Title: Vehicles as a Mobile Cloud: Leveraging Mobility for Content Storage and Dissemination

Advisors:

Thrasyvoulos Spyropoulos
Mobile Communications, Eurecom
Thrasyvoulos.Spyropoulos@eurecom.fr
http://www.eurecom.fr/~spyropou/

Chadi Barakat
Diana team, Inria
Chadi.Barakat@inria.fr
http://www-sop.inria.fr/members/Chadi.Barakat/

Description:

The volume of content (video, audio, pictures, news) in the Internet is increasing at an unprecedented rate. As a result, we are moving from an Internet principally meant to connect us to well defined hosts to an Internet where the main purpose is to access content and share it without caring much about its location [1]. In addition, Internet users want to access this content from different locations and with different devices, like smartphones, pads, and notebooks. This new way of using the Internet has pushed towards centralized solutions for data storage and sharing in the form of data centers, clouds or social networks. However, these solutions remain proprietary, limiting the volume of resources one can freely obtain, or raising privacy issues around the posted content and the way end users interact with it.

In parallel to these developments, the increasing access of content with mobile devices has stressed existing cellular networks beyond their capacity [2]. Forecasts suggest that the exponential growth of mobile data demand will quickly surpass the additional capacity offered by 4G technologies [3]. As a result, cellular network operators are looking to offload as much data traffic as possible over femtocells, WiFi hotspots, and even device-to-device communications [4]. Furthermore, to take advantage of the content-centric shift in user demand, researchers and operators are investigating the storage of popular content as close to the mobile user as possible, to avoid wasting resources and decrease access delays (in a spirit similar to traditional CDNs) [5].

Motivated by the above concerns, in this proposal we plan to explore the possibility to build an infrastructure for content storage and sharing around vehicles (cars, buses, etc). With their mobility and their capability to carry on storage and communication devices without serious power limitations, vehicles can form an interesting networking infrastructure able to store content, relay it between each other, and deliver it to end users upon request. This content might be both end-user generated (e.g. pictures and videos taken with a mobile device) or downloaded at some point from the Internet. Our goal is to *bring the cloud closer to the user, optimize resource usage to access cloud content, and provide the means for an open, non-profit oriented data sharing and dissemination.* Yet, we stress that this architecture is not meant to bypass altogether the existing wireless and cloud infrastructure, but rather to extend it and/or provide a more cost-efficient alternative. Traditional infrastructure could still act as a fallback solution, when content cannot be found locally or fast enough. We believe such an architecture offers advantages along the following directions:

- ***Data offloading***: Popular and location-dependent content could be easily retrieved from nearby nodes of the mobile cloud, without the need to congest valuable cellular resources transferring the same content multiple times, from end-servers deep in the core network.
- ***Coverage extension*****:** Vehicles, through their mobility, can extend “cheap” coverage beyond the usually sporadic WiFi hotspot connectivity, or in locations with weak or no 4G signal.
- ***Performance improvement*****:** In a recent survey, the reduction of distance between the base station and the terminal has been identified as the main driver of performance improvement for cellular networks [6]. Connecting a wireless device to a nearby vehicle can further reduce this distance and potentially lead to important performance gains.
- ***Reduced dependence on proprietary solution*****s:** Finally, such an architecture provides to end users cloud functionality without the need to pass by other proprietary solutions as the social networks and the clouds of Google, Apple and DropBox.

Communication in the mobile cloud we envision will be built around two main recent paradigms:

Information-Centric Networks, is a departure from traditional communication architectures, to make content rather than endpoints the central element. Instead, end users publish their content with their mobiles, the content name is then made available to all users of the infrastructure via a search engine associated to the infrastructure (either directly or indirectly via a dedicated web site for example). Then, other end users of the infrastructure, if authorized, can request this content to passing vehicles and count on their collaboration to look for it and bring it back to them.

Device-to-device (or hop-by-hop) *opportunistic* communications, are a natural match to the above content-centric approach and the increased mobility of vehicles involved. Here, communication may occur only when two nodes are in proximity, between (i) end devices and vehicles in the mobile cloud, for querying the cloud or for retrieving content, and (ii) between vehicles in the mobile cloud, for updating or moving the stored (cached) content of different vehicles and make the content “float” near the location(s) it will be mostly needed.

Nevertheless, in order to materialize the many potential benefits from the envisioned architecture, a number of significant challenges must be faced, stemming from the mobility of vehicles and the non-centralized nature of the system. Below we summarize these challenges and sketch our intended approach in dealing with each one:

- **Request and content dissemination:** When an end device does not find the content it is looking for in the mobile cloud node it happens to be connected at, the request must be efficiently disseminated over the cloud, and the content sent back from its closest storage point. *Approach*: *we will devise algorithms that can predict the location(s) of future demand, and ensure content “floats” (is “cached”) in parts of the mobile cloud near these areas. This location-based component is novel with respect to traditional architectures. Furthermore, we will leverage efficient opportunistic anycast and multicast algorithms to disseminate content and requests in case of “misses”.*

- **Storage management**: The volume of available content will easily exceed the storage capacity in each vehicle. Hence, it is important that vehicles participating in the mobile cloud are equipped with algorithms to decide on what content to keep and what to discard. *Approach: we will formulate this as a resource allocation problem; the optimal solution* *for this problem will be a subtle interplay between the popularity of different contents, the mobility pattern of involved vehicles, and the local or global metric(s) the mobile cloud is optimizing. In practice, we will identify per content utilities related to the optimal solution, that will allow nodes to rank different contents and take decisions in a local, distributed manner.*

- **Scheduling of content:** When two vehicles come within transmission range, their encounter will not last long. As a result, they might not be able to exchange all the content that our caching and dissemination algorithms might otherwise require. It is thus necessary for nodes to start exchanging the most useful content for the community, before the connection is lost. *Approach: In a manner similar to the storage management problem, we will use the optimal per content utilities to decide which content to transmit first. Furthermore, to ensure the available contact duration is not wasted on metadata exchange (each node will have a very long list of available content and descriptions), we will explore the use of Bloom filters and other data compression/reduction techniques.*

- **Distributed estimation of demand and availability:** The infrastructure will be used to store different kind of content, some more popular than others. Furthermore, the popularity of different contents might change over time. Finally, new content will be inserted in the network and old content might be removed. As a result, appropriate algorithms will be needed to estimate and maintain, in a distributed manner, the global state of different contents, namely the approximate demand for each content and the availability in the mobile cloud. These two parameters are key inputs to the optimal management of cloud resources. *Approach: Nodes will sample demand and content availability in a lightweight and distributed way. Appropriate unbiased estimators will be sought for, and be used to approximate the optimal allocation of storage and bandwidth resources, in a manner that will minimize the expected long-term discrepancies between the two.*

This research proposal aims at addressing the above challenges and proposing algorithms and protocols for the best collaboration of vehicles (private or public) towards the deployment of a mobile cloud infrastructure. The topic selected relates to the LABEX domains of “data-centric networking” and “heterogeneity and efficiency”, with the application domain targeted being that of Intelligent Transportation Systems. We will depart from several years of collaboration around end-to-end communications in Delay Tolerant Networks and leverage our understanding of the effect of mobility patterns (key element in this project as well) on optimal resource allocation policies. However, the abundant resources of vehicles (compared to phones or sensors, usually the subject of DTN research), the intricacies of their mobility patterns, and the content-centric nature of this novel architecture will require a radical rethinking of the problem and solutions. Finally, we are well aware that such an architecture might require appropriate security and incentive mechanisms. While we do not intend to provide new solutions for these, we will stay informed and leverage the state-of-the-art in these directions, where necessary, to further solidify our architecture.

The research will start by a careful high-level design of the system identifying the main functions and building blocks. We will then formulate the key problem(s) assuming all parameters of interest are known. This will allows to obtain ideal optimal solutions, that can be used as a baseline for comparison, but also to perform a feasibility study, and understand the limitations of such an architecture. We will then move into proposing practical and distributed algorithms for handling content within this mobile infrastructure with the aim to statistically approximate the optimal policies. Along this process, we will pay particular attention to the validation of all algorithms, using realistic data as much as possible (e.g. mobility traces, traces of content demand, etc.) and experimenting the algorithms within well-known frameworks for content-centric networking and opportunistic networking as the CCNx framework or the DTN2 one.

References:

[1] T. Koponen, M. Chawla, B.-G. Chun, A. Ermolinskiy, K. H. Kim, S. Shenker, and I. Stoica, “A Data-Oriented (and Beyond) Network Architecture,” In Proc. of ACM SIGCOMM, 2007.

[2] Cisco Visual Networking Index, White Paper, Feb. 2012, available online at http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-520862.pdf, 2012.

[3] “Growing Data Demands Are Trouble For Verizon, LTE Capacity Nearing Limits” http://www.talkandroid.com/97125-growing-data-demands-are-trouble-for-verizon-lte-capacity-nearing-limits/

[4] “Mobile data offloading through WiFi,” 2010. Proxim Wireless white paper.

[5] N. Golrezaei, K. Shanmugam, A.G. Dimakis, A.F. Molisch and G. Caire, “FemtoCaching: Wireless Video Content Delivery through Distributed Caching Helpers,” in Proc. of IEEE INFOCOM 2012.

[6] V. Chandrasekhar, J. G. Andrews and A. Gatherer, "Femtocell Networks: a Survey", IEEE Communications Magazine, Vol. 46, No. 9, pp. 59-67, September 2008

Selected past publications:

[7] Amir Krifa, Chadi Barakat, Thrasyvoulos Spyropoulos, “Message Drop and Scheduling in DTNs: Theory and Practice”, in IEEE Transactions on Mobile Computing, vol. 11, no. 9, September 2012.

[8] Amir Krifa, Chadi Barakat, Thrasyvoulos Spyropoulos, “MobiTrade: Trading Content in Disruption Tolerant Networks”, in proceedings of ACM Mobicom Workshop on Challenged Networks (CHANTS), Las Vegas, September 2011.

[9] Amir Krifa, Chadi Barakat, Thrasyvoulos Spyropoulos, “Optimal Buffer Management Policies for Delay Tolerant Networks”, in proceedings IEEE SECON 2008. Best paper award.