Computing Power of Hybrid Models in Synchronous Networks

Pierre Fraigniaud[1] ✉
IRIF, Université Paris Cité and CNRS, France.

Pedro Montealegre[2] ✉
Facultad de Ingeniería y Ciencias, Universidad Adolfo Ibáñez, Santiago, Chile.

Pablo Paredes ✉
Departamento de Ingeniería Matemática, Universidad de Chile, Chile

Ivan Rapaport ✉
DIM-CMM (UMI 2807 CNRS), Universidad de Chile, Chile.

Martín Ríos-Wilson ✉
Facultad de Ingeniería y Ciencias, Universidad Adolfo Ibáñez, Santiago, Chile.

Ioan Todinca ✉
LIFO, Université d'Orléans and INSA Centre-Val de Loire, France.

Abstract

During the last two decades, a small set of distributed computing models for networks have emerged, among which LOCAL, CONGEST, and Broadcast Congested Clique (BCC) play a prominent role. We consider *hybrid* models resulting from combining these three models. That is, we analyze the computing power of models allowing to, say, perform a constant number of rounds of CONGEST, then a constant number of rounds of LOCAL, then a constant number of rounds of BCC, possibly repeating this figure a constant number of times. We specifically focus on 2-round models, and we establish the complete picture of the relative powers of these models. That is, for every pair of such models, we determine whether one is (strictly) stronger than the other, or whether the two models are incomparable.

2012 ACM Subject Classification Theory of computation → Distributed algorithms

Keywords and phrases hybrid model, synchronous networks, LOCAL, CONGEST, Broadcast Congested Clique

Digital Object Identifier 10.4230/LIPIcs.CVIT.2016.23

1 Introduction

This paper analyzes the relative power of distributed computing models for networks, all resulting from the combination of standard synchronous models such as LOCAL and CONGEST [4], as well as Broadcast Congested Clique (BCC) [1]. Each of these three models has its strengths and limitations. We investigate the power of models resulting from combining these three models, in order to take advantage of their positive aspects without suffering from their negative ones.

For the sake of comparing models, we focus on the standard framework of distributed *decision* problems on labeled graphs (see [2]). Such problems are defined by a collection $\mathcal{L}$ of pairs (G, ℓ), where $G = (V, E)$ is a graph, and $\ell : V \to \{0, 1\}^*$ is a function assigning

* Additional support for ANR projects QuData and DUCAT.
† This work was supported by Centro de Modelamiento Matemático (CMM), ACE210010 and FB210005, BASAL funds for centers of excellence from ANID-Chile, FONDECYT 11190482, FONDECYT 1220142 and PAI77170068

42nd Conference on Very Important Topics (CVIT 2016).
Editors: John Q. Open and Joan R. Access; Article No. 23; pp. 23:1–23:3
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

a label $\ell(u) \in \{0,1\}^*$ to every $u \in V$. Such a set $\mathcal{L}$ is called a distributed *language*. A distributed algorithm A *decides* $\mathcal{L}$ if every node running A eventually accepts or rejects, and the following condition is satisfied: for every labeled graph (G,ℓ), every node should accept in a yes-instance (i.e., an instance $(G,\ell) \in \mathcal{L}$), and, in a no-instance (i.e., an instance $(G,\ell) \notin \mathcal{L}$), at least one node must reject.

For every $t \geq 0$, let us denote by $\mathbf{L}^t$ the set of distributed languages $\mathcal{L}$ for which there is a t-round algorithm in the LOCAL model deciding $\mathcal{L}$. The sets $\mathbf{C}^t$ and $\mathbf{B}^t$ are defined similarly, for the CONGEST and BCC models, respectively. Note that while it is easy to show, using indistinguishability arguments, that, for every $t \geq 1$, $\mathbf{L}^t \smallsetminus \mathbf{L}^{t-1} \neq \varnothing$ and $\mathbf{C}^t \smallsetminus \mathbf{C}^{t-1} \neq \varnothing$, establishing that there is indeed a decision problem in $\mathbf{B}^t \smallsetminus \mathbf{B}^{t-1}$ requires significantly more work [3]. Also, we define $\mathbf{L}^* = \cup_{t\geq 0}\mathbf{L}^t$, $\mathbf{C}^* = \cup_{t\geq 0}\mathbf{C}^t$, and $\mathbf{B}^* = \cup_{t\geq 0}\mathbf{B}^t$. So, in particular, $\mathbf{L}^*$ is the class of distributed languages that can be decided in a constant number of rounds in the LOCAL model.

2 Our Results

On the negative side, we provide a series of separation results between 2-round hybrid models. In particular, we show that **BC** and **CB** are incomparable. That is, there are languages in $\mathbf{BC} \smallsetminus \mathbf{CB}$, and languages in $\mathbf{CB} \smallsetminus \mathbf{BC}$. In fact, we show stronger separation results, by establishing that $\mathbf{BC} \smallsetminus \mathbf{C}^*\mathbf{B} \neq \varnothing$, and $\mathbf{CB} \smallsetminus \mathbf{BL}^* \neq \varnothing$. That is, in particular, there are languages that can be decided by a 2-round algorithm performing a single BCC round followed by one CONGEST round, which cannot be decided by any algorithm performing k CONGEST rounds followed by a single BCC round, for any $k \geq 1$.

On the positive side, we show that, for any non-negative integers $\alpha_1,\dots,\alpha_k,\beta_1,\dots,\beta_k$,

$$\prod_{i=1}^{k} \mathbf{L}^{\alpha_i}\mathbf{B}^{\beta_i} \subseteq \mathbf{L}^{\sum_{i=1}^{k}\alpha_i}\mathbf{B}^{\sum_{i=1}^{k}\beta_i}. \tag{1}$$

That is, if a language $\mathcal{L}$ can be decided by a t-round algorithm alternating LOCAL and BCC rounds, then $\mathcal{L}$ can be decided by a t-round algorithm performing all its LOCAL rounds first, and then all its BCC rounds — with the notations of Eq. (1), $t = \sum_{i=1}^{k}(\alpha_i + \beta_i)$. So, in particular $\mathbf{BL} \subseteq \mathbf{LB}$. This inclusion is strict, since, as said before, $\mathbf{CB} \smallsetminus \mathbf{BL}^* \neq \varnothing$. In fact, this separation holds even if the number of LOCAL rounds depends on the number of nodes n in the network, as long as the algorithm performs $o(n)$ LOCAL rounds after its BCC round. Another consequence of Eq. (1) is that the largest class of languages among all the ones considered in this paper is $\mathbf{L}^*\mathbf{B}^*$, that is, languages that can be decided by algorithms performing k LOCAL rounds followed by k' BCC rounds, for some $k \geq 0$ and $k' \geq 0$.

Interestingly, our separation results hold even for randomized protocols, which can err with probability at most $\epsilon \leq 1/5$. That is, in particular, there is a language $\mathcal{L} \in \mathbf{CB}$ (i.e., that can be decided by a deterministic 2-round algorithm) which cannot be decided with error probability at most $1/5$ by any randomized algorithm performing one BCC round first, followed by k LOCAL rounds, for any $k \geq 1$. All our results about 2-rounds hybrid models are summarized on Figure 1.

3 Our Techniques.

All our separation results are obtained by reductions from communication complexity lower bounds. However, we had to revisit several known communication complexity results for adapting them to the setting of distributed decision, in which no-instances may be rejected

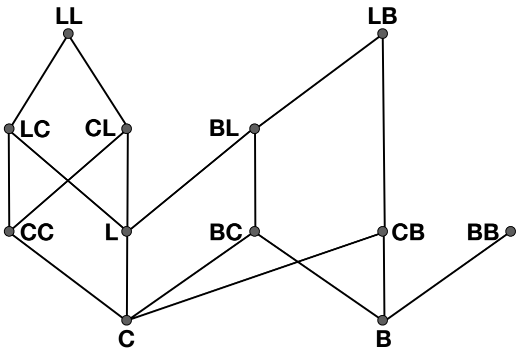

Figure 1 *The lattice of 2-round hybrid models. An edge between a set of languages* $\mathbf{S}_1$ *and a set* $\mathbf{S}_2$*, where* $\mathbf{S}_1$ *is at a level lower than* $\mathbf{S}_2$*, indicates that* $\mathbf{S}_1 \subseteq \mathbf{S}_2$*. In fact, all inclusions are strict. Transitive edges are not displayed. Two sets that are not connected by a monotone path are incomparable. For instance,* **CB** *and* **BL** *are incomparable, while* $\mathbf{BC} \subseteq \mathbf{LB}$.

by a single node, and non necessarily by all the nodes. In particular, we revisit the classical `Index` problem. Recall that, in this problem, Alice is given a binary vector $x \in \{0,1\}^n$, Bob is given an index $i \in [n]$, and Bob must output x_i based on a single message received from Alice (1-way communication). We define the `XOR-Index` problem, in which Alice is given a binary vector $x \in \{0,1\}^n$ together with an index $i \in [n]$, Bob is given a binary vector $y \in \{0,1\}^n$ together with an index $j \in [n]$, and, after a single round of 2-way communication, Alice must output a boolean out_A and Bob must output a boolean out_B, such that

$$\mathrm{out}_A \wedge \mathrm{out}_B = x_j \oplus y_i.$$

That is, if $x_j \neq y_i$ then Alice and Bob must both accept (i.e., output *true*), and if $x_j = y_i$ then *at least* one of these two players must reject (i.e., output *false*). We show that the sum of the sizes of the message sent by Alice to Bob and the message sent by Bob to Alice is $\Omega(n)$ bits. This bound holds even if the communication protocol is randomized and may err with probability at most $1/5$, and even if the two players have access to shared random coins.

The fact that only one of the two players may reject a no-instance (i.e., an instance where $x_j \oplus y_i = 0$), and not necessarily both, while a yes-instance must be accepted by both players, yields an asymmetry which complicates the analysis. We use information theoretic tools for establishing our lower bound. Specifically, we identify a way to decorrelate the behaviors of Alice and Bob, so that to analyze separately the distribution of decisions taken by each player, and then to recombine them for lower bounding the probability of error in case the messages exchanged between the players are small, contradicting the fact that this error probability is supposed to be small.

References

1. Andrew Drucker, Fabian Kuhn, and Rotem Oshman. On the power of the congested clique model. In *Proceedings of the 2014 ACM Symposium on Principles of Distributed Computing*, pages 367–376, 2014.
2. Laurent Feuilloley and Pierre Fraigniaud. Survey of distributed decision. *Bull. EATCS*, 119, 2016.
3. Noam Nisan and Avi Widgerson. Rounds in communication complexity revisited. In *Proceedings of the twenty-third annual ACM symposium on Theory of computing*, pages 419–429, 1991.
4. David Peleg. *Distributed Computing: A Locality-Sensitive Approach.* SIAM, 2000.