Dark Encounter Computations
(and Musings on Biological Networks)
(Extended Abstract)

René Peralta

Information Technology Laboratory, National Institute of Standards and Technology
`rene.peralta@nist.gov`

1 Introduction

Two models of distributed computation are described in which the agents are anonymous finite-state sensors interacting through a communication network whose dynamics (in case the agents are mobile) and/or topology they do not control or even know about. These models were recently introduced in a series of papers by Angluin et al. [3,4,2,5].

In one model, the agents are stationary but form a network of unknown topology. Furthermore, the pair-wise communication links in the network become available in a random and unpredictable infinite sequence. The sensors would like to compute a function of their inputs. While doing so they are, by design, or inherent limitation, not endowed with unique identities (hence the term "dark encounter computations" in the title of this note). Angluin et al. developed a theory about such networks that, among other results, showed that their computational power is dependent on topology. Perhaps contrary to expectation, the power is inversely proportional to the connectivity, with a straight-line configuration being the most powerful and the complete graph configuration being the weakest. This observation leads naturally to a number of questions related to the structure of the network. What topological properties of the network can be computed? (e.g., is the network a ring?). What properties can individual sensor nodes compute? (e.g., am I in a cycle?). These and related questions are answered in the affirmative in [2].

In the second model, agents are randomly mobile and communicate via pairwise interactions. As with the first model, the agents are anonymous and have no control of the sequence of interactions. A stochastic assumption ensures that every pair of agents eventually encounter each other.[1] Angluin et al. show that this model of computation turns out to be surprisingly powerful. In a type of computation called "population protocol" the networks are able to compute predicates such as parity and majority. Population protocols can also perform modular arithmetic (for details, see [4]).

There are, of course, a number of possible definitions of what it means for these networks to "compute" something. In the work of Angluin et al. the network is

---

[1] An example of such a network is RFID devices installed in cars. In this setup, sensors get to exchange information when two cars get sufficiently close to each other.

M. Burmester and A. Yasinsac (Eds.): MADNES 2005, LNCS 4074, pp. 182–185, 2006.

said to *stably compute* a function if all sensors eventually converge to the correct value even when their inputs are allowed to vary an unbounded but finite number of times.

It is natural to wonder how contrived this paradigm is. The question arises because most of these problems become uninteresting if the finite-state and/or the anonymity constraints are removed. In particular, if the sensors are endowed with enough memory, the nodes can assign each other unique identifiers and are then able to simulate a Turing Machine. Anonymity, however, is a plausible constraint in a variety of scenarios. Some of these are listed below. Each is deserving of a lengthy discussion. In this short note, however, I only elaborate on one of them.

- Biological networks (see discussion below);
- Nanotechnology: molecular-sized engineered devices for distributed computation (in a patient's bloodstream, for example) may well have extremely limited computational and (usable) memory resources;
- Limited memory use may be a desirable security property. For example, a captured spying device would ideally "know" very little about what it is doing, where it has been, and what it has seen. In this way, the purpose of the "mission" may remain secret;
- Even when network agents, such as RFID devices, have enough memory to hold unique identifiers, privacy concerns have been raised which may be addressed by designing the devices to be anonymous. This issue is being debated in the context of RFID chips on vehicles (see, for example, the controversy surrounding Texas bill HB2893 `http://www.atsnn.com/story/132110.html`).

2 Biological Networks

In biological networks of lower organisms, and despite the fact that the state space of member agents is typically huge (much larger than the number of agents), we can reasonably expect anonymous communication as a result of evolutionary processes. In their random evolutionary walk through the space of possible configurations and behaviors, it seems highly unlikely that lower organisms such as bacteria would develop the ability to call each other by name. Yet bacteria have been shown to exhibit a remarkable range of social behaviors. For example, they are able to coordinate sporulation in such a way that the collective behavior only occurs if enough of the members "think" it is a good idea. It is unclear to what extent the different known communication mechanisms of bacteria can be simulated (in a way that does not do away with fundamental properties of biological systems) by the networks of Angluin et al. The type of distributed problem-solving we are interested in here has been shown to occur via chemotactic signaling and quorum-sensing. However, these communication mechanisms do not, strictly speaking, involve pairwise interaction. Another communication mechanism, plasmid exchange, does involve pairwise interaction. But I am not aware of this mechanism being used in coordinating real-time bacterial behavior. Nevertheless, distributed computation via pair-wise communication has been documented in neurons (quite surprisingly, as most studied neuronal communication involves broadcasting to all neighboring neurons) [8]. This gives reason to suspect the phenomenon may not be uncommon in biological networks of lower organisms (it is clearly an important component of the social behavior of higher organisms such as humans). In my own observations of pond life through a microscope, I have observed many "networks" of algae which clearly coordinate behavior and communicate via pairwise interactions. I have no expertise, however, to investigate the question of whether the two observed phenomena (coordination and communication through pairwise interaction) are related.

Emerging areas of research aim at describing the fundamental properties of biological networks. Among the most prolific of these efforts are those of Barabasi's group (see `http://www.nd.edu/ alb/`). Here are, without elaboration, some morsels of wisdom from this body of work:

- the World Wide Web exhibits a topology which is not what one would expect. A computer scientist might expect it to be similar to Erdös's random graphs. A social scientist would expect it to be similar to "small-world" networks. As shown in [1], it is neither, and apparently the differences have important implications (with hindsight, one quickly realizes that random graphs in the sense of Erdös and Rényi [6] are not likely to be endowed with life-sustaining properties such as Maturana and Varela's autopoiesis [2]).
- the same structural properties that make the World Wide Web and the Internet resilient to local failures also makes them vulnerable to attack.

A different methodological approach to understanding biological networks seems implicit in Harvard's Bauer Center writings (see `http://cgr.harvard.edu/research/biological.html`). Part of its research mission statement reads as follows

> "What subset of the space of possible networks do biological systems occupy? Can we explain why they occupy this subset? Possible answers to the last question include recycling of historical accidents, strong evolutionary constraints from the combined requirement for robustness and adaptability, and the requirement that any functional module can be constructed by a series of incremental improvements. "

In a recent report by the National Research Council's Committee on Network Science for Future Army Applications (executive summary is available at `http://darwin.nap.edu/execsumm_pdf/11516.pdf`), the Committee concludes that abstract networks are poorly understood and that such understanding is necessary to address the problem of securing computer networks. The report goes on to recommend that a new field, "network science", be funded.

Although the above discussion may seem far removed from network security issues, I contend that it is not. Autopoiesis, for example, views as a defining

---

[2] For a review of autopoiesis see, for example, [7].

characteristic of life that of being able to maintain structure (and hence functionality) while interacting with an environment that is constantly damaging, or outright destroying, components of said structure.

Acknowledgements

I am indebted to Kevin Mills for many helpful suggestions on this manuscript.

References

1. R. Albert, A. Barabasi, and H. Jeong. Scale-free characteristics of random networks: The topology of the world wide web. *Physica A: Statistical Mechanics and its Applications*, 281:69–77, June 2000.
2. Dana Angluin, James Aspnes, Melody Chan, Michael J. Fischer, Hong Jiang, and René Peralta. Stably computable properties of network graphs. In Viktor K. Prasanna, Sitharama Iyengar, Paul Spirakis, and Matt Welsh, editors, *Distributed Computing in Sensor Systems: First IEEE International Conference, DCOSS 2005, Marina del Rey, CA, USE, June/July, 2005, Proceedings*, volume 3560 of *Lecture Notes in Computer Science*, pages 63–74. Springer-Verlag, June 2005.
3. Dana Angluin, James Aspnes, Zoë Diamadi, Michael J. Fischer, and René Peralta. Urn automata. Technical Report YALEU/DCS/TR–1280, Yale University Department of Computer Science, November 2003.
4. Dana Angluin, James Aspnes, Zoë Diamadi, Michael J. Fischer, and René 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. ACM Press, 2004.
5. Dana Angluin, James Aspnes, Michael J. Fischer, and Hong Jiang. Self-stabilizing population protocols. In *Ninth International Conference on Principles of Distributed Systems (pre-proceedings)*, pages 79–90, December 2005.
6. Erdös and Rényi. On the evolution of random graphs. *Publ. Math. Inst. Hung. Acad. Sci.*, 17(5), 1960.
7. Pier Luigi Luisi. Autopoiesis: a review and a reappraisal. *Naturwissenschaften*, 90:49–59, Feb 2003.
8. Elad Schneidman, Michael J. Berry, Ronen Segev, and William Bialek. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature*, April 2006.