

Population Protocols on Real Social Networks *

Luca Becchetti Lorenzo Bergamini Francesco Ficarola Andrea Vitaletti

Sapienza University of Rome

becchett,bergamini,ficarola,vitale@dis.uniroma1.it

Abstract

In this paper we present an experimental analysis to assess the performance of Population Protocols, a well-known fully decentralized computational model, on a real, evolving social network.

Categories and Subject Descriptors H.4 [Information Systems Applications]: Miscellaneous

General Terms Measurement, Performance, Experiments

Keywords Social Networks, Population Protocols

1. Introduction

Only a few attempts to collect data from real social interactions have been made ([3], [1], [6]), mainly because of logistical difficulties. In fact, distributing suitable devices to a significant population of individuals is an expensive and time-consuming task.

In this paper we present experimental work on a real, evolving social network we collected over one week. We used the data we recorded to create a real movement pattern on an agent simulator (NetLogo) to compare the behavior of population protocols on a random topology (as done in the original work) to those obtained when using our realistic mobility pattern. Results showed that even though the full convergence time is greater in case of the realistic mobility, the vast majority of agents ($\approx 80\%$) converge well before and in any case the population protocols work fine even with a mobility model taken from reality.

2. Population Protocols

Population protocols (we often use the acronym PP in the following) provide a theoretical model of a collection of tiny,

*Partially supported by the Italian MIUR-PRIN'08 project COGENT (COmputational and GamE-theoretic aspects of uncoordinated NeTworks)

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright is held by the author/owner(s).

SNS'12 April 10, 2012, Bern, Switzerland

Copyright © 2012 ACM ACM 978-1-4503-1164-9/12/04...\$10.00

possibly mobile agents that interact to carry out a computation. Agents are assumed to be identical finite state machines. Each agent initially possesses an input value that determines its initial state, while pairs of agents can exchange state information whenever they interact. Interaction patterns of the agents are unpredictable, but they are subject to some fairness constraints, so that any computation eventually converges to a stable output.

It has also been shown in [4] that the following three predicates are sufficient to define all those that can be computed by Population Protocols (the set of semilinear predicates):

- *Threshold Predicate*. For any given $a \in \mathcal{X}$, this predicate is true if $\#a \geq T$ where T is a given threshold¹.
- *Modulo Predicate*. This predicate is true whenever $\#a \equiv j \pmod{k}$ for given j and k .
- *Comparison Predicate*. This predicate is true whenever $\#a \geq \#b$, for $a, b \in \mathcal{X}$.

We considered the PP model since it makes minimalistic, bottom-line assumptions as to the computational, communication and storage capabilities of the agents, thus increasing the robustness of results obtained.

3. Experimental analysis

Our experimental analysis aimed at assessing the convergence properties of the PP model over a real, evolving social network.

Social network. For a period of 1 week (working days - Mon-Fri) we tracked in a completely anonymous way the movement of people at the first floor of our Sapienza university department of computer and systems science, site in Rome. We asked each of ≈ 120 students to wear an active RFI badge (tag in the following, that periodically broadcasts information about contacts with similar tags (i.e., whenever the person wearing the tag came close to another member of the volunteer group).² An example of the kind

¹In the remainder, for any $a \in \mathcal{X}$, $\#a$ denotes the number of agents whose input value is the symbol a .

²A high spatial resolution of less than 1.2 meters is attained by using very low radio power levels for the contact sensing. Furthermore, assuming that the subjects wear the tags on their chest, the body effectively acts as a shield for the sensing signals. This way, contacts are detected only when

of information we can get is for instance "Tag **xxxx** met tag **yyyy** at time **tttt**". The timestamp information is extremely important, as the social network we will produce is the network of contacts that took place in a given time-interval T .

At the end of the experiment we collected $\approx 250\text{MB}$ of data logs; then we parsed the collected data so that they could be used as an input to NetLogo simulator ([2]) to reproduce the movement of agents (each agent representing a student in this case).

Application to population protocols. We used the data collected to analyze the convergence properties of standard PP algorithms to compute the predicates described in Section 2, when interactions are determined by real movement patterns of the agents. We measured the convergence time (i.e. time required for all the agents to stabilize on the same correct output) of population protocols when the standard adversary model is adopted (i.e. a probabilistic adversary that chooses with the same probability a couple of agents to interact) and compared it to the time required when communication occurs according to the patterns determined by social contact graph we collected.

Experimental results. In this paragraph, we discuss the results for Modulo and Threshold Predicates, those for Comparison have a similar flavour. Results for this experiment are drawn in Figures 1 and 2. As can be noticed by comparing the 2 graphs, the convergence is faster on the random topology than on the social one for both predicates, with the modulo requiring a bit more ticks to converge. This is an expected behavior and is due to the fact that in the first case we are considering a complete graph, thus every node can be selected to interact with every other node in the network, which fastens the convergence toward the right result. The same consideration does not apply to the social network case, where a node can be selected to start an interaction only when it is currently transmitting data in the time step under consideration (≈ 50 seconds) and it can select a node to interact with only among its neighbors. In any case, a very interesting aspect to remark is that even if the total convergence in the social network case is reached after ≈ 1 million and 2 millions ticks for the 2 predicates, in both cases a significant percentage of nodes converges well before (around at 1/3 of the total convergence time): this means that the real mobility has an obvious impact on the full convergence (i.e. all nodes converge), but this is most probably due to the fact that the few missing nodes (the ones that must still converge) are selected with lower probability because they probably have a lower degree (i.e. they had less contacts).

References

[1] Reality mining project. <http://reality.media.mit.edu/>.

participants actually face one another. If a sensed contact persists for a few seconds, then given the short range and the face-to-face requirement, it is reasonable to assume that the experiment is able to detect an ongoing social contact (as e.g. a conversation)[5]




Figure 1. Modulo and threshold: random topology

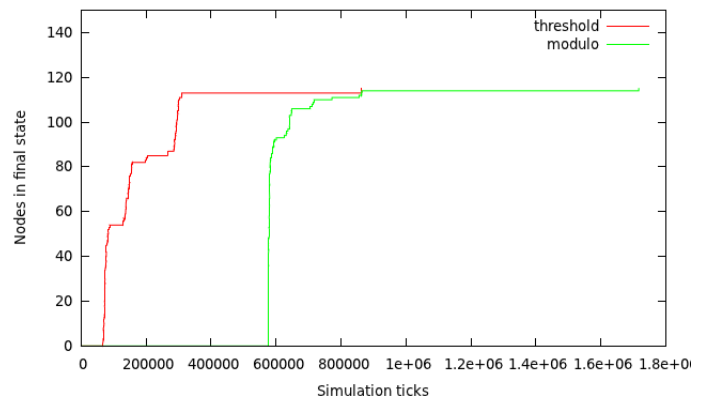


Figure 2. Modulo and threshold: real social topology

[2] <http://ccl.northwestern.edu/netlogo/>.

[3] <http://www.sociopatterns.org/>.

[4] D. Angluin, J. Aspnes, Z. Diamadi, M. J. Fischer, and R. Peralta. Computation in networks of passively mobile finite-state sensors. In *PODC '04: Proceedings of the twenty-third annual ACM symposium on Principles of distributed computing*, pages 290–299, New York, NY, USA, 2004. ACM. ISBN 1-58113-802-4. doi: <http://doi.acm.org/10.1145/1011767.1011810>.

[5] A. Barrat, C. Cattuto, V. Colizza, J. Pinton, W. Van den Broeck, and A. Vespignani. High resolution dynamical mapping of social interactions with active RFID. *ArXiv e-prints*, Nov. 2008.

[6] P. Hui, A. Chaintreau, J. Scott, R. Gass, J. Crowcroft, and C. Diot. Pocket switched networks and the consequences of human mobility in conference environments. In *Proceedings of ACM SIGCOMM first workshop on delay tolerant networking and related topics*, 2005.