Sporadic Cloud Computing over a Virtualization Layer
A New Paradigm to Support Mobile Multi-hop Ad-hoc Networks

Esteban F. Ordóñez-Morales[1], Yolanda Blanco-Fernández[2] and Martín López-Nores[2]
[1]*Grupo de Investigación e Innovación en Ingenierías, Universidad Politécnica Salesiana, Cuenca, Ecuador*
[2]*AtlantTIC Research Center, Department of Telematics Engineering, University of Vigo, Vigo, Spain*

In our doctoral proposal we deploy Sporadic Ad-hoc Networks (SANs) over the devices of a group of always-on users who happen to meet in a place. The goal is to develop tailor-made services that exploit the possible similarities among the preferences of the users and the technological capabilities of their terminals to establish direct and hop-by-hop ad-hoc communications. In order to overcome the intrinsic limitations of mobile devices, we explore the new concept of Sporadic Cloud Computing (SCC) that is aimed at providing each terminal with additional resources by exploiting the (computational, networking, storing...) capabilities of the rest of devices connected to the SAN. In order to abstract the complexity stemmed from the mobility scenarios, SCC works with an enhanced Virtualization Layer that deals with a few static virtual nodes instead of a higher number of mobile real nodes. This allows to turn our SANs into reliable and stable communication environments to promote interactions among potentially like-minded strangers in a great diversity of mobility scenarios, involving both pedestrians and cars in vehicular environments.

1 RESEARCH PROBLEM

After the irruption of the Web 2.0 and the smartphone revolution, most of the on-move users carry with them handheld devices for actively interacting daily with their friends/followers/followees/... in the context of the virtual world of the Internet. Sociologist have already advised about the negative effects derived from some of these behaviours, which might lead the users to immersing themselves in a virtual communication burble with their contacts, by giving up interacting face-to-face with nearby individuals (Kuss and Griffiths, 2011). In the same line, other experts have analyzed the consequences of the so-called FOMO (Fear Of Missing Out) effect which denotes the fear of users of losing events, news and situations that happen in their social networks, thus not taking an eye off their devices (Przybylski et al., 2013). In order to fight these situations, we are working in the development of the SPORANGIUM (SPORAdic networks in the Next-Generation Information services for Users on the Move) platform which deployes Sporadic Ad-hoc Networks (SANs) among always-on users who happen to be in a place, by establishing multi-hop ad-hoc connections over their respective mobile terminals. The goal is to promote more direct interactions among strangers who happen to meet in spaces like cinemas, stadiums, museums, concert halls, etc., who might have potentially common interests (cinema, sport, art, music, etc.) that would be convenient to explore.

For that purpose, SPORANGIUM must orchestrate activities and tailor-made services that bring together the particular context of the users, their potentially common preferences and the capabilities of their devices for establishing ad-hoc connections. However, these services far exceed traditional mobile devices capabilities, which suffer from computational limitations, as well as battery restrictions and processing time. To face this situation, the new Mobile Cloud Computing (MCC) paradigm has arisen that takes inspiration from the well-known Cloud Computing (CC), which is based on delivering computing as a service whereby shared resources, software and information are provided as a utility over a network (typically the Internet). In MCC the goal is to enable to process a large amount of data on demand anytime from anywhere, so that mobile devices connect to the Internet to use an environment that integrates diverse platforms and technologies (Dinh et al., 2013). Specifically, MCC promotes to move the computing power and data storage away from mobile devices and into the cloud, bringing multiple service models (IaaS, PaaS, SaaS...) and mobile computing to a wide range of on-move always-on users.

Recently, the MCC paradigm has been exported to vehicular communication environments where some researchers have proposed the so-called Vehicular Cloud Computing (VCC). The goal here is to exploit

F. Ordóñez-Morales E., Blanco-Fernández Y. and López-Nores M..
Sporadic Cloud Computing over a Virtualization Layer - A New Paradigm to Support Mobile Multi-hop Ad-hoc Networks.

both the physical data center units that are in charge of performing the data computation and storage (like in MCC) and the on-board resources of the own vehicles (Olariu et al., 2013). In our proposal of PhD work, we want to extend the VCC paradigm to support mobility scenarios beyond the vehicular environment, by involving both pedestrians and vehicles on the road. Specifically, we explore a new paradigm – named *Sporadic Cloud Computing* (SCC)– aimed to allow the users' devices to exploit both the (computing, storing, networking, sensing...) resources available in the rest of terminals connected to the SANs, and those provided from external data centers. In the deployment of our SCC paradigm in diverse mobile communication environments, one of the main research challenges has to do with the high mobility of the nodes connected to the ad-hoc network (e.g. cars in a vehicular ad-hoc network), and therefore, with its frequenlty changing topology. This causes that communication fails often as these nodes move fast and are out of the range of the ad-hoc network, which also hampers the routing tasks when forwarding information (Gerla, 2012).

2 OUTLINE OF OBJECTIVES

Our doctoral proposal is aimed at designing, developing and validating the mechanisms necessary to:

1. turn our SANs networks into reliable and stable communication environments with good performance in terms of overhead, packet delivery ratios and scalability, covering vehicular, pedestrian and mixed environments, and
2. deploy enhanced "X"aaS service models (e.g. CaaS, NaaS, STaaS, SEaaS...) through our SCC paradigm, so that the devices connected to the SAN can collaborate and share their respective Computing, Networking, STоring and SEnsing capabilities in the deployment of advanced communication services.

As introduced before, the main research challenges derived from both objectives has to do with (i) the frequent topology changes happened in certain communication environments (e.g. in a vehicular adhoc network due to the fast movements of the cars), and (ii) the fact that the capabilities available in the SANs vary on the time, as the location of the users' devices change.

To deal with the high mobility of the nodes connected to our SANs, our proposal takes advantage of the improvements proposed in the realm of Mobile Ad-hoc Networks (MANETs), which have been envisaged to face the problems derived from (i) the wireless transmission mediums, (ii) the high variability of the network topology due to unpredictable nodes' movements, and (iii) the existence of severe restrictions in terms of processing capabilities, memory and battery consumption. In particular, we start from the work presented in (Dolev et al., 2004) where the authors described a *virtualization layer* named VNLayer (*Virtual Node Layer*). Specifically, the VNLayer is a cluster-based approach where the mobile nodes collaboratively create an infrastructure of static *virtual nodes* to ease the routing problem and the maintenance of persistent state information in the area covered by an ad-hoc wireless network of mobile devices (as our SANs), notwithstanding the mobility of the (real) physical nodes. Actually, the VNLayer resides between the link layer and the Internet Layer, so that the virtual nodes can be addressed as if they were static server devices. This helps to mask the uncertainty that arises from the MANETs' varying topology and from the fact that the physical devices can fail unpredictably. Consequently, it is easier for developers to work at the nodes' upper layers, since they can deploy applications on mobile devices and virtual servers with greater ease and efficiency. Besides, virtualisation creates a level of hierarchy in the otherwise flat MANETs, which brings in opportunities to re-design MANET protocols to operate more efficiently and reliably.

Since the virtualization layer by Dolev et al. has been developed to handle communications in MANETs, the first objective of our doctoral proposal consists of **extending and adapting the working of the VNLayer to the restrictions and peculiarites of more demanding mobility scenarios**, including, for instance, communication environments where pedestrians and occupants of vehicles are involved. To this aim, we need to envisage refinements in the VNLayer (resulting in our VNLayer+) to take into account, for instance, the comparatively faster movements of vehicles, the freedom of movement of the pedestrians, as well as the fact that these nodes are not subject to the strict energy, space and computing capabilities restrictions of MANETs. These restrictions must be considered to turn our SANs into reliable communication environments, covering a wide diversity of application scenarios beyond the generic MANETs explored in Dolev et al.'s approach.

Our SCC paradigm must develop **transport-layer coordination mechanisms among the devices that are connected to the SANs in order to enable an efficient sharing and allocation of their available resources by working over the virtualization layer**, which is the second objetive of the doctoral proposal. This fact causes that, differently from the traditional approaches envisaged in CC, MCC and VCC, our “X”aaS service models need to deal with the (static) virtual nodes of the VNLayer+, which are emulated/supported by the devices of the users on the move. Specifically, while the SAN is established among the users’ terminals, the messages to request resources from the ad-hoc network (or to advertise resources that are left to other devices’ disposal) are managed by virtual nodes. The tandem SCC-VNLayer+ contributes to (i) fight/alleviate the communication errors and data loss noticeable in mobile ad-hoc networks (Dinh et al., 2013), and (ii) to orchestrate advanced applications to improve the experience of the users, by taking advantage of the reliable data exchange over the SAN (thanks to the VNLayer+) and the availability of additional resources in each terminal (thanks to the service models of SCC).

The possible applications to be deployed in the realm of our SANs cover a wide spectrum, ranging from the orchestation of activities bounded to an event where a group of like-minded users happen to meet (e.g. in a museum, theater or stadium), to the provision of both improved applications for intervehicular communication (e.g. optimization of traffic flows, chats among drivers, proactive organization of ride-sharing opportunities or selective distribution of personalized advertising in nearby places), and refinements in the context of the *smart cities* through the planification of people mobility and urban games, among others.

3 STATE OF THE ART

In this section, we review related works in the two main research fields of our doctoral proposal: the use of virtualization in MANETs and the exploitation of resource sharing among devices in mobile communication environments.

3.1 Virtualization in Mobile Ad-hoc Networks

The VNLayer was presented in (Dolev et al., 2004) as a set of procedures to turn ad-hoc networks of mobile devices into more predictable environments for communications. The main idea is to engage the mobile physical nodes (PNs) in collaboration to emulate virtual nodes (VNs) that remain in known grid locations, as shown in Figure 1 (where back circles and white squares denote PNs and VNs, respectively).

The VNLayer divides the geographical area of an ad-hoc network into square regions, whose size is chosen so that every PN in a region can reliably send and receive data from every other physical node in that region and neighboring ones. Each VN (one per region) is emulated by the PNs located in the corresponding region, so that when all the physical nodes leave this region, the virtual node stops to work. In each region, one PN is chosen as the *leader* in the region and becomes the primary responsible for packet reception, buffering and forwarding. Meanwhile, a subset of non-leader nodes are designated as *backups* to maintain information consistent with the leader’s version (specifically, replicas of the virtualization-related state information and the routing tables tackled by the routing protocols working on the virtualization layer). This way, the VNs can maintain persistent state and be fault tolerant even when individual PNs fail or leave the region.

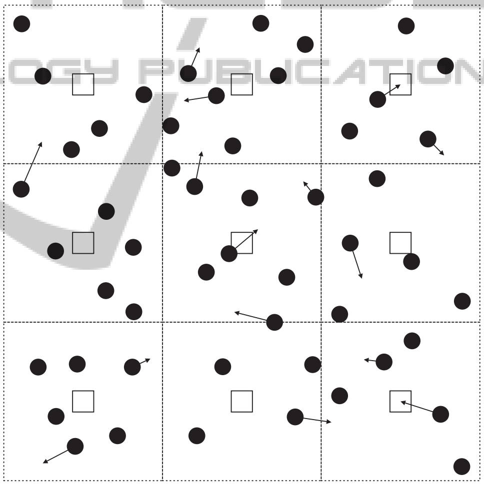

Figure 1: Static virtual nodes (white squares) overlaying the mobile nodes of a MANET (black circles).

An exhaustive analysis of the VNLayer allows to detect certain sources of inefficiency in its functioning, which are mainly related to:

- *The Procedure used to Identify New Backup Nodes*: The approach adopted by Dolev et al. in order to designate backup nodes among non-leaders is a probabilistic one and it is driven by a Coin Tosser Function, according to which the greater the number of nodes in the region, the lower the probability that a PN will choose as backup. This avoids having many backups in dense MANETs, and thus reduces the overhead due to state synchronisations, i.e. to the exchange of messages aimed at ensuring that the replicas of the state information from the upper layers (e.g. routing tables) kept in the backup nodes are consistent with the leader's version.

- *The Procedure adopted in the Leader Election*: This process suffers from two main problems: (i) it involves a great number of messages to be exchanged –which contributes to increase the duration of the virtual nodes' downtimes–, and (ii) the selection process does not prioritize the backup nodes that have an updated version of the outgoing leader node (which would be the best candidates to become a new leader). In particular, this procedure could designate as new leader either a node that was not acting as a backup or a non-synchronised backup node, even in cases that there were synchronised backups in the region. This causes that the state information from the upper layers would be lost unnecessarily and new synchronisations would be triggered immediately, thus increasing the overhead.
- *The Procedures adopted Once a Region Becomes Empty*: The VNLayer does not preserve information about the states of the VNs corresponding to regions that become empty after the withdrawal of all the PNs located in them. This degrades the performance of the virtualization layer, thus causing unnecessary delays in the recovery of the VNs.

Besides the above limitations (which have been identified in MANET scenarios), the VNLayer requires additional refinements to improve its performance in more restrictive and demanding mobile communication environments, such as the pedestrian/vehicular/ mixed scenarios we want to explore in our doctoral proposal.

3.2 Resource Sharing among Mobile Devices

The idea of taking advantage of the resources available in the handheld devices that are located around an on-move user is not new at all. Specifically, the capabilities that have gained more momentum are those related to the networking resources. In this regard, the so-called *spontaneous networking* has arisen in the last years, where wireless mobile nodes opportunistically exploit multi-hop ad-hoc paths toward peers to share content and available resources in an impromptu way. The idea is to take benefit from the available bandwidth in many handheld devices (which is often underutilized) to be shared with other peers in current vicinity, thus better exploiting the increasing availability of computing/memory/bandwidth-related capabilities at portable wireless terminals. In this line, we found the approach proposed in (Bellavista and Giannelli, 2010) where a middleware named RAMP (*Real Ad-hoc Multi-hop Peer-to-peer*) is described. In particular, RAMP combines network-layer solutions and application-layer approaches to support Internet connectivity sharing in spontaneous networks. Specifically, RAMP creates multi-hop paths toward border nodes (i.e., nodes directly connected to the traditional Internet and offering part of their underutilized connectivity to nearby peers), so that each node can use the path currently deemed as the most suitable, e.g., because it provides largest bandwidth or requires lowest power comsumption. This approach leads to a significant routing overhead when exploiting different multi-hop heterogenous paths traversing the same node.

At the application layer we found other approaches that go beyond the solutions proposed in RAMP, which have been designed for vehicular ad-hoc networks taking inspiration from BitTorrent-style P2P file sharing systems (Nandan et al., 2005; Lee et al., 2007; Chen and Chan, 2009; Lee et al., 2006; Eriksson et al., 2008). The goal is not only to exploit the connectivity of one terminal from the rest of devices, but to collaboratively download different chunks of the same content during periods of connectedness. All these application-layer protocols are aimed at enabling (collaborative) downloads of contents that typically are appealing to all (or most of) the vehicles connected to the VANET. The contributions that we are pursuing in SCC are located in a lower layer, where the goal is to aggregate the connections of several nodes in a transparent way, without conditioning besides the usage of the downloaded contents by those nodes (covering, e.g., scenarios where the accessed information is useful for just one node in the SAN). To this aim, we must envisage tranport-layer solutions that deal with multiple connections and multi-hop communications in diverse mobile ad-hoc networks, by working on the top of our enhanced virtualization layer, which, to the best of our knowledge, is approach completely novel in literature.

Beyond the networking capabilities, in the vehicular environment it is also possible to find new service models that allow vehicles to share to each other on-board storage facilities (STaaS: Storage as a Service), computing power (CaaS: Computing as a Service), services (about traffic information, driver safety or weather and road conditions) that are assembled from the information collected by other vehicles (COaaS: Collaboration as a Service), and advanced functionalities related to provision of entertainment as a service on the road (ENaaS: ENtertainment as a Service) or taking of photos and recording of videos in particular places and at specific times (PicWaaS: Pictures on a Wheel as a Service), as described in (Arif et al., 2012). These services can be deployed over diverse vehicular clouds, ranging from static clouds (which aggregate the capabilities of parked cars) and semi-static clouds (involving vehicles stopped for a moment because of a traffic jam) to mobile clouds (the most common option where a large amount of vehicles travel on the road). The most sophisticated approaches have been designed in static and semi-static clouds, while the challenges derived from the frequently changing topologies of mobile clouds have not received the same attention (Gerla et al., 2014). The goal of our proposal is to handle the mobility of the nodes connected to our SANs (both pedestrians and vehicles), by exploiting the virtualization refinements and the mechanisms envisaged in the SCC to face the communication errors and data loss noticeable in highly dynamic ad-hoc communication environments (Arif et al., 2012).

4 METHODOLOGY AND STAGE OF THE RESEARCH

After an in-depth review of the state-of-the-art, the first step to tackle our research problem has been the development of a simulator (whose high-level design is sketched in Section 4.1), which is aimed at validating the procedures of the VNLayer+ (Section 4.2) and the “X”aaS service models of the SCC paradigm (Section 4.3). While the simulator has been totally implemented, our ongoing work is focused on the foundations of the VNLayer+ and the specific mechanisms of some “X”aaS models of the SCC.

4.1 A SAN Simulator

Covering vehicular, pedestrian and mixed environments requires that our simulator (i) models the mobility requiremens of each application scenario, and (ii) deals with the communications among the mobile nodes. To this aim, we have revised diverse pedestrian and vehicular mobility models defined in literature (Sharma and Singh, 2013), with the goal of selecting the ones that represent realistically the behaviours of the moving nodes that are connected to our SANs. Regarding the pedestrians, as depicted in Figure 2, we have chosen three different models to generate diverse types of mobility traces, referred both to individuals and groups.

- The *Random Walk Mobility Model* allows the nodes to move randomly without restrictions (e.g. a pedestrian walking a street in a city), so that the destination, speed and direction are all chosen independently of other nodes.
- The *Nomadic Community Mobility Model* considers groups of nodes that collectively move from one point to another. This model is especially appropriate for scenarios where a group of pedestrians move together, but each individual could also roam around a particular location individually (e.g. a group of tourists who visit together the historical centre of a city). By adjusting the corresponding parameters, it is possible to control how far each node can roam from each reference point, thus resulting into very realistic movements.
- Finally, we adopt the *Reference Point Group Mobility Model* where each group is composed of a number of members and one leader, so that the movements of the leader determine the mobility behaviour of the entire set (e.g in a mobility scenario where a group of students visit a museum, being guided by an expert).

As depicted in Figure 2, the above mobility models have been integrated via the existing simulation tool MobiSim[1], whose modular architecture allows to easily add extra models and trace formats. Regarding the generation of vehicular mobility traces, we have resorted to SUMO[2], due to the possibility of adding new mobility models and submitting realistic (vehicular) mobility traces in NS-2 format.

Also, we have decided to adopt NS-3 because this simulator greatly improves NS-2 in terms of efficiency, memory management and kernel architecture, besides making easier the integration of third-group software and the definition of new mobility models by using C++. Certainly, the models implemented in NS-3 are too simple in order to fulfill a wide diversity of mobility requirements. However, our simulator overcomes this limitation thanks to the (pedestrian and vehicular) mobility behaviours modeled by the external simulators MobiSim and SUMO. As these behaviours are modeled as NS-2 traces, NS-3 uses a NS2MobilityHelper module to convert them to NS-3 mobility events. As seen in Figure 2, NS-3 supports protocols such as UDP, TCP, IP and multiple routing protocols for mobile ad-hoc networks. Considering our virtualization mechanisms requires to include two additional modules aimed at implementing the virtual layer level (VNLayer+) and a virtualized routing protocol grounded on it (VNRouting), thus greatly improving the performance of the ad-hoc network. Lastly, the lowest level hosts diverse versions of the IEEE 802.11 protocol (such as IEEE 802.11p specifically developed for vehicular networks).

---

[1] http://www.masoudmoshref.com/old/myworks/documentpages

[2] http://sumo.sourceforge.net/

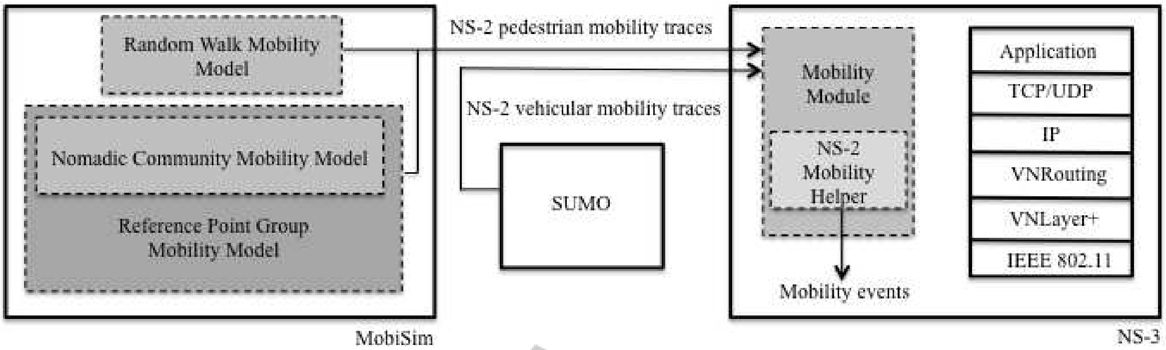

Figure 2: High-level design of our SAN simulator.

In conclusion, the exploitation of synergies among MobiSim, SUMO and NS-3 makes it possible to easily define multiple and diverse mobility scenarios where we will explore the potential of the SANs contributed in our proposal.

4.2 The VNLayer+

As depicted in Figure 3, the VNLayer+ divides the geographical area of the SAN into regions (denoted as cubes), so that a virtual node is located in each region. Analogously to what we commented for the VNLayer by Dolev et al., each virtual node is supported by several moving nodes (both pedestrians and vehicles in our scenarios), so that the virtual node stops to work after all the physical nodes leave the region. Also, there exist in each region leader and backup nodes in order to maintain replicas of the virtualization-related state information and the routing tables tackled by the routing protocols working on the VNLayer+.

We have focused on the envisage of procedures aimed to face the inefficiency sources identified in the traditional VNLayer (recall Section 3.1). Next, we sketch the ideas considered in our refinements, whose details can be found in (Bravo-Torres et al., 2015):

- First, we have developed a new leader election procedure to prevent from the slow reaction of the VNLayer to leader withdrawls, which impinged heavily on the communications in scenarios of high mobility, since the VNs were down during a non-negligible portion of the average time that the vehicles would remain in the respective region. Briefly, the leader election procedure is driven by different types of events (message receipts, timeouts and region changes), which take each PN to multiple states. Our approach consists of removing some of these states and reorganizing the transitions between the remaining ones in order to speed up the discovery of the new leader.

  Besides, we are also interested in sophisticating the election of VN leaders so that the role is dynamically transferred to the physical node that is most likely to remain longest within the corresponding region (as inferred from information coming from either the link layer or the applications layer).

- The backup designation procedure proposed in the VNLayer needs improvements too. In particular, our goal is to ensure that the number of backup nodes in a region stays, whenever possible, within a given minimum (to guarantee the resilience of the virtual nodes) and a given maximum (to avoid excessive synchronisation overhead). To this aim, our idea is that the leader reports the number of backups in the region, so that other non-leaders can learn whether they should offer themselves to further support the VN. This way, becoming a backup is no longer a fortuitous and autonomous decision as in Dolev et al.'s approach, but rather an informed and supportive one.

- Last, we have also developed procedures to improve the management of empty regions designed in the VNLayer. In this regard, our approach is based on defining a new table (named *B_table*) whose entries contain the physical addresses of the backup nodes along with its state (synchronised or not). Specifically, *B_table* replicas are stored in all the nodes of a region, which is the key to avoid losing state information from the upper layers when a newcomer assumes leadership shortly after the previous leader has left. In this scenario, the upstart (the newcomer) does not have the state information of the virtual VN of the region, but the synchronised backup do. Combining the *B_table* and the information from the backup node, the upstart can start operating just as well as the former leader in a very short time. These mechanisms are complemented with other procedures aimed to keep the structure of our SANs stable, thus enabling that the communica-

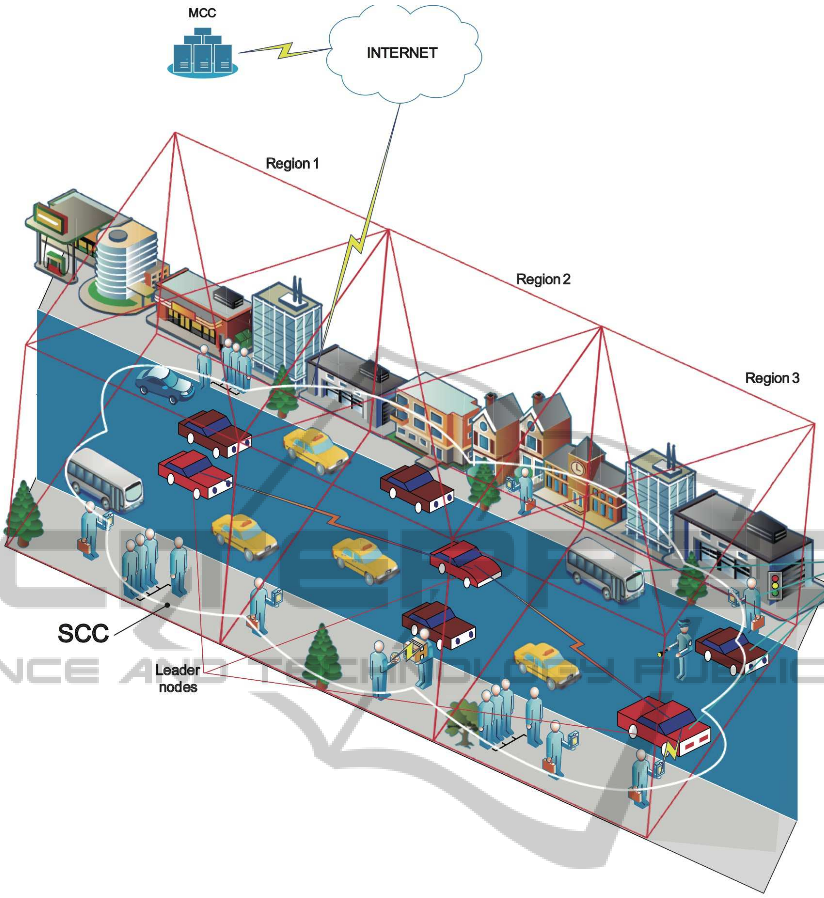

Figure 3: Sporadic Ad-hoc Network (SAN) deployed in a generic mobility scenario involving pedestrians and vehicles.

tions and services over the ad-hoc networks are not interrupted as users move. For that purpose, when a virtual node is about to become inactive (because there are very few devices in that region), the leader sends its buffered information to the leader node of a neighbor region. This information is stored until the original virtual node is operative again (i.e. until new terminals enter that region). At this moment, the just-restored virtual node requests the information and processes it by resorting to the capabilities provided by the new terminals that are located in its region, thus preventing from losing information as nodes move.

4.3 The SCC Paradigm

The devices connected to the SANS require coordination mechanisms to orchestrate the sharing and allocation of their available resources (and even of extra resources available from the Internet). To this aim, the virtual nodes of the VNLayer+ must establish communications to each other in order for the users' devices (i) to request resources to other terminals and (ii) to advertise those that they are willing to provide to the SAN. We have designed a transport-layer approach aimed at sharing and allocating resources, on which the multiple "X"aaS services of our SCC paradigm will be grounded. Each of these service models will require particular refinements (on which we are focusing our ongoing research work) that will be develop on the common substratrum described in this section.

For our descriptions, we assume that the geographic area where the SAN has been deployed is divided into regions (recall Figure 3), whose leaders and backup nodes have been selected by the mechanisms of the VNLayer+ mentioned in Section 4.2. In this scenario, we suppose that a terminal connected to the SAN (hereafter, application node) needs extra resources for running an application and asks for them to the sporadic cloud. This process is organized as follows, as depicted in Figure 4:

- Firstly, the application node broadcasts a Resource Discovery Message (`MRDiscovery`) in its region.

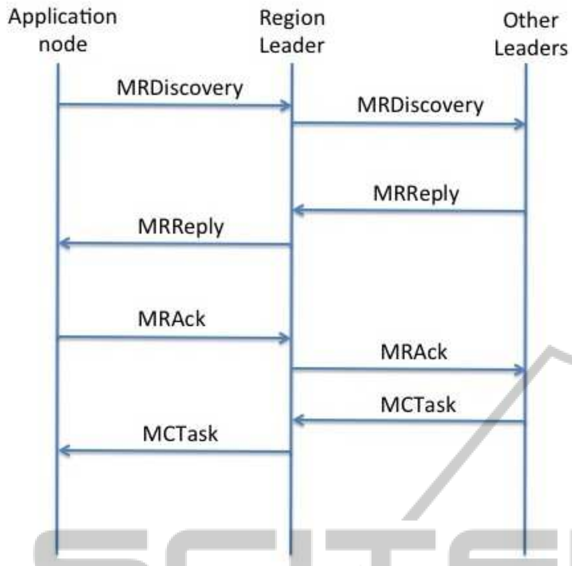

Figure 4: Messages exchanged among the application node and the leaders of the regions identified within a SAN.

- Upon the reception of the `MRDiscovery` message, the leader node of the region sends it to the leaders of its adjacent regions. This process is repeated by the remaining leaders until reaching the last region.
- In each region, the leader includes in the `MRDiscovery` message information about the resources available in the devices supporting that virtual node. The leader of the last region aggregates all the responses and sends this information back to the application node via a Resource Reply Message (`MRReply`) through the leaders of the intermediate regions.
- After receiving the `MRReply` message, the application node distributes tasks among the virtual nodes, considering the availability of resources reported by the leaders of their respective regions. The corresponding task assignment is notified to each leader via a Resource Acknowledgement Message (`MRAck`). This allocation process changes as per the specific “X”aaS service deployed, which at the same time depends on the kind of resources required for the application node (computing, storing, networking...).
- Since virtual nodes might need to cooperate when it comes to getting the information required by the application node, each leader records the tasks to be done by the rest of leaders. This information is also reported to the backup nodes of each leader in order to ensure a correct synchronization among them, so that a backup can take the place of the current leader when this node leaves the region.
- After receiving the `MRAck` message, each leader distributes its task assigment among the physical nodes of the region, by considering the capabilities that these devices put at disposal of the SAN at this moment. Once the terminals have finished their job (which depends obviously on the capabilities shared in each “X”aaS service), the leader of this region is notified. This node finally informs to the application node by sending a Completed Task Message (`MCTask`).
- This way, the cooperation among the virtual nodes enables to get the information required by the application node. Depending on the application to be runned in this node, our approach allows to replicate this information in special virtual nodes of the SAN, so that other application nodes can also access it. Specifically, to this aim, we resort to very stable virtual nodes which are supported by a huge amount of connected devices that are located, for example, in crowded squares or avenues and intersections with a lot of vehicular traffic. In order to access these contents, each node just needs to send a Content Request Message (`MCRequest`) to the leader of its region to trigger the process. The information is finally received via a Content Response Message (`MCResponse`).

Note that above descriptions are also valid for a scenario where multiple application nodes ask for resources simultaneously. In this case, the capabilities available in each virtual node will be considered when deciding about new requests, and, once the resources provided by the terminals are about to finish, the requests will be queued until the resources blocked by other application nodes are released.

As introduced before, having developed the common substratrum of SCC, we need to envisage the specific mechanisms required in each “X”aaS service. Currently, we are working on the N(etworking)aaS model with the goal of allowing the mobile nodes to collaboratively download contents from the Internet and to share them with the rest of members over the SAN. To this aim, we require transport-layer solutions that manage multiple connections over the multi-hop ad-hoc network in a transparent way. These mechanisms should lead to significant improvements in terms of download time, thanks to the simultaneous exploitation of the Internet connections provided from the terminals of several individuals (e.g. if we have 3 terminals with a bandwidth of 1 Mbps each one to download a content of 3 MB, the download would last 3 seconds using just one connection, and 1 second splitting the content in three chunks of 1 MB each one and aggregating the 3 connections).

5 EXPECTED OUTCOME

In our doctoral proposal, wireless mobile nodes opportunistically harness multi-hop ad-hoc paths to exploit the availability of a great amount of (often underutilized) resources that are provided by the terminals of nearby "always-on" users in the context of the SANs.

On the one hand, such SANs promote shared experiences among potentially like-minded users who happen to be close to each other, by relying on direct or hop-by-hop ad-hoc communications. On the other one, in the realm of these ad-hoc networks, our new SCC paradigm enables the deployment of multiple context-aware applications and "X"aaS service models whose novelty is grounded on two features. First, the fact of working with static virtual nodes instead of mobile real nodes, which makes easier the routing tasks and ensures reliable and stable communication environments. Second, the possibility of bringing together in a transparent way (i) the resources provided by each terminal that is connected to the SAN and (ii) the capabilities of the traditional MCC (if available) in diverse ad-hoc mobility scenarios, working at the transport layer and on the top of the virtualization mechanisms.

To put it from another angle, the resource sharing and allocation pursued in our SCC allows to improve the experience of each individual thanks to the capabilities contributed by the rest of members of the SAN, notwithstanding their mobility. This way, the users might, for instance, enjoy connectivity of the Internet and even reduce download times by aggregating the 3G/4G connections of other terminals, use extra storage space (in these devices or even in the cloud) and also exploit additional computational resources provided by more powerful terminals in order to run, for example, (demanding) personalization strategies. Such strategies would allow to provide the SAN users with contents of their interest which might have been proactively selected as per their preferences (and even collected, enhanced and shared by other individuals) within tailor-made sporadic communities.

ACKNOWLEDGEMENT

This work is being funded by the Ministerio de Educación, Cultura y Deporte (Gobierno de España) research project TIN2013-42774-R.

REFERENCES

Arif, S., Olariu, S., Wang, J., Yan, G., and Khalil, I. (2012). Datacenters at the airport: reasoning about time-dependent parking lot occupancy. *IEEE Transactions on Parallel and Distributed Systems*, 23:2067–2080.

Bellavista, P. and Giannelli, C. (2010). Internet connectivity sharing in multi-path spontaneous networks: Comparing and integrating network- and application-layer approaches. In *MOBILWARE'10, 3rd International ICST Conference on MOBILe Wireless MiddleWARE, Operating Systems, and Applications*, pages 84–99.

Bravo-Torres, J. F., López-Nores, M., Blanco-Fernández, Y., Pazos-Arias, J. J., and Ordóñez-Morales, E. F. (2015). Vanetlayer: A virtualization layer supporting access to web contents from within vehicular networks. *Journal of Computational Science*, In press.

Chen, B. B. and Chan, M. C. (2009). MobTorrent: A framework for mobile internet access from vehicles. In *INFOCOM'09, 24TH IEEE Conference on Computer Communications*, pages 1404–1412.

Dinh, H. T., Lee, C., Niyato, D., and Wang, P. (2013). A survey of mobile cloud computing: architecture, applications, and approaches. *Wireless Communications and Mobile Computing*, 13(18):1587–1611.

Dolev, S., Gilbert, S., Lynch, N. A., E., E. S., Shvartsman, A. A., and Welch, J. L. (2004). Virtual mobile nodes for mobile ad hoc networks. *Distributed Computing*, 3274:230–244.

Eriksson, J., Balakrishnan, H., and Madden, S. (2008). Cabernet: Vehicular content delivery using WiFi. In *MobiCom'08, 14th ACM International Conference on Mobile Computing and Networking*, pages 199–210.

Gerla, M. (2012). Vehicular cloud computing. In *MED-HOC-NET'12, 11th Annual Mediterranean Ad-Hoc Networking Workshop*. IEEE Press.

Gerla, M., Wu, C., Pau, G., and Zhu, X. (2014). Content distribution in VANETs. *Vehicular Communications*, 1(1):23–12.

Kuss, D. J. and Griffiths, M. D. (2011). Excessive online social networking: Can adolescents become addicted to facebook. *Education and Heath*, 29(4):68–71.

Lee, K., Lee, S. H., Cheung, R., Lee, U., and Gerla, M. (2007). First experience with CarTorrent in a real vehicular ad-hoc network testbed. In *2007 Mobile Networking for Vehicular Environments*, pages 109–114.

Lee, U., Park, J. S., Yeh, J., Pau, G., and Gerla, M. (2006). Code torrent: Content distribution using network coding in VANET. In *DIWANS'06, Workshop on Dependability Issues in Wireless Ad Hoc Networks and Sensor Networks 2006 (part of Mobicom 2006)*.

Nandan, A., Das, S., Pau, G., Gerla, M., and Sanadidi, M. (2005). Cooperative downloading in vehicular ad-hoc networks. In *WONS'05, 2nd International Conference on Wireless On-Demand Network Systems and Services*, pages 32–41.

Olariu, S., Hristov, T., and Yan, G. (2013). The next paradigm shift: from vehicular networks to vehicular clouds. In Basagni, S., Giordano, S., and Stojmenovic, I., editors, *Mobile ad-hoc networking: cutting edge directions*. John Wiley & Sons.

Przybylski, A. K., Murayama, K., DeHaan, C. R., and Gladwell, V. (2013). Motivational, emotional, and behavioral correlates of fear of missing out. *Computers in Human Behavior*, 29(4):1841–1848.

Sharma, A. and Singh, E. J. (2013). Mobility models for manet: Mathematical perspective. *International Journal of Advanced Research in Engineering and Applied Sciences*, 2(5):59–68.