

Landscape of Locality

Jukka Suomela 

Aalto University, Finland

<https://jukkasuomela.fi/>

jukka.suomela@aalto.fi

Abstract

The *theory of distributed computing* aims at understanding which tasks can be solved efficiently in large distributed systems. This forms the basis for our understanding of the modern world, which heavily depends on world-wide communication networks and large-scale distributed computer systems.

In distributed computing the key computational resource is *communication*, and we seek to find out which computational problems can be solved with only a few communication steps. This is directly connected to the concept of *locality*: in T synchronous communication rounds, all nodes in a network can gather all information in their radius- T neighborhoods, but not any further. Hence the *distributed time complexity* of a graph problem can be defined in two equivalent ways: it is the number of communication rounds needed to solve the problem, and it is the distance up to which individual nodes need to see in order to choose their own part of the solution.

While the locality of graph problems has been studied already since the 1980s, only in the past four years we have started to take big leaps in understanding what the landscape of distributed time complexity looks like and with what kind of tools and techniques we can study it.

One concept that has been a driving force in the recent developments is the notion of *locally verifiable* problems. These are graph problems in which a solution is feasible if and only if it looks valid in all constant-radius neighborhoods; put otherwise, these are problems that could be solved efficiently with a *nondeterministic* distributed algorithm, and hence they form a natural distributed analogue of class NP. Now the key question is this: if a problem is locally verifiable, is it also locally solvable, and if not, what can we say about its distributed time complexity?

Naor and Stockmeyer [SIAM J. Comput. 1995] formalized the idea of locally verifiable problems by introducing the class of LCL problems (locally checkable labeling problems). While the concept is old, and over the years we have seen results related to the locality of many *specific* LCLs, little was known about the distributed complexity of LCLs *in general*. By 2015, we had only seen examples of LCLs with localities $O(1)$, $\Theta(\log^* n)$, and $\Theta(n)$, and it was wide open whether these three classes are all that there is.

All this started to change rapidly after we proved [Brandt et al., STOC 2016] that there are natural examples of LCLs that have a locality strictly between $\omega(\log^* n)$ and $o(n)$. The same paper also paved the way for the development of a new general-purpose proof technique for analyzing the locality of locally verifiable problems, namely *round elimination*.

Now after four years of work and a number of papers by several research teams working in the area, we have reached a point in which there is a *near-complete picture of the landscape of LCL problems* – and it looks nothing like what we would have expected.

2012 ACM Subject Classification Theory of computation → Distributed computing models; Theory of computation → Complexity classes

Keywords and phrases Theory of distributed computing, Network algorithms, Locality, Distributed time complexity

Digital Object Identifier 10.4230/LIPIcs.SWAT.2020.2

Category Invited Talk



© Jukka Suomela;

licensed under Creative Commons License CC-BY

17th Scandinavian Symposium and Workshops on Algorithm Theory (SWAT 2020).

Editor: Susanne Albers; Article No. 2; pp. 2:1–2:1

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany