Journal of Network and Computer Applications 87 (2017) 73–86

Contents lists available at ScienceDirect

Journal of Network and Computer Applications

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

NDNrel: A mechanism based on relations among objects to improve the performance of NDN

Rodolfo S. Antunes*, Matheus B. Lehmann, Rodrigo B. Mansilha, Luciano P. Gaspary, Marinho P. Barcellos

Institute of Informatics – Federal Unifersity of Rio Grande do Sul (INF/UFRGS), Av. Bento Gonçalves 9500, Setor 4, Porto Alegre, RS, Brazil

ARTICLE INFO

Keywords:
Information-centric networks
Content-centric networks
Content distribution

ABSTRACT

Named-Data Networking (NDN) is a promising architectural approach that focuses on the efficient distribution of data objects. Each of these objects represents an individual piece of content that is uniquely named and can be cached by network nodes. Recent work on Information-Centric Networking (ICN) explores new uses of the data object concept to enable advanced applications based on content distribution. Such work includes the use of metadata and manifests to describe contents as sets of multiple correlated objects. In this paper, we extend our previous work on a model and a backward-compatible mechanism, named NDNrel, that allows publishers to distribute contents as related objects. The new version of NDNrel reduces the cost of distributing relation information to clients by leveraging the usage of metadata and manifests on ICN architectures. Also, it considers additional aspects such as the authentication of relations and data that comprise the content. The evaluation of NDNrel is revisited and extended to analyze the performance and overhead caused by the usage of the proposed relation model in two scenarios: distribution of multimedia and Web content. Results demonstrate that using NDNrel to structure the contents as objects and relations outperforms default NDN, reducing the average client download time and the global network traffic in at least 28% and 34% respectively.

1. Introduction

Information-Centric Networking (ICN) proposes a communication architecture with mechanisms tailored to improve the performance of content distribution (Zhang et al., 2013; Bari et al., 2012; Xylomenos et al., 2013). The efficiency of these mechanisms is directly impacted by the structure of the contents. More specifically, how the concept of uniquely identified data objects is applied to different types of information.

Some dominant types of content, such as video (Cha et al., 2007) and Web pages (Wang et al., 2014), can be regarded as sets of individual data elements. To trivially distribute such contents through an ICN, a publisher naively stores all associated elements in a single object. However, as shown later, this method negatively impacts the performance of content distribution, because of oversized objects and excessive data redundancy. More advanced distribution mechanisms could help publishers distribute more efficiently sets of individual data elements.

Manifests and metadata enable the design of new distribution mechanisms for ICN applications, and this is currently a highly debated theme (Mosko et al., 2015) in the research community. This study explores how they can make the distribution of multiple-object contents more efficient. More specifically, our proposal allows publishers to describe a content as a set of individual objects, by establishing *relations* among them. In a nutshell, a relation is a link between two objects indicating that the data from one complements in some way the data from the other. Following this concept, a content becomes a collection of multiple objects and the respective relations that characterize the interactions among their data. These definitions are employed to formulate a model that publishers can use to distribute contents as sets of related objects.

This paper expands our study, first introduced in Antunes et al. (2015), on a *backward-compatible* extension to the Named-Data Networking (NDN) (Zhang et al., 2014) architecture that enables the distribution of contents as related objects. Compared to our previous work, this paper includes two main contributions. First, it describes the new version of the relations mechanism for the NDN architecture, henceforth called NDNrel. This new version reduces the cost of distributing relation information to clients by incorporating recent advances in the use of metadata and manifests on ICN architectures.

* Corresponding author.
E-mail addresses: rsantunes@inf.ufrgs.br (R.S. Antunes), mblehmann@inf.ufrgs.br (M.B. Lehmann), rbmansilha@inf.ufrgs.br (R.B. Mansilha), paschoal@inf.ufrgs.br (L.P. Gaspary), marinho@inf.ufrgs.br (M.P. Barcellos).

http://dx.doi.org/10.1016/j.jnca.2017.03.014
Received 13 June 2016; Received in revised form 28 February 2017; Accepted 14 March 2017
Available online 16 March 2017

The new version also considers additional design aspects, such as the authentication of relation information and the data that comprises a content. Second, this paper revisits the evaluation of the proposed mechanism. It adds a second scenario, based on the Web. This scenario was chosen to stress the relations mechanism, to determine potential gains and measure traffic overhead when objects are smaller and induce larger overheads.

Throughout this paper, we show that the concept of relations behind NDNrel is flexible, since it does not impose any restriction on the way publishers can model contents. Further, the mechanism implementation benefits from recent advances in ICN research, such as manifest objects, to improve the performance of applications that employ relations. The comparison with default NDN demonstrates that the achieved data redundancy elimination by NDNrel leads to promising increases in user and network performance, even in an unfavorable scenario with significant overhead from relations. Comparatively to the state-of-the-art, NDNrel enables relations to be used in arbitrary novel ways beyond those already explored. Furthermore, it stands out as a proposal tailored to NDN, presenting a backward-compatible approach that does not require modification to routers.

The remainder of the paper is organized as follows. Section 2 addresses work related to NDNrel. Section 3 introduces the fundamental concepts of the relations model instantiated by NDNrel. Section 4 discusses the design aspects considered on the implementation of the relation model in the NDN architecture. The evaluation methodology and main results are presented in Section 5 for the multimedia content scenario and Section 6 for the Web content scenario. Section 7 discusses qualitative benefits of NDN regarding its flexibility and generalization. Section 8 briefly discusses the role of NDNrel in the comparison of NDN with traditional content delivery schemes. Section 9 closes the paper with final remarks and future work considerations.

2. Related work

This section first discusses related work that applies the concept of content and semantic relations at the application layer. Next, it focuses on relation mechanisms proposed in the context for NDN and other ICN architectures.

The concept of relations is explored in different systems for content distribution. The most popular example are Web pages, which employ relations as a fundamental concept for the creation of complex documents with the use of multiple objects. In the context of multimedia, relations are applied to enable the creation of complex contents based on multiple audio and video channels. The DASH standard from ISO/IEC (2014) utilizes a manifest file to specify a multimedia content with various components available in different HTTP URLs. The client can select a subset of these components for playback depending on its quality requirements. Also related to multimedia, SMIL (Bulterman et al., 2008) is a markup language used for the composition of rich multimedia presentations based on audiovisual elements stored in different objects. The above examples make use of a specific format to describe how contents are formed with data acquired from different source objects. However, these formats are created targeting specific application domains and they are not applicable to other types of contents.

The literature of ICN includes three pieces of work (ICN-RE (Perino et al., 2012), NetInf (Dannewitz et al., 2013), PSIRP/Pursuit Blackadder (Tagger et al., 2013)) related to our NDNrel proposal.

ICN-RE (Perino et al. 2012) employs a concept similar to relations for implicit redundancy elimination of object data. In a nutshell, ICN-RE identifies, isolates and publishes byte-identical portions of different contents as a single object. The remaining parts are published individually. The mechanism uses a meta-object that lists the names of all objects that should be downloaded to rebuild the original content. ICN-RE uses the concept of relations in the meta-object to enable redundancy elimination of objects data. However, the format of this meta-object is strictly designed for the mechanism of implicit redundancy elimination. More importantly, ICN-RE requires modifications in ICN routers.

NetInf (Dannewitz et al. 2013) introduces the concept of information objects (IOs), which are collections of metadata and pointers to actual data objects. The metadata contained in IOs allows clients to perform semantic queries about published contents. The PURSUIT project proposes a publish/subscribe architecture including an information-centric middleware (Tagger et al., 2013) for the Blackadder prototype based on semantic technologies and metadata. The proposed middleware enables, among other things, establishing relations through common semantic attributes. The fundamental difference of these works on ICN is the focus on ICN architectures other than NDN. Thus, their findings cannot be directly extended to NDN due to specificities in naming (e.g., flat IDs under nested scopes) and routing mechanisms (e.g. separated control and data planes (Jokela et al., 2015). Further, those studies do not provide a detailed evaluation on the potential network performance gains when leveraging content and semantic relations among data objects.

Compared to the state-of-the-art, NDNrel can achieve at least the same results by using a similar content structure. Further, NDNrel provides additional features, such as high-level content description, content sharing between multiple publishers, and flexible data redundancy. These advantages are made possible by a *relations model*, proposed in Section 3, used to describe the contents distributed with NDNrel. The model is intended to be generic and flexible, hence it can be employed in various scenarios besides identifying duplicate content data. For instance, Section 7 discusses how NDNrel can be used to describe sensor data from multiple devices in IoT. In this scenario, NDNrel would allow administrators to structure the existing data efficiently.

NDNrel enables publishers to use objects from third-party sources to describe their contents. This allows NDNrel to provide higher gains (when compared with ICN-RE). Lastly, NDNrel has a more flexible concept of data redundancy. ICN-RE can only identify redundancy when the contents are bitwise equal. NDNrel, in its turn, allows publishers to represent their content as they wish. For example, a movie content can have in its description several video objects that are not bitwise equal. Because of the content representation, consumers know that they need only one of the video objects, avoiding the transmission of unnecessary data and redundancy.

In summary, our work is novel in exploring the design aspects of introducing backwards-compatible relations mechanisms for NDN to allow publishers modeling their contents in innovative ways. We argue that allowing publishers to use their knowledge about contents to define relations among objects can bring benefits beyond those achieved by mechanisms such as implicit decomposition. While current proposals for relations in content distribution are specific to their application domains (e.g. multimedia, P2P), NDNrel is designed to be generic and allow current and future ICN applications to natively benefit from their intrinsic semantic relations.

3. A Model for relations among objects in ICN

In this section, we first discuss the fundamental concept of relations employed in NDNrel, our backward-compatible relations mechanism proposal for NDN. Next, we exemplify the use of relations when modeling contents.

3.1. Fundamental concepts

The proposed NDNrel model employs three fundamental elements: object, relation, and content. An *object* represents an individual piece of information that is uniquely identified in the network. The model also assumes that objects are the basic data element that clients can request to the network. This representation is equivalent to the one