What Makes a Distributed Problem Truly Local?

Adrian Kosowski

► To cite this version:

Adrian Kosowski. What Makes a Distributed Problem Truly Local?. SIROCCO 2016 - 23rd International Colloquium on Structural Information and Communication Complexity, Jul 2016, Helsinki, Finland. pp.3, 10.1007/978-3-319-48314-6 . hal-01415090

HAL Id: hal-01415090

<https://hal.inria.fr/hal-01415090>

Submitted on 2 Jan 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

What Makes a Distributed Problem Truly Local?*

Adrian Kosowski

Inria Paris and IRIF, Université Paris Diderot, France

Abstract

In this talk we attempt to identify the characteristics of a task of distributed network computing, which make it easy (or hard) to solve by means of fast local algorithms. We look at specific combinatorial tasks within the LOCAL model of distributed computation, and rephrase some recent algorithmic results in a framework of constraint satisfaction. Finally, we discuss the issue of efficient computability for relaxed variants of the LOCAL model, involving the so-called non-signaling property.

In distributed network computing, autonomous computational entities are represented by the nodes of an undirected *system graph*, and exchange information by sending messages along its edges. A major line of research in this area concerns the notion of *locality*, and asks how much information about its neighborhood a node needs to collect in order to solve a given computational task. In particular, in the seminal LOCAL model [19], the complexity of a distributed algorithm is measured in term of number of *rounds*, where in each round all nodes synchronously exchange data along network links, and subsequently perform individual computations. A t -round algorithm is thus one in which every node exchanges data with nodes at distance at most t (i.e., at most t hops away) from it.

Arguably, the most important class of local computational tasks concerns *symmetry breaking*, and several forms of such tasks have been considered, including the construction of proper *graph colorings* [3–9, 11, 15, 17, 18, 22], of *maximal independent sets* (MIS) [1, 4, 5, 14, 16, 18], as well as edge-based variants of these problems (cf. e.g. [21]). In this talk we address the following question: What makes some symmetry-breaking problems in the LOCAL model easier than others?

We note that the LOCAL model has two flavors, involving the design of deterministic and randomized algorithms, which are clearly distinct [8]. When considering randomized algorithms, for n -node graphs of maximum degree Δ , a hardness separation between the randomized complexities of the specific problems of MIS and $(\Delta + 1)$ -coloring has recently been observed [11, 14]. No analogous separation is as yet known when considering deterministic solutions to these problems. We look at some partial evidence in this direction, making use of the recently introduced framework of *conflict coloring* representations [9]

*This talk includes results of joint work with: P. Fraigniaud, C. Gavoille, M. Heinrich, and M. Markiewicz.

for local combinatorial problems. A conflict coloring representation captures a distributed task through a set of local constraints on edges of the system graph, thus constituting a special case of the much broader class of constraint satisfaction problems (CSP) with binary constraints. Whereas all local tasks are amenable to a conflict coloring formulation, one may introduce a natural *constraint density* parameter, which turns out to be inherently smaller for some problems than for others. For example, for the natural representation of the $(\Delta + 1)$ -coloring task, the constraint density is $1/(\Delta + 1)$, while for any accurate representation of MIS, the constraint density is at least $1/2$. We discuss implications of how low constraint density (notably, much smaller than $1/\Delta$) may be helpful when finding solutions to a distributed task, especially when applying the so-called *shattering method* [20] in a randomized setting, and more directly, when designing faster deterministic algorithms through a direct attack on the conflict coloring representation of the task [9].

We close this talk with a discussion of relaxed variants of the LOCAL model, inspired by the physical concept of non-signaling. In a computational framework, the *non-signaling property* can be stated as the following necessary (but not sufficient) property of the LOCAL model: for any $t > 0$, given two subsets of nodes S_1 and S_2 of the system graph, such that the distance between the nearest nodes of S_1 and S_2 is greater than t , in any t -round LOCAL algorithm, the outputs of nodes from S_1 must be (probabilistically) independent of the inputs of nodes from S_2 . We point out that for a number of symmetry breaking tasks in the LOCAL model, the currently best known asymptotic lower bounds can be deduced solely by exploiting the non-signaling property. This is the case for problems such as MIS [10, 14] or 2-coloring of the ring [10]. On the other hand, such an implication is not true for, e.g., the $\Omega(\log^* n)$ lower bound on the number of rounds required to 3-color the ring [15] — this lower bound follows from different (stronger) properties of the LOCAL model [12, 13]. This leads us to look at the converse question: How to identify conditions under which non-signaling solutions to a distributed task can be converted into an algorithm in the LOCAL model? We note some progress in this respect for quantum analogues of the LOCAL model [2].

References

- [1] N. Alon, L. Babai, and A. Itai. A fast and simple randomized parallel algorithm for the maximal independent set problem. *J. Algorithms* **7**(4):567-583 (1986)
- [2] P. Arrighi, V. Nesme, and R.F. Werner. Unitarity plus causality implies localizability. *J. Comput. Syst. Sci.* **77**(2):372-378 (2011)
- [3] L. Barenboim. Deterministic $(\Delta + 1)$ -coloring in sublinear (in Δ) Time in Static, Dynamic and Faulty Networks. In *Proc. 34th ACM Symp. on Principles of Distributed Computing (PODC)*, pp. 345-354, 2015.
- [4] L. Barenboim and M. Elkin. Distributed $(\Delta + 1)$ -coloring in linear (in Δ) time. In *Proc. 41th ACM Symp. on Theory of Computing (STOC)*, pp. 111-120, 2009.

- [5] L. Barenboim and M. Elkin. *Distributed Graph Coloring: Fundamentals and Recent Developments*. Synthesis Lectures on Distributed Computing Theory, Morgan & Claypool Publishers, 2013.
- [6] L. Barenboim, M. Elkin, and F. Kuhn. Distributed $(\Delta + 1)$ -coloring in linear (in Δ) time. *SIAM J. Comput.* **43**(1):72-95 (2014)
- [7] L. Barenboim, M. Elkin, S. Pettie, and J. Schneider. The locality of distributed symmetry breaking. In Proc. *53rd IEEE Symp. on Foundations of Computer Science* (FOCS), pp. 321-330, 2012.
- [8] Y-J. Chang, T. Kopelowitz, and S. Pettie. An Exponential Separation Between Randomized and Deterministic Complexity in the LOCAL Model, In Proc. *57th IEEE Symp. on Foundations of Computer Science* (FOCS), to appear, 2016, <http://arxiv.org/abs/1602.08166>.
- [9] P. Fraigniaud, M. Heinrich, and A. Kosowski. Local Conflict Coloring, In Proc. *57th IEEE Symp. on Foundations of Computer Science* (FOCS), to appear, 2016, <http://arxiv.org/abs/1511.01287>.
- [10] C. Gavoille, A. Kosowski, and M. Markiewicz. What Can Be Observed Locally? In Proc. *23rd International Symposium on Distributed Computing* (DISC), pp. 243-257, 2009.
- [11] D.G. Harris, J. Schneider, and H-H. Su. Distributed $(\Delta + 1)$ -Coloring in Sublogarithmic Rounds, In Proc. *48th Annual Symposium on the Theory of Computing* (STOC), pp. 465-478, 2016.
- [12] A.E. Holroyd and T.M. Liggett. Finitely dependent coloring. Submitted preprint: <http://arxiv.org/abs/1403.2448>.
- [13] A.E. Holroyd and T.M. Liggett. Symmetric 1-dependent colorings of the integers. *Electron. Commun. Probab.* **20**(31) (2015).
- [14] F. Kuhn, T. Moscibroda, and R. Wattenhofer. What cannot be computed locally! In Proc. *23rd ACM Symp. on Principles of Distributed Computing* (PODC), pp. 300-309, 2004.
- [15] N. Linial. Locality in Distributed Graph Algorithms. *SIAM J. Comput.* **21**(1): 193-201 (1992)
- [16] M. Luby. A simple parallel algorithm for the maximal independent set problem. *SIAM J. on Computing* **15**:1036-1053 (1986)
- [17] M. Naor. A lower bound on probabilistic algorithms for distributive ring coloring. *SIAM J. on Discrete Mathematics* **4**(3):409-412 (1991)
- [18] A. Panconesi and A. Srinivasan. Improved Distributed Algorithms for Coloring and Network Decomposition Problems. In Proc. *24th ACM Symp. on Theory of Computing* (STOC), pp. 581-592, 1992.
- [19] D. Peleg. *Distributed Computing: A Locality-Sensitive Approach*. SIAM, Philadelphia, PA, 2000.

- [20] J. Schneider, R. Wattenhofer. A new technique for distributed symmetry breaking. In *Proc. 29th ACM Symp. on Principles of Distributed Computing (PODC)*, pp. 257-266, 2010.
- [21] J. Suomela. Survey of local algorithms. *ACM Comput. Surv.* **45**(2):24 (2013)
- [22] M. Szegedy and S. Vishwanathan: Locality based graph coloring. In *Proc. 25th ACM Symp. on Theory of Computing (STOC)*, pp. 201-207, 1993.