Mastering Emergent Behavior in Large-Scale Networks

Andrei Pruteanu
Embedded Software Group
Delft University of Technology
a.s.pruteanu@tudelft.nl

Abstract

With the ever increasing scale of mobile wireless networks (such as MANETs, WSNs and VANETs), there is a growing need for performing aggregate computations via distributed, robust and scalable algorithms. Although traditional research initiatives have already addressed these goals, high network dynamics such as node mobility have been largely ignored. Given our initial positive results regarding algorithms that cope well with mobility in large-scale networks, we want to further validate our ideas on how to address the shortcomings induced by node mobility or any other network dynamics. Our approach was to start with finding promising mechanisms. Then, as a second step, we plan to find ways of describing the desired system behavior while letting the network adapt and evolve the local interactions.

Keywords

emergent behavior, local interaction, adaptive behavior, mobile networks

1 Introduction and Related Work

Recent years have seen a significant increase in the number and the diversity of the devices that form the wireless networks around us. The networks of the future are characterized by tremendous size (orders of magnitude larger compared to the current experiments), high topological dynamics (node mobility), heterogeneity (different device capabilities), presence of actuation and numerous sensors. The research problem that we see arising is finding a means to harness the un-controllable emergent behavior exhibited by mobile large-scale networks. As a direct application, in future transport and logistics warehouses, mobile robots will drive the mobile infrastructure as predicted by Gartner (see Figure 1). Through the SenterNovem Free Project, by which this research is funded, we look into new ways of dealing with the challenges of the future transport and logistics infrastructure. Networks of wireless sensor networks that track pallets, boxes or any other large containers will have to cope with increased scale of the network and node mobility. Research domains such as *mobile ad-hoc networks* (MANETs) and *wireless sensor networks* (WSNs) have studied the corresponding scalability issues, for example, by providing theoretical boundaries [3, 4]. However, large collections of networked devices also bring in the problem of mobility - the larger the network, the higher the probability that individual or groups of devices become mobile.

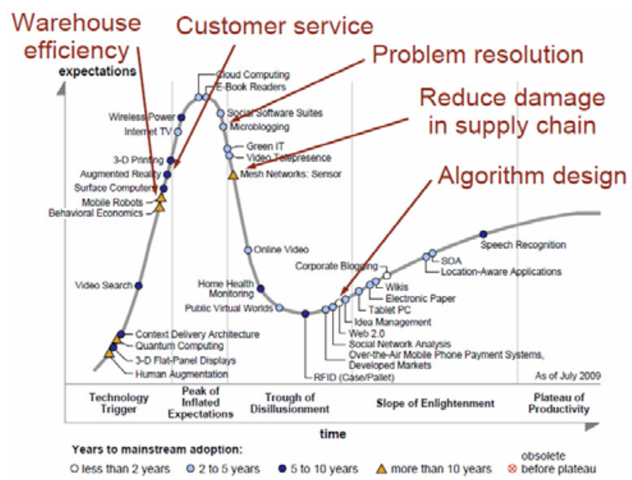

Figure 1. Gartner predictions for future technologies

Research projects targeting the development of large-scale cyber-physical systems, including *programmable matter* [2], swarms of tiny robots [6] and *amorphous computing* [1], take mobility as a default assumption. The scalability aspect, in particular, the need to program the network as a whole rather than as an individual set of nodes [5, 8, 9] is one of the most difficult and challenging research aspects.

The proposed ideas in this call for contributions will be explored during a PhD track spanning from June 2009 to June 2013.

2 Problem Description

We find that a rather unexplored, yet important research direction is to address scalability issues in the context of truly large-scale deployments that exhibit considerable dynamics (e.g. node mobility). Layer-based communication protocols are one of the building blocks of Internet technologies today. The problem is that such approaches are not feasible for large-scale distributed systems that are also characterized by node mobility and lack of static infrastructure (e.g. the mobile telephony base-stations). Additionally, layer-based protocols suffer from sub-optimality. Very difficult problems to cope with are the undesired emergent behavior, difficulty of programming and debugging node-level code, lack or limited time synchronization, unpredictable and highly variable network topology. There is a need for re-thinking the way we program and interact with the large-scale systems with respect to their extreme properties. We argue that in order to harness their power, these ensembles have to be designed and programmed with a rather radical new approach in terms of distributed algorithms of choice and communication paradigms. What is novel in our research is the fact that we use the emergent behavior as a design tool.

2.1 Problem Domain

How does one program a large-scale network to act as a coherent ensemble? What are the techniques or mechanisms that cope well with topology dynamics (node mobility), allow the control of emergent behavior without any coordination and rely only on local communication for all this to happen? Is there a way to ensure self-healing, self-organizing properties in a decentralized manner?

The answers to these questions will have to incorporate key features and properties of the algorithms that do not rely solely on the properties of the classic models such as random geometric graphs, cellular automata, evolutionary algorithms etc. Properties such as node density, computational aggregates through distributed variables, their stability and convergence speed are important building blocks and determine some important theoretical aspects of our algorithms.

The research will be centered on dealing with the previously mentioned problems that large-scale networks of spatially distributed devices usually exhibit. We will develop algorithms that scale well and are resilient to node mobility and we will create programming techniques such that these systems will showcase self-organizing properties through local interaction between the devices. The overall context into which we place our work is somewhere in between several research fields such as distributed systems, software technologies, biologically inspired algorithms, information technology, etc. The proposed solutions will borrow techniques inspired by natural or artificial systems that exhibit similar properties in terms of scale and dynamics.

2.2 Specific Problem Addressed

Imagine a large transport and logistics warehouse. Packets and pallets are moved constantly. The task of tracking their status is a tedious one and requires a lot of resources in terms of equipment and human activity needed to supervise and maintain the well-functioning of the systems. Current solutions address the case of static infrastructure with reduced or no node mobility at all. This simplifying scenario is not tracking the real-life dynamics found in logistics warehouses. There is a need to improve on that.

3 Solutions - Research Ideas

A lot of the research initiatives for the case of distributed systems follow a bottom-up approach to explore potential benefits of various node interactions. While this is feasible for small systems, given the exponential increase of connected devices that networks of the future will be constituted of, the opposite top-down approach is emerging as a viable solution. We chose to explore this design path. Some related work presented in Section 1 is already considering the network as an ensemble that behaves coherently. This shift in paradigm is no longer trying to obtain complex behavior of the system by programming individual devices and hope that their interaction will lead to the desired behavior. Instead, the top-down approach is considering the network as a single device (entity) that has to be programmed to perform different functionalities. Through local interaction the nodes exchange information and compute (update) the system state according to global specifications.

3.1 The Original Idea of the Thesis

So far we identified several promising mechanisms. The first one is the gossiping algorithm that was developed for Internet applications like peer-to-peer overlay systems. It proves itself highly valuable when knowledge about the topology is lacking or the network exhibits high variation in size. In a distributed manner, it allows applications to compute aggregates such as sums or averages. We used this mechanism to compute network size estimation and a new clustering algorithm as described in Section 4.

The second mechanism is the Cellular Automata. It has proven its validity as a model for several physical phenomena such as gas dynamics, plasma physics, oil-water partitioning. They are characterized by a large collection of interacting elements by means of local interaction. We plan to take this mechanism and adapt the assumption space, considering that time synchronization is not always a given and the topology is dynamic, and create new models for dynamic large-scale networks. We will use this idea in a future paper.

3.2 Expected Contribution to the Field of WSN

Our work is positioned such that the transition to large scale-networks is facilitated through new methodologies and classes of algorithms created specifically for these systems. Simply borrowing tools and techniques from the current generation of sensor networks research will most likely cause the failure of harnessing the benefits of the large-scale systems. There is also a high difficulty in trying to predict how the interaction of individual nodes results in a coherent, complex, stable, self-healing network-wise behavior. Additionally, dealing with software bugs that occur only-after certain scale of the network is reached, is a notoriously difficult problem. From all these consideration, arises an obvious need for techniques and mechanisms designed specifically for this class of systems. Given that diffusion of information is a very important aspect, node mobility can actually be beneficial with regards to this aspect. Similarly, the unpredicted behavior resulting out of the interaction of a large number of devices can be harnessed and used to create novel approaches to classic applications like software updates, clustering and routing of information in a mobile network. Our contribution to the field of wireless sensor networks is the fact we use the emergent behavior that is associated with large-scale networks to achieve better adaptability to system dynamics. In other words, the software on the nodes evolves and, by means of local interactions, achieves the desired global behavior while being resilient to node mobility, links failures etc. Reliable links do not always characterize networks with high dynamics. Assigning functionality to groups (domains)

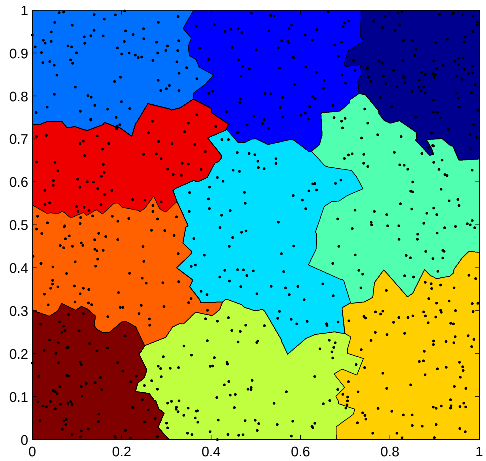

Figure 2. ASH: Quasy-static overlay on top of a mobile network

of nodes, although at a loss of granularity, opens a whole range of applications. Different sections of the network perform different tasks given some specific properties such as location, event detection, power levels, etc.

3.3 The Methodological Approach

No longer the topology dynamics, scalability issues and the control of emergent behavior through new programming paradigms can be ignored or deferred for future work. The means we use to address these challenges will first consist of identifying mechanism such as gossiping, cellular automata, biologically inspired mechanisms etc. that are promising solutions. Secondly, we will create algorithms that not only offer solutions for strict class of problems, but also evolve and adapt to various network properties changes such as topology, node count and quality of the links. Then we will compare our solutions to existing state-of-the-art to check the validity of our ideas. The goal is to identify algorithms that scale well for the case of networks deployed in transport and logistics applications, such as large warehouses.

4 Results and future work

In a first paper, we introduced *ASH* [7], a novel mechanism for constructing a quasi-static overlay network on top of a mobile infrastructure (see Figure 2). By means of information diffusion and a mechanism inspired by the equilibrium of gases inside a container we are able to partition the network into stable domains that "hide" the node mobility for certain classes of applications.

In a second paper, that is still under review, we introduced the *NetSize* algorithm to estimate the size of various static and dynamic networks. Derived from gossiping algorithms, our extension incorporates the notion of resets for being able to provide, within certain boundaries, an estimation of the network size.

A third paper, which is work in progress, extends the cellular automata models to random geometric graphs. The goal is, by means of visualization methodologies, to obtain stable geometric patterns given a dynamic network. They can uncover potentially useful applications such as distributed software updates (e.g. the patterns has a wave-like distribution), clustering (e.g. quasi-static patches of homogenous color) and gradients for routing information etc. The search through the cellular automata rule space will have to consider complex entropy measures of the system-state given a space and time evolution of the network.

In the future we plan to include even more autonomy and means of evolving desired behavior through genetic programming. The system (network) needs to adapt (evolve) and self-organize such that the goal is achievable starting from any random condition and is resilient to various dynamics. In other words, although randomness is defining most of the properties, given the necessary guidance in the evolution the network maintains the desired overall behavior while constantly adapting the local interactions of the nodes. Most of the engineering approaches limit themselves to strict protocol design based on application/deployment-specific requirements.

We argue that sometimes, no matter how elaborate and imaginative the human imagination can be, there exists solutions that cannot be simply created by means of classical approaches. Different solutions in the design space have to be evolved and tested given less strict requirements. In other words, sometimes its better to leave the system evolve and, with some guidance, find the desired solution.

5 Acknowledgments

The research is financed by the Dutch government through the SenterNovem-Free Project.

6 References

[1] Abelson et al. Amorphous computing. *Commun. ACM*, 43(5):74–82, 2000.

[2] Goldstein et al. Programmable matter. *IEEE Computer*, 38(6):99–101, June 2005.

[3] M. Grossglauser and D. N. C. Tse. Mobility increases the capacity of ad hoc wireless networks. *IEEE/ACM Trans. Netw.*, 10(4):477–486, 2002.

[4] P. Gupta and P. Kumar. The capacity of wireless networks. *Information Theory, IEEE Transactions on*, 46(2):388–404, mar 2000.

[5] Hnat et al. Macrolab:a vector-based macroprogramming framework for cps. In *SenSys 2008*.

[6] M. Karpelson et al. Milligram-scale high-voltage power electronics for piezoelectric microrobots. In *ICRA 2009*.

[7] A. Pruteanu, S. Dulman, and K. Langendoen. Ash: Tackling node mobility in large-scale networks. In *IEEE SASO 2010*, sep 2010.

[8] D. Rosa et al. Programming modular robots with locally distributed predicates. In *ICRA 2008*, 2008.

[9] D. Yamins. *A theory of local-to-global algorithms for one-dimensional spatial multi-agent systems*. PhD thesis, Harvard, Cambridge, MA, USA, 2008. Adviser-Nagpal, Radhika.