). ICN would protect information at the source in a more flexible and robust way than delegating this function to the channel or the applications [4]. In addition, this is a necessary requirement for an ICN: in-network caching requires to embed security information in the content data-unit, because content may arrive from any node and we cannot trust the serving node; thus, end-users must be able to verify the validity of the received data and caching nodes must do the same to avoid caching fake contents.

Support for time/space-decoupled model of communications: simplifying implementations of publish/subscribe service models, and allowing "pieces" of network, or sets of devices to operate even when disconnected from the main Internet (e.g. sensors networks, ad-hoc networks, vehicle networks, delay-tolerant-networks, social gatherings, mobile networks on board vehicles, trains, planes).

On the cons side, ICN has some drawbacks and challenges. A first, obvious, con is that it requires *changes in the basic network operation,* which per se is already a big obstacle to take-up of this approach. A second con is that it raises *scalability* concerns: (i) the number of different contents and corresponding names is much bigger than the number of host addresses; this has implications on the size of routing tables and on the complexity of lookup functions and (ii) in some proposed ICN architectures [3],