Urban Scale Dissemination in Mobile Pervasive Computing Environments

Alistair Morris, Mélanie Bouroche and Vinny Cahill
Distributed Systems Group, Trinity College Dublin, Dublin 2, Ireland

1 INTRODUCTION

The unprecedented urban growth (Capello, 2013) has become a main issue characterising the 21st Century. This growth has created new problems and exacerbated others. Urban sprawl, air pollution, traffic bottlenecks, insufficient infrastructure form just small subset of problems potentially remedied through the collection of context data at an urban scale. This context data enables the realisation of better decisions, improved service delivery and greater accountability through greater transparency.

To fully realise this concept and thus the potential of Smart Cities the comprehension of facts (Paskaleva, 2009) and thus context-awareness, defined as the ability to provide services with full awareness of current execution environment, remains fundamental to the building of modern mobile and ubiquitous systems. Several research efforts have previously devoted themselves to attempting to transparently implement the features required by context-aware applications including production, processing, storage, and distribution of context data. This allows the deployment of context-aware services in mobile ubiquitous computing environments.

Context distribution, defined as how we gather and deliver relevant context about the environment to all interested entities connected to the mobile ubiquitous middleware, remains an unsolved problem and continues to emerge as a fundamental research area where existing approaches do not scale. The presence of non-negligible overhead hinders both the scalability and reliability of a middleware (Corradi et al., 2010).

Detailing our arguments, this paper has the following structure: Initially we outline the *Stage of research* and an *Outline of objectives*. Following this we define the *Research Problem* and how we determine this through conducting a review of *Existing Work*. Finally we propose the *Methodology* that we will apply to test the *Hypothesis* and comment on what *Future Work* remains.

2 STAGE OF RESEARCH

Research at present considers how we deliver context data. Even if a range of solutions exist, it becomes possible to compare the suitability of each particular deployment scenario. For instance, in an ad-hoc network deployment, flooding/gossip-based protocols become more suitable where a lack of structured dissemination overlay ensures adaptability.

At present we remain focused in the process of extensively evaluating the main approaches that disseminate context data. We hope to discover and define the main benefits and shortcomings of these. This becomes crucial when determining the scalability of any context-aware application. Our principle research focus remains on how this affects context provisioning to mobile pervasive devices in urban scale environments.

Differing from already existing work, our research focuses on context data distribution to deeply study the main requirements, implementation primitives, and identify the main research challenges.

3 OUTLINE OF OBJECTIVES

A gap in research exists that hinders the realisation of context-aware distribution in the wide-area wireless networks that span the Smart City. We see that past research has mainly focused on small-scale deployments, often limited to homes and commercial buildings, where context data distribution has affordable run-time overhead. As a consequence, previous research work has mainly addressed local middleware issues in order to support context provisioning to supported applications. These approaches have also used simple and centralised approaches to implement the distribution process. Such approaches remain poorly suited for the scalable middleware required in the Smart City.

In addition, context data distribution involves different protocol layers (from OSI network to application layer), and covers various emerging research

Morris A., Bouroche M. and Cahill V. (2014).
Urban Scale Dissemination in Mobile Pervasive Computing Environments.
In *Doctoral Consortium*, pages 18-22

fields, including traditional (infrastructure) mobile ubiquitous systems, Mobile Ad-hoc NETworks (MANET), Vehicular Ad-hoc NETworks (VANET), and Delay Tolerant Networks (DTN) (Conti and Giordano, 2007). As different network and middleware deployments influence context provisioning additional research remains a requirement as we do not yet fully understand how they influence and limit the scalability of context data distribution in the Smart City.

4 RESEARCH PROBLEM

Context data distribution, namely the gathering and delivery of relevant context data to all interested entities connected in the mobile ubiquitous middleware, has emerged as a new research area in context-aware systems (Baldauf et al., 2007). In fact, context data distribution remains significant from both the perspective of the application and the middleware in the Smart City. Service adaptation becomes triggered by received context data: hence, the timely delivery of context data enables services that promptly adapt to the current execution context.

The middleware also has to transparently manage and route huge amounts of time sensitive context data to mobile nodes: especially in wide-area mobile networks, that can lead to non-negligible overhead, thus hindering both system scalability and reliability. Seminal surveys already argue the significance of context data distribution in context-aware middleware. (Baldauf et al., 2007) (Sinderen et al., 2006).

5 EXISTING WORK

Dissemination enables data flow from the source to the sinks of the context data. Dissemination solutions belong to three different categories: flooding-based, overlay-based, and gossip-based.

The first two categories form deterministic approaches where a sink receives matching context data produced by sources belonging to the same context data distribution method. The last category takes the form of a probabilistic approach meaning that a sink has a likelihood to miss some matching data. Systems adopting a hybrid approach that mixes these three main dissemination solutions also exist. We present examples of the use of these context data dissemination methods by citing well known systems.

5.1 Flooding-based

Flooding-based algorithms realise context data dissemination via flooding operations. Here each node broadcasts known context data across the entire network. This means that receiver nodes must locally select context data as required. For instance, Adaptive Traffic Lights exchanges those context data useful to co-ordinate red/yellow/green timings with vehicles at an intersection (Gradinescu et al., 2007). HiBOp and CAR floods neighbours with data useful for the maintenance of the DTN routing infrastructure (Boldrini et al., 2007) (Musolesi and Mascolo, 2009). In MANIP, each data comes with a physical locality tag that limits the physical locality where data becomes flooded (Macedo et al., 2009).

5.2 Overlay-based

Overlay-based algorithms build dissemination “backbones” in a deterministic manner by using context data subscriptions. This causes all data dissemination to take place over a networked overlay.

To build the overlay nodes must exchange control information and this introduces overhead through the additional communication that must take place. Overlay-based approaches offer two visibility scopes to each subscription for scalability: system wide scope and limited scope. In network wide scope, the dissemination process ensures that each subscription remains visible to the whole network causing the retrieval of matching data to always remains possible. With limited scope, the dissemination process limits subscription visibility to a subset of nodes; unfortunately this means that there exists a possibility that the method will not find the entire set of matching context data.

The COPAL middleware employs centralised brokers in which all subscriptions have network-wide visibility scopes (Li et al., 2010). CMF similarly uses distributed brokers that coordinate to supply data to mobile nodes according to specific requirements (Kranenburg et al., 2006). Pervaho uses a Location-based modification of the Publish/Subscribe (LPSS) paradigm by imposing location-based constraints: each publication and each subscription has a visibility scope, and the delivery of published context data to an active subscription only occurs if publisher and subscriber lie in the intersection of these two scopes (Eugster et al., 2009). Thus, Pervaho performs context-/location-based filtering to limit the amount of events received. Here the adopted network layer ensures that subscriptions have network wide visibility scopes. Finally, MobiSoC and MiddleWhere use overlay-based methods to disseminate context data to mobile nodes; at the same time, central servers provide network-wide visibility to context subscriptions (Grupta et al., 2009).

We see a limited scope approach in EgoSpaces that introduces a overlay-based approach based on agents (Julien and Roman, 2006). Each agent operates over multiple views that includes context data associated with hosts/agents in the physical locality. Each view imposes constraints on contained data/resources meta-data as each view spans close neighbours. This leads to a limited visibility scope approach. Habit realises data dissemination in a limited scope overlay-based manner by exploiting the physical proximity of nodes and the social relationships of users (Mashhadi et al., 2009). It uses a regularity graph that keeps trace of when and how often two nodes communicate, and an interest graph that keeps trace of nodes interests, to build dissemination paths. Mobile Gaia groups nodes into clusters, and it uses an event service to disseminate context data: because every cluster has its own event service, the final approach remains overlay-based with limited visibility scope (Chetan et al., 2005).

5.3 Gossip-based

Gossip-based algorithms disseminate data in a probabilistic manner by requiring each node to resend the context data to a random set of neighbours. These approaches suit fast-changing and unstable networks, such as MANETs (Friedman et al., 2007) as they do not attempt to construct and maintain complex routing infrastructure. Gossip-based protocols fall into two board categories: context-oblivious and context-aware approaches (Friedman et al., 2009).

Context-oblivious protocols use random re-transmission probabilities. This means that no external context information becomes used to tailor their operation. These include pure probabilistic gossip techniques that simply resend each received data with a re-transmission probability that differs for each neighbour node and depends either on the local node density or on neighbourhood information. No heavy computation becomes necessary on traversed nodes as the protocol selects nodes randomly. Unfortunately this does not guarantee performance or scalability; such protocols waste networking resources through the gossiping of data unnecessarily.

Context-aware protocols select neighbouring nodes for gossiping context data by using some form of external context data. For instance, some approaches use node distance to position context data replicas; other approaches use social similarity, for example the membership of a group, to select the nodes for gossiping data to. Put simply, context-aware approaches minimises the amount of uselessly gossiped context data but they require more co-ordination to exchange the context data used to make gossiping decisions. For example, HiBOp and CAR exploit context-awareness for message routing purpose: above all, to select the best forwarder during message routing (Boldrini et al., 2007) (Musolesi and Mascolo, 2009).

5.4 Hybrid-based

Finally, hybrid approaches have become popular. Active Highways collects context data flooded from sensors local/remote to the vehicles, and relies upon fixed servers to assist travelling vehicles; a overlay-based approach with limited visibility scopes disseminates data to servers (Iftode et al., 2008). Similarly, HiCon exploits direct sensor access to produce data, and a overlay-based approach to disseminate context data by imposing locality scope (Cho et al., 2008). Finally, in MobEyes, data becomes acquired either through deployed sensor hardware or by other vehicles during contacts: vehicles in proximity exchange data by using a flooding-based approach where each vehicle broadcasts context data either locally sensed or received by other vehicles according to the system configuration (Lee et al., 2006).

5.5 Conclusive Overview

Flooding-based and gossip-based methods do not present good candidates for our pervasive Smart City scenario: it appears neither can achieve both scalability and reliability for context data delivery. Overlay-based methods present an appealing approach but in a pervasive mobile heterogeneous environment it lacks feasibility where we must always rebuild dissemination backbones as the network exhibits mobility. Hence we look closer at the hybrid based approach where flooding/overlay-based context data dissemination methods enable the delivery of control messages. With this an urban-scalable overlay-based context data approach becomes possible.

6 METHODOLOGY

The NS3 network simulator will enable the examination of performance for each of the hybrid context data dissemination protocols. NS-3 has become regraded as an efficient and scalable simulator (Chaudray et al., 2012). This makes it possible to model the large heterogeneous networks encountered in the Smart City. Therefore we aim to investigate performance using mobility models deployed at the IP layer with nodes densities that vary.

The simulation of mobility at the IP layer infers the non-availability of end-points and the variation of latency. We concentrate on two performance metrics in particular *Delivery Ratio* and *Dissemination Overhead*. This enables us to determine how mobility affects performance and thus scalability.

We define respectively the *Delivery Ratio* as the received context data packets divided by context packets sent by the application and *Dissemination Overhead* as the fraction of context data used by the protocol for the dissemination of control messages. Similar research efforts have previously employed *Delivery Ratio* as a benchmark for efficiency (Tsirigos and Haas, 2004) (Peng and Lu, 2000). Similarly previous work also shows *Dissemination Overhead* can provide an effect indicator (Dubois-Ferriere et al., 2004) (Zhao et al., 2006).

To model a realistic Smart City scenario node densities should exceed 10,000 nodes with the aim to simulate a context data dissemination network that has a size comparable to 1% of the population of city of 1 million citizens. This will test the scalability and performance of the context data dissemination method and, importantly, the simulator itself.

7 HYPOTHESIS

We hypothesise that hybrid solutions that use different dissemination algorithms together ensures better performance and therefore scalability. For instance, when overlay-based and flooding/gossip-based methods become used at the same time: while the overlay-based approach ensures context access, flooding/gossip based approaches can disseminate data in a probabilistic manner thus reducing context access time and increasing reliability.

Overall flooding-based and gossip-based algorithms appear promising in their performance and thus their scalability. Even if flooding-based methods have scalability issues, flooding remains scalable if constrained by scope.

Gossip-based approaches improve scalability by reducing the delivery guarantees. In contrast to context-oblivious approaches that waste networking resources, context-aware gossip-based protocols represent a more efficient way of building a context data dissemination overlay. For example, HiBOp and Habit uses social relationships as good hints to drive gossip decisions (Boldrini et al., 2007) (Mashhadi et al., 2009). Also, CAR demonstrates that the utilisation of low-level time context information, contact frequency in particular, produces viable results (Musolesi and Mascolo, 2009).

8 FUTURE WORK

Toward the main goal of realising how middleware might adapt different context data dissemination algorithms at run-time to foster greater performance and urban scalability, additional research remains planned that aims to define attributes that:

- Drives the selection of the proper context data dissemination algorithms at the run-time
- Enable the adaption of the run-time behaviour for a specific context dissemination algorithm to maximise performance and thus scalability

These attributes represent contextual modelled and processed aspects useful for achieving scalability. For example, (Taherkordi et al., 2008), models the context attributes of a wireless sensor network that then prompts adaption.

No solution adapts or switches different context data dissemination algorithms at run-time to maximise scalability depending on the current status. Research therefore remains to ensure that the adaption of dissemination algorithms depending on run-time conditions has the potential to enable urban scalable mobile ubiquitous context data dissemination.

At the time of writing a thorough analysis of existing work has been conducted. The design and implementation of the network simulations remains an ongoing progress. The main challenge remains in how simulations can become scalable while still accurately reflecting the complex nature of the mobile ubiquitous networks in the Smart City. To conclude research also remains in how the specific implementation of a hybrid method may influence performance and thus scalability. For example the notion of scope in flooding-overlay based hybrids or the choice of context that drives context-aware gossip-overlay hybrids.

REFERENCES

Matthias Baldauf, Schahram Dustdar, and Florian Rosenberg. A survey on context-aware systems. *International Journal of Ad Hoc and Ubiquitous Computing*, 2(4):263–277, 2007.

Chiara Boldrini, Marco Conti, Jacopo Jacopini, and Andrea Passarella. Hibop: a history based routing protocol for opportunistic networks. In *World of Wireless,* *Mobile and Multimedia Networks, 2007. WoWMoM 2007. IEEE International Symposium on a*, pages 1–12. IEEE, 2007.

Roberta Capello. Recent theoretical paradigms in urban growth. *European Planning Studies*, 21(3):316–333, 2013.

Rachna Chaudhary, Shweta Sethi, Rita Keshari, and Sakshi Goel. A study of comparison of network simulator-3 and network simulator-2. *IJCSIT) International Journal of Computer Science and Information Technologies*, 3(1):3085–3092, 2012.

Shiva Chetan, Jalal Al-Muhtadi, Roy Campbell, and M Dennis Mickunas. Mobile gaia: a middleware for ad-hoc pervasive computing. In *Consumer Communications and Networking Conference, 2005. CCNC. 2005 Second IEEE*, pages 223–228. IEEE, 2005.

Kyungmin Cho, Inseok Hwang, Seungwoo Kang, Byoungjip Kim, Jinwon Lee, SangJeong Lee, Souneil Park, Junehwa Song, and Yunseok Rhee. Hicon: a hierarchical context monitoring and composition framework for next-generation context-aware services. *Network, IEEE*, 22(4):34–42, 2008.

Marco Conti and Silvia Giordano. Multihop ad hoc networking: The reality. *Communications Magazine, IEEE*, 45(4):88–95, 2007.

A. Corradi, M. Fanelli, and L. Foschini. Adaptive context data distribution with guaranteed quality for mobile environments. In *Wireless Pervasive Computing (ISWPC), 2010 5th IEEE International Symposium on*, pages 373–380, 2010.

Henri Dubois-Ferriere, Deborah Estrin, and Thanos Stathopoulos. Efficient and practical query scoping in sensor networks. In *Mobile Ad-hoc and Sensor Systems, 2004 IEEE International Conference on*, pages 564–566. IEEE, 2004.

Patrick Eugster, Benoît Garbinato, and Adrian Holzer. Pervaho: A specialized middleware for mobile context-aware applications. *Electronic Commerce Research*, 9(4):245–268, 2009.

Roy Friedman, Daniela Gavidia, Luis Rodrigues, Aline Carneiro Viana, and Spyros Voulgaris. Gossiping on manets: the beauty and the beast. *ACM SIGOPS Operating Systems Review*, 41(5):67–74, 2007.

Roy Friedman, Anne-Marie Kermarrec, Hugo Miranda, and Luís Rodrigues. Gossip-based dissemination. In *Middleware for Network Eccentric and Mobile Applications*, pages 169–190. Springer, 2009.

Victor Gradinescu, Cristian Gorgorin, Raluca Diaconescu, Valentin Cristea, and Liviu Iftode. Adaptive traffic lights using car-to-car communication. In *Vehicular Technology Conference, 2007. VTC2007-Spring. IEEE 65th*, pages 21–25. IEEE, 2007.

Ankur Gupta, Achir Kalra, Daniel Boston, and Cristian Borcea. Mobisoc: a middleware for mobile social computing applications. *Mobile Networks and Applications*, 14(1):35–52, 2009.

Liviu Iftode, Stephen Smaldone, Mario Gerla, and James Misener. Active highways (position paper). In *Personal, Indoor and Mobile Radio Communications, 2008. PIMRC 2008. IEEE 19th International Symposium on*, pages 1–5. IEEE, 2008.

Christine Julien and G-C Roman. Egospaces: Facilitating rapid development of context-aware mobile applications. *Software Engineering, IEEE Transactions on*, 32(5):281–298, 2006.

Uichin Lee, Biao Zhou, Mario Gerla, Eugenio Magistretti, Paolo Bellavista, and Antonio Corradi. Mobeyes: smart mobs for urban monitoring with a vehicular sensor network. *Wireless Communications, IEEE*, 13(5):52–57, 2006.

Fei Li, Sanjin Sehic, and Schahram Dustdar. Copal: An adaptive approach to context provisioning. In *Wireless and Mobile Computing, Networking and Communications (WiMob), 2010 IEEE 6th International Conference on*, pages 286–293. IEEE, 2010.

D Macedo, A Dos Santos, José Marcos S Nogueira, and Guy Pujolle. A distributed information repository for autonomic context-aware manets. *Network and Service Management, IEEE Transactions on*, 6(1):45–55, 2009.

Afra J Mashhadi, S Ben Mokhtar, and Licia Capra. Habit: Leveraging human mobility and social network for efficient content dissemination in delay tolerant networks. In *World of Wireless, Mobile and Multimedia Networks & Workshops, 2009. WoWMoM 2009. IEEE International Symposium on a*, pages 1–6. IEEE, 2009.

Mirco Musolesi and Cecilia Mascolo. Car: context-aware adaptive routing for delay-tolerant mobile networks. *Mobile Computing, IEEE Transactions on*, 8(2):246–260, 2009.

Krassimira Antonova Paskaleva. Enabling the smart city: The progress of city e-governance in europe. *International Journal of Innovation and Regional Development*, 1(4):405–422, 2009.

Wei Peng and Xi-Cheng Lu. On the reduction of broadcast redundancy in mobile ad hoc networks. In *Proceedings of the 1st ACM international symposium on Mobile ad hoc networking & computing*, pages 129–130. IEEE Press, 2000.

Amirhosein Taherkordi, Romain Rouvoy, Quan Le-Trung, and Frank Eliassen. A self-adaptive context processing framework for wireless sensor networks. In *Proceedings of the 3rd international workshop on Middleware for sensor networks*, pages 7–12. ACM, 2008.

Aristotelis Tsirigos and Zygmunt J Haas. Analysis of multipath routing-part i: The effect on the packet delivery ratio. *Wireless Communications, IEEE Transactions on*, 3(1):138–146, 2004.

Herma Van Kranenburg, Mortaza S Bargh, Sorin Iacob, and Arjan Peddemors. A context management framework for supporting context-aware distributed applications. *Communications Magazine, IEEE*, 44(8):67–74, 2006.

Marten J van Sinderen, Aart T van Halteren, Maarten Wegdam, Hendrik B Meeuwissen, and E Henk Eertink. Supporting context-aware mobile applications: an infrastructure approach. *Communications Magazine, IEEE*, 44(9):96–104, 2006.

Youping Zhao, Jeffrey H Reed, Shiwen Mao, and Kyung K Bae. Overhead analysis for radio environment mapenabled cognitive radio networks. In *Networking Technologies for Software Defined Radio Networks, 2006. SDR'06.1 st IEEE Workshop on*, pages 18–25. IEEE, 2006.