Network Cost Services for Network-Aware FI Applications

Georgios Gardikis, George Xilouris
National Centre for Scientific Research "Demokritos", Media Networks Laboratory
Aghia Paraskevi 153 10, Greece
emails: {gardikis; xilouris}@iit.demokritos.gr

***Abstract*— While the explicit provision of Network Awareness to applications towards their self-optimisation has attracted significant attention during the past decade, no universal, feasible and scalable solution has been proposed. In this context, this position paper proposes the deployment of universal Network Cost Services in the FI scene, explicitly providing the "network cost" between two or more endpoints for a specific service and thus assisting the optimization of current and emerging communication paradigms, such as peer-to-peer communication, distributed caching, information-centric networking and in-network cloud services. A generic Network Cost provision framework is described, accompanied with a preliminary proof-of-concept implementation. Deployment considerations and open research issues are also discussed.**

***Index Terms*— Network-Aware Applications, Network Cost, traffic optimisation**

I. Introduction

The explicit provision of Network Awareness to applications is considered a promising aspect within the Future Internet scene. It involves the on-demand signaling of specific network-level metrics to querying applications, with the aim of optimizing application behavior and performance.

This application-network coupling falls within the generic "Resources Awareness" FI design principle, as identified in the latest FIArch document [1], which refers to the explicit interchange of resource availability information between different Internet architectural layers.

While such an interplay has been extensively studied by several research efforts, especially during the past decade, no universal solution has been proposed which can be scalable, widely applicable, and which respects the independence of the different architectural layers and also the privacy of the respective business actors.

In this context, this position paper suggests a framework for providing explicit network awareness in the form of a network cost vector between two endpoints. The rationale and logic of this framework –based on an evolutionary architecture and leveraging existing trends in standardization bodies– aims at assisting traffic optimization at application layer, exploiting network status information, while at the same time addressing scalability, privacy and security issues.

The next sections of the paper better illustrate the addressed problem, present the proposed framework along with a preliminary implementation, and discuss a number of associated issues.

II. Problem Statement

For many decades, the success of the current Internet has been founded on the simple principle of host-to-host communication based on the functional logic of (mainly) isolated architectural layers. Network-awareness is implicitly derived, mostly using in-band mechanisms, such as the TCP congestion control.

While several mechanisms have been proposed for network-side explicit notification of network status and/or topology, most of them have been considered inapplicable, either because they are not scalable and/or they impose considerable additional computing/communication overhead on core network elements. Another issue is the privacy of the network operator and the general unwillingness to publicly expose detailed network topology and utilization information. For these reasons, existing application-layer mechanisms, as described above, are still dominant for implicit network assessment.

However, as the Internet community is currently witnessing a gradual paradigm shift from host-centric to content-centric communication, existing implicit mechanisms may be not as efficient. In many current and emerging communication architectures (e.g. peer-to-peer, CDNs, in-network clouds, content-centric networks etc.), the content is distributed into the network, rather than concentrated in a single server. Therefore, its retrieval requires the simultaneous communication with tens, hundreds, or even thousands of peers. Assessing the communication with each of these peers at application level via the usual in-band methods would pose significant overhead in the network, besides requiring a considerable amount of time.

A more efficient alternative would be the establishment of an explicit Network Cost Service (NCS), offered by the network operator or a third-party entity and providing on-demand a "Network Cost" (NC) assessment for any given pair or group of network endpoints. Such an assessment could be instantly exploited at application level in order to optimize operations such as in-network content placement, peer selection, prioritization and load balancing.

A similar approach is promoted by the Application Layer Traffic Optimisation (ALTO) WG of the IETF, which, in addition to identifying the aforementioned need for a network cost service [2] is currently standardizing a protocol for communicating network cost maps, mainly focusing on optimizing peer-to-peer transfers [3].

The present proposal considers the ALTO approach as a starting point and a) considers several enhancements such as a more generic representation of the network cost, serving a much wider spectrum of use cases besides p2p and b) discusses and implements a complete framework to enable such a service on an operational network leveraging active and passive measurement techniques.

Candidate use cases for the proposed NCS include, but should not be restricted to:

- peer-to-peer file sharing and streaming - for optimizing p2p communication via selection of "best peers" i.e. peers with lower NC
- one-to-one and one-to-many realtime audiovisual communication – for a priori optimizing stream format and rate according to NC
- distributed caching, Content-Delivery Networks (CDN) and Information Centric Networking (ICN) – for choosing optimal locations for caching content within the network
- in-network clouds – for efficient load balancing, traffic distribution and resource transfer to optimal locations, including Virtual Machine (VM) migration.

III. A Network Cost Service Framework

As aforementioned, Network Cost can be provided as a service either by the Network Operator or a third party. However, it seems that the more natural solution is that the NOs provide the service themselves, so that the can control the entire network measurement chain, from the network elements' monitoring agents up to the interface with the querying applications.

A high-level conceptual diagram is shown in Fig.1.

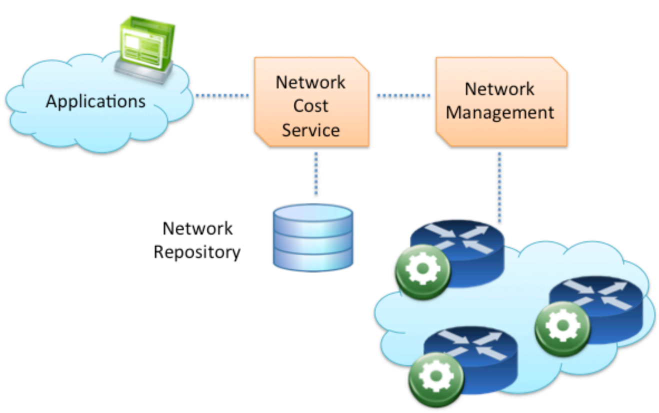

Fig.1. High-level structure of a network cost provision framework

The operation of a network cost framework would normally be restricted within a single administrative network domain (e.g. an ISP network or a backbone provider), although multi-domain scenarios could also be envisaged, as will be described. Within the network, existing management approaches can be used for collecting network metrics. To assist network cost estimation, it would be desired that active measurements, in addition to passive ones, are employed. Active measurements include the controlled injection of artificial (probe) traffic across network paths so as to measure them. This task could be undertaken by dedicated network probe modules, strategically placed within the network. Another option would be to enhance existing monitoring agents within the network elements so as to support the establishment of active measurement sessions on-demand. Our proof-of-concept implementation (cf. next section) includes properly modified SNMP agents so as to perform active measurements on-demand, conforming to the OWAMP (One-Way Active Measurement) Protocol, and supported by a custom-made MIB.

In this context, the Network Management entity should undertake the additional task of invoking end-to-end active measurements and providing measurement results to the Network Cost Server.

The Network Cost Service itself receives raw network measurement data and assembles end-to-end costs, which are cached in a Network Repository for instant provision. The Repository should also host other network data, such as policies and topology, which assist in the cost estimation procedure. It must be noted Network Costs are calculated exploiting recently acquired measurement data already in the Repository; the invocation of a new in-network end-to-end measurement at each NC request would result in considerable overhead in the network and put the scalability of the proposed approach into question – in addition to raising security and stability issues.

The NCS exposes a public interface –commonly a SOAP-based or RESTful one- for providing costs on-demand, in e.g. XML or JSON format (For reference, the ALTO protocol uses JSON for more lightweight representation). Authentication and/or encryption mechanisms may also apply at this stage. Requests originating from applications should contain the addresses of the end-points between which the Network Cost is requested. For content/service-aware networks, which provide different treatment to various services (e.g. prioritise realtime streams against background traffic), Service Type information will also need to be explicitly signaled, since a given pair of endpoints will exhibit different costs for different services.

For maximum flexibility, it is proposed that the network cost itself is provided as a vector, than a single scalar value:

$$\vec{C}^{(a,b)} = \left(c_1^{(a,b)}, c_2^{(a,b)}, c_3^{(a,b)}, \ldots, c_n^{(a,b)} \right)$$

where $c_i^{(a,b)}$, as a real number, denotes a certain network metric between endpoints a and b. For example, c_1 can correspond to hop count, c_2 to one-way delay, c_3 to average packet loss, c_4 to jitter/delay variation etc.

In order to eventually derive a single scalar value corresponding to the application-specific network cost, each application should employ a weight vector W:

$$\bar{W} = (w_1, w_2, w_3, \dots w_n)$$

Each element of the weight vector corresponds to the sensitivity of the specific application against a specific network metric, i.e. how the application experience degrades as the metric increases. Keeping the aforementioned correspondence, w_1 denotes the sensitivity of the application to hop count, w_2 to one-way delay etc. Thus, the application-specific network cost, as a scalar value, can be expressed as the product of the two vectors:

$$NC^{(a,b)} = \bar{C}^{(a,b)} \cdot \bar{W} = c_1^{(a,b)} w_1 + c_2^{(a,b)} w_2 + \dots + c_n^{(a,b)} w_n$$

The above expression for simplicity assumes linear sensitivity to network conditions. This of course is not always true; for example, in UDP video streaming applications, it is known that image quality degrades exponentially with packet loss. In this case, more complex models could be employed. However, since the purpose of the network cost is to be a quick and rough estimate of the communication quality between two endpoints and not a realistic approximation to the actual Quality of Experience (QoE), as perceived by the user, the aforementioned linear model should be adequate in most cases.

IV. Preliminary Proof-of-Concept Implementation

Fig. 2 shows a preliminary proof-of-concept implementation of the proposed Network Cost Service framework, as it was carried out in the frame of the ICT-ALICANTE [4] project. Use cases to be tested mainly focus on p2p streaming (best peer selection) and distributed caching. The technologies and protocols which were used for producing and communicating the cost metrics are only indicative, and can be replaced by equivalent ones.

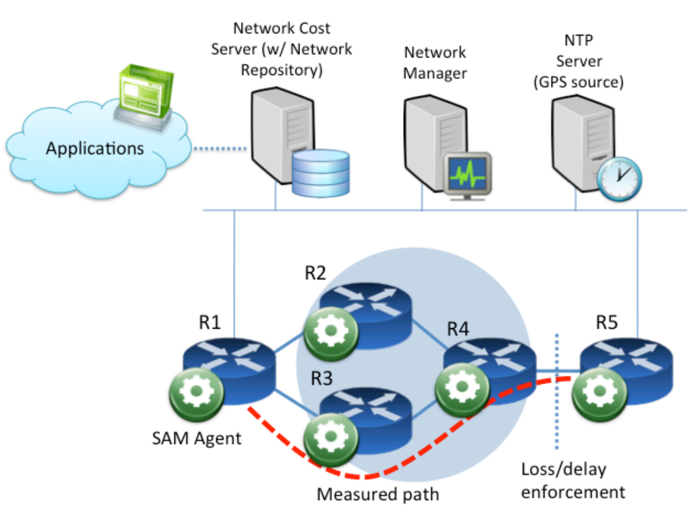

Fig.2. Proof-of-concept NCS implementation

The experimental network domain of the testbed consists of five Linux-based routers (R1-R5), which constitute a DiffServ/MPLS domain. Each router featured a specially customized SNMP agent (SNMP for Active Measurements – “SAM agent”) which supports the establishment of active measurement sessions conformant to the OWAMP protocol [5]. SAM agents utilize a custom-built MIB (SAM-MIB) and are configured via SNMP to periodically conduct active measurements, and communicate the measured metrics to the Network Manager also over SNMP. The SAM Agent has been released as open-source [6] and more technical details are to be found in [7].

The Network Manager configures the SAM agents to periodically conduct measurements to other peers in the network. The Network Cost Server retrieves the results and forms the end-to-end Network Costs to be communicated to applications via an exposed WebService. NCs are provided over the SOAP protocol in the form:

$$\bar{C}^{(a,b)} = \left(c_1^{(a,b)}, c_2^{(a,b)}, c_3^{(a,b)}, c_4^{(a,b)}, c_5^{(a,b)} \right)$$

where:

- $c_1^{(a,b)}$ is the number of hops traversed
- $c_2^{(a,b)}$ is the average one-way delay (OWD) experienced by the probe traffic packets (in μsec)
- $c_3^{(a,b)}$ is the average packet loss
- $c_4^{(a,b)}$ is the average jitter (in μsec)
- $c_5^{(a,b)}$ is the percentage of duplicate packets received

Accurate one-way delay (OWD) measurements [7] require tight synchronization between the sender and the receiver of the probe traffic. In our testbed, this is made possible thanks to a dedicated Stratum 1 NTP server, driven by a GPS clock and achieving μsec-order synchronization among the network elements.

Since the ALICANTE network is a content/service-aware one, Network Cost queries also include a Service Type field, which is automatically mapped to a specific traffic class to be used in the DiffServ network. In order to provide realistic measurements corresponding to the specific service, SAM agents mark the DSCP field accordingly in the probe traffic packets. Thus, the same pair of endpoints may yield different cost vectors for different service types.

Preliminary experimentation focused on the estimation of the Network Cost between R1 and R5 over the path shown in Fig.2. In the R4/R5 link, specific amount of loss and delay was enforced via a network emulator. Lab trials [8] show that the components of the framework, as designed and implemented, respond with sufficient accuracy to the network conditions, properly measuring the Network Cost elements (specifically one-way delay and packet loss).

V. Other Considerations

When it comes to a real-worlds implementation and deployment of a Network Cost Service, there are a number of considerations which must be taken into account:

A. Security and Network Operator Privacy

In addition to employing anti-DoS policies to avoid a denial-of-service attack to the NCS, it might be desirable that encryption and authentication mechanisms are employed in the NCS public interface, so as network costs are communicated only to authorized entities. Furthermore, it is true that batch, automated NC requests would allow the recipient to gradually form a thorough snapshot of the network status, even exploiting techniques such as network tomography. This could be a privacy breach for the Network Operator, since it would publicly expose sensitive network information. For this reason, certain restriction policies should apply even to authorized NC requests.

B. Inter-domain operation

The aforementioned architectural description, along with the proof-of-concept implementation correspond to NC provision within a single administrative network domain. For multi-domain operation, the most realistic solution involves the cooperation between NC Servers of adjacent domains. Thus, a NC estimation to a peer belonging to a foreign domain should result from the communication of the NC Servers from all intermediate domains involved in the network path. This communication should probably take place in a cascading form.

C. Access network issues

The discussion which preceded included NC estimation between core, edge and access network elements (routers). It is however essential to include the access network in the measurement procedure, since bottlenecks often occur in the access part. However, performing end-to-end measurements among thousands or millions of user terminals (mostly CPE/home gateways) would be unscalable. For this purpose, it could be proposed that the link between the CPE and the access router is evaluated separately, via a WAN management protocol such as [9] and then combined with in-network path measurements so as to form the end-to-end cost. This of course complicates the functionality of the NC Service, but could yield a feasible solution.

All the aforementioned issues constitute interesting fields for further research.

VI. Conclusion

Within the FI scene, which is shifting from host-centric to content-centric communication, a network cost service seems a promising approach, offering applications the capability to automatically self-optimise according to network conditions. We presented a generic yet feasible framework for providing a Network Cost Service and described a proof-of-concept implementation, designed and deployed in the frame of the ICT ALICANTE project. Preliminary results show that Network Costs can be provided with relatively minimal overhead in the network operation, and exhibit remarkable accuracy. Next research tasks should include experimenting in wider-scale networks and also in simulation platforms so as to assess the scalability of the mechanism, optimizing the network cost representation (also using nonlinear models) and also dealing with security, privacy, inter-domain and access network issues, as mentioned in the previous section.

References

[1] European Commission, FIArch group, "Future Internet Design Principles", January 2012, http://ec.europa.eu/information_society/activities/foi/docs/fiarchdesignprinciples-v1.pdf

[2] J. Seedorf, E. Burger, "Application-Layer Traffic Optimization (ALTO) Problem Statement", RFC 5693, October 2009, http://tools.ietf.org/html/rfc5693

[3] R. Alimi, R. Penno, Y. Yang, "ALTO Protocol", Internet Draft (Work in Progress), http://tools.ietf.org/html/draft-ietf-alto-protocol-12, July 2012

[4] FP7/ICT ALICANTE: Media Ecosystem Deployment through Ubiquitous Content-Aware Network Environments, http://www.ict-alicante.eu/public/

[5] S. Shalunov et al, "A One-way Active Measurement Protocol", RFC 4656, September 2006, http://tools.ietf.org/html/rfc4656

[6] SNMP for Active Measurements (SAM) Framework, http://www.medianetlab.gr/opensource/

[7] L. de Vito, S. Rapuano, L. Tomaciello, "One-Way Delay Measurement: State of the Art", IEEE Trans. On Instrumentation and Measurement, 57(12), December 2008, pp. 2742-2750

[8] G. Gardikis, K. Sarsembagieva, G. Xilouris, A. Kourtis, "An SNMP Agent for Active In-Network Measurements", in Proc. ICUMT 2012, October 3-5, 2012, St. Petersburg, Russia (to appear)

[9] Broadband Forum, TR-069 CPE WAN Management Protocol, May 2004, http://www.broadband-forum.org/technical/download/TR-069.pdf

Dr. Georgios Gardikis received his Diploma in Electrical and Computer Engineering from the National Technical University of Athens in 2000 and his PhD from the same University in 2004. His expertise lies in the fields of digital terrestrial and satellite broadcasting, distribution networks for multimedia service provisioning, Quality of Experience (QoE) assessment and novel mechanisms for application/network coupling. Since 2004, Dr. Gardikis has been collaborating with the National Center of Scientific Research "Demokritos" and the Centre for Technological research of Crete in the frame of national and EU-funded R&D projects, including FP6/IST IMOSAN, in which he served as Technical Coordinator. He has also collaborated with the University of the Aegean, conducting a national research project regarding experimental DVB-H deployment, and with the Ministry of Transport and Communications as consultant on digital broadcasting issues. At present, he is Associate Researcher in Institute of Informatics & Telecommunications in the NCSR "Demokritos" and Visiting Assistant Professor at Dept. of Applied Informatics and Multimedia, TEI of Crete. He is member of the IEEE/Communications and Broadcast Technology Societies, member of W3C/"Web and TV" Interest Group and member of the Technical Chamber of Greece since 2001. He has more than 40 publications in international journals and refereed conferences/workshops.

George Xilouris received his BS in Physics from University of Ioannina, Greece in 1999 and his MSc in Automation Control Systems and Robotics from National Technical University of Athens in 2001. He joined the Institute of Informatics and Telecommunications of NCSR "Demokritos" in 2000 and since then, he has participated in numerous EU-funded projects, as well as in National funded projects. His current research interests include, Next Generation Networks, Network Management, Network Performance, Cloud Computing Network Security, Digital Television (DVB), Satellite Communications and IT services.