

# Brief Announcement: Relations Between Space-Bounded and Adaptive Massively Parallel Computations

Michael Chen

Iowa State University, Ames, IA, USA

A. Pavan

Iowa State University, Ames, IA, USA

N. V. Vinodchandran

University of Nebraska–Lincoln, NE, USA

---

## Abstract

In this work, we study the class of problems solvable by (deterministic) Adaptive Massively Parallel Computations in constant rounds from a computational complexity theory perspective. A language  $L$  is in the class  $\text{AMPC}^0$  if, for every  $\varepsilon > 0$ , there is a deterministic AMPC algorithm running in constant rounds with a polynomial number of processors, where the local memory of each machine  $s = O(N^\varepsilon)$ . We prove that the space-bounded complexity class  $\text{ReachUL}$  is a *proper subclass* of  $\text{AMPC}^0$ . The complexity class  $\text{ReachUL}$  lies between the well-known space-bounded complexity classes Deterministic Logspace (DLOG) and Nondeterministic Logspace (NLOG). In contrast, we establish that it is unlikely that PSPACE admits AMPC algorithms, even with polynomially many rounds. We also establish that showing PSPACE is a subclass of *nonuniform-AMPC* with polynomially many rounds leads to a significant separation result in complexity theory, namely PSPACE is a proper subclass of  $\text{EXP}^{\Sigma_2^b}$ .

**2012 ACM Subject Classification** Computing methodologies → Massively parallel algorithms; Theory of computation → Complexity classes

**Keywords and phrases** Massively Parallel Computation, AMPC, Complexity Classes, LogSpace, NL, PSPACE

**Digital Object Identifier** 10.4230/LIPIcs.DISC.2023.37

**Funding** This work is supported in part by NSF grants 1934884, 2130536, and 2130608.

**Acknowledgements** The authors thank Sriram Pemmaraju and Meena Mahajan for their helpful discussions. We thank anonymous reviewers for their valuable comments and pointers to some critical references.

## 1 Introduction

The *Massively Parallel Computation* (MPC) model is widely accepted as the standard theoretical model for distributed computation frameworks such as MapReduce, Spark, Hadoop, FlumeJava, Beame, Pregel, and Gigraph [7, 9]. It was defined in [5], and it captures computation on large data: data is adversarially distributed to processors, and each processor has local memory  $s = O(N^\varepsilon)$  ( $0 < \varepsilon < 1$  where  $N$  is the input size). Computation occurs in rounds, and in each round, every machine performs computation based on its local data and then communicates with other machines with the constraint that the amount of communication by a process is equal to that of its local memory  $s$ . A salient feature of the MPC model is that no computational restriction is placed on the processor, except that each processor has local memory  $s$ , and a key objective is to minimize the number of rounds.



© Michael Chen, A. Pavan, and N. V. Vinodchandran;  
licensed under Creative Commons License CC-BY 4.0  
37th International Symposium on Distributed Computing (DISC 2023).  
Editor: Rotem Oshman; Article No. 37; pp. 37:1–37:7



Leibniz International Proceedings in Informatics  
LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Ideally, one would like to design an algorithm with constant rounds with a small number of processors. The MPC model has been extensively studied in the context of designing algorithms as well as its relationship with complexity classes [2, 3, 4, 7, 5, 6, 16].

Recent work of [7] introduced an adaptive extension of the MPC model called *Adaptive Massively Parallel Computation* model (AMPC). In the AMPC model, the processors communicate via a shared memory called Distributed Data Stores (DDS) by reading from and writing to the DDS. In a single round, a machine can adaptively query the DDS to obtain  $s$  words and write at most  $s$  words, and as in the case of MPC,  $s$  is  $O(N^\varepsilon)$ . In [7], authors designed a constant round randomized AMPC algorithm for 1V2-CYCLE as well as a few other graph problems.

In this work, we report progress on the power and limitation of the AMPC model from a computational complexity theory viewpoint. Towards this, we define a robust complexity class which we denote by  $\text{AMPC}^0$ . A language  $L$  is in  $\text{AMPC}^0$  if for every  $\varepsilon$  there is a constant round (depending on  $\varepsilon$ ) AMPC algorithm with  $P = p(N)$  many processors (where  $p(\cdot)$  is a polynomial) each with  $s = O(N^\varepsilon)$  memory. We define a similar complexity class  $\text{AMPC}^{\text{poly}}$  where the number of rounds is polynomial. We study the relationship of these AMPC complexity classes with respect to the standard space-bounded complexity classes DLOG, NLOG, and PSPACE. The starting point of our work is that the ideas from the randomized AMPC algorithm for 1V2-CYCLE from [7] can be used to show that the complexity class DLOG is a subset of (uniform)  $\text{AMPC}^0$ . Motivated by this, we explore whether NLOG is a subset of  $\text{AMPC}^0$ . We make progress toward this question by studying a complexity class ReachUL [8, 1, 11]. This is a natural complexity class that lies between DLOG and NLOG and has been studied earlier in the context of designing space-efficient algorithms for reachability that beat the Savitch's bound [1]. We prove that ReachUL is a subset of (uniform)  $\text{AMPC}^0$ . More interestingly, we show that ReachUL is a *proper subset* of (uniform)  $\text{AMPC}^0$ . On the contrary, we observe that it is unlikely that the whole of PSPACE (or even NP) can be solved even in  $\text{AMPC}^{\text{poly}}$ . This is because every language that admits (uniform)  $\text{AMPC}^{\text{poly}}$  algorithm can be solved in subexponential time. Since we do not believe that PSPACE can be solved in subexponential time, we obtain that it is unlikely that  $\text{PSPACE} \subseteq \text{AMPC}^{\text{poly}}$ . We also consider the limitation of nonuniform  $\text{AMPC}^{\text{poly}}$ . We unconditionally show that there exist languages in  $\text{E}^{\Sigma_2^P}$  that are not in  $\text{AMPC}^{\text{poly}}$ . We note that the work reported in [15] also considered the relations of complexity classes such as DLOG and NLOG to MPC model.

► **Remark.** In an algorithmic setting, it is typically desired that the total memory of an AMPC algorithm  $P \cdot s$  to be  $N \cdot \text{poly log}(N)$ . However, to define a robust complexity class (closed under reductions), we allow  $P$  to be polynomial and require that for every  $0 < \varepsilon < 1$ , there is an algorithm with  $s$  local memory per processor.

## 2 Preliminaries

We now give the formal description of the AMPC model [7, 9]. Let  $p(\cdot)$ ,  $s(\cdot)$  and  $r(\cdot)$  are functions from  $\mathbb{N}$  to  $\mathbb{N}$ . An  $\text{AMPC}[p(N), s(N), r(N)]$  algorithm for length  $N$ , is a collection of processors  $M_{i,j}$ ,  $1 \leq i \leq p(N)$  and  $1 \leq j \leq r(N)$  where each processor has a memory bound of  $s(N)$ . In addition to the processors there is a collection of *Distributed Data Stores* (DDS) denoted by  $\mathcal{D}_0, \mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_{r(N)}$ . For each DDS, the data is stored in a bit addressable manner (as done in [9]) i.e., a collection of key-value pairs in the form of  $(i, i^{\text{th}}$  bit of DDS). The input string  $x = x_1 \dots x_N$  is stored in  $\mathcal{D}_0$  in the form of  $\{(i, x_i)\}_{i=1}^N$ . The computation occurs in rounds. The processors  $M_{i,j}$ ,  $1 \leq i \leq p(N)$  participate in the  $j$ th round. In the  $j$ th round, each of these processors is allowed to make  $s(N)$  adaptive queries to read from

$\mathcal{D}_{j-1}$  and each processor is allowed to write up to  $s(N)$  bits to  $\mathcal{D}_j$ . The computation stops after  $r(N)$  rounds, and we say that the algorithm accepts string  $x$  if the value of key 1 in  $\mathcal{D}_{r(N)}$  is 1.

Inherently this is a nonuniform model of computation. A language  $L$  is in the class (nonuniform)  $\text{AMPC}^0$  if for every  $0 < \varepsilon < 1$ , there exists a polynomial  $p(\cdot)$ , and a constant  $r = r(N) > 0$  such that for every input length  $N \geq 0$ , there is a  $\text{AMPC}[p(N), N^\varepsilon, r]$  algorithm that accepts  $L$  on strings at length  $N$ . We define the uniform  $\text{AMPC}$  model. This definition is similar to the uniform MRC model as defined in [10]. For an algorithm  $P$ , we use  $P_{i,j}$  to denote a processor whose behavior is the same as  $P$  on inputs  $i$  and  $j$ . A language  $L$  is in the class (uniform)  $\text{AMPC}^0$  if for every  $\varepsilon > 0$ , there exists a polynomial  $p(\cdot)$  and a constant  $r = r(N)$ , and a logspace bounded algorithm  $U$  that on input  $1^N$  outputs the code of a processor  $P$  with the following properties: the processors  $P_{i,j}$   $1 \leq i \leq p(N)$ ,  $1 \leq j \leq r$  constitute a  $\text{AMPC}[p(N), N^\varepsilon, r]$  algorithm that accepts  $L$  at strings of length  $N$ . Analogously we define uniform and nonuniform versions of the class  $\text{AMPC}^{\text{poly}}$  where the number of rounds is allowed to be a polynomial. In the rest of the document, we write  $\text{AMPC}^0$  to denote (uniform)  $\text{AMPC}^0$ .

We use  $\text{DLOG}$  (resp.  $\text{PSPACE}$ ) to denote the class of languages accepted by deterministic logspace (resp. polynomial-space) Turing machines. The complexity class  $\text{E}$  is the class of languages that are accepted by deterministic  $2^{O(N)}$ -time bounded machines, and  $\Sigma_2^{\text{P}}$  denote the class of languages in the second level of the polynomial-hierarchy. A language  $L$  is in the class  $\text{SubEXP}$ , if for every  $\varepsilon > 0$ , there is a  $O(2^{N^\varepsilon})$ -time bounded machine that accepts  $L$ . A language  $L$  is in  $\text{NC}^1$  if  $L$  can be decided by a family of circuits  $\{C_N\}_{N \in \mathbb{N}}$  where  $C_N$  has  $\text{poly}(N)$  size and  $O(\log N)$  depth.

► **Definition 1** ([8, 1]). *A nondeterministic machine is called reach-unambiguous if for every configuration  $C$ , there is at most one path from the start configuration to  $C$ . The class  $\text{ReachUL}$  is the class of languages that are accepted by  $O(\log N)$ -space-bounded reach-unambiguous machines.*

► **Definition 2** ( $\text{REACH-UNAMBIGUOUS}$ ).  *$\text{REACH-UNAMBIGUOUS}$  is the language consisting of tuples  $\langle G, a, b \rangle$  such that (1)  $G = (V, E)$  is a directed graph, (2) for all  $u \in V$  there exists at most 1 directed path from  $a$  to  $u$  and, (3) there exists a directed path from  $a$  to  $b$ .*

It is known that  $\text{REACH-UNAMBIGUOUS}$  is complete for  $\text{ReachUL}$  with respect to logspace reductions [13, 8, 1].

## 3 Results

### 3.1 ReachUL in $\text{AMPC}^0$

We show that  $\text{ReachUL}$  is a proper subset of  $\text{AMPC}^0$ . We start with the following theorem.

► **Theorem 3.**  $\text{DLOG} \subsetneq \text{AMPC}^0$ . *That is,  $\text{DLOG}$  is a proper subset of  $\text{AMPC}^0$ .*

► **Corollary 4.**  $\text{AMPC}^0$  is closed under logspace reductions.

Inclusion of Theorem 3 follows from [7]. The authors showed a randomized constant round  $\text{AMPC}$  algorithm for the  $\text{DLOG}$ -complete problem,  $1\text{V}2\text{-CYCLE}$ . Their algorithm can be modified to obtain a deterministic algorithm by allowing for  $O(N)$  processors and using  $O(N^{1+\varepsilon})$  total memory for any  $\varepsilon$ . The strictness of inclusion follows from Theorem 8 which we prove.

Let  $\langle G, a, b \rangle$  be an input instance of REACH-UNAMBIGUOUS where  $G = (V, E)$  is a reach-unambiguous graph such that  $V = \{v_1, \dots, v_N\}$  and  $a, b \in V$ . Without loss of generality, we assume that the out-degree of each vertex is at most 2. For  $u \in V$  and  $s \in \mathbb{N}$ , define  $T_s(u)$  to be the tree resulting from a *Breadth First Search* (BFS) starting from  $u$  in  $G$  upto  $s$  nodes such that no node is partially visited, i.e., either all the children of any vertex are in the tree, or none of them are. We shall call every node other than  $u$  in  $T_s(u)$  a descendant of  $u$ . The main ingredient in our proof is Algorithm 1 that constructs a compressed version of  $T_s(u)$ . This algorithm is based on the tree contraction idea [1]. We note that recently tree contraction has been studied in the context AMPC in [12]. For any graph  $G$ , we often overload the notation  $G$  to also refer to the vertex set of the graph.

► **Definition 5.** Let  $u \in V$ , and  $v$  be a descendant of  $u$  in  $T_s(u)$ .  $v$  is said to be an *intermediate vertex* for  $T_s(u)$  if there exists an edge  $(v, w) \in E$  such that  $w \notin T_s(u)$ . Define  $I_s(u) \subseteq T_s(u)$  as the set of vertices that are intermediate for  $T_s(u)$ . We say that  $T_s(u)$  is complete if  $I_s(u) = \emptyset$ , otherwise  $T_s(u)$  is incomplete.

Intermediate vertices capture the idea of vertices that can still be explored. If  $v$  is an intermediate vertex for  $T_s(u)$ , that means  $v$  can still be further explored. But due to the BFS parameter  $s$ , it could not explore  $v$  any further. We shall assume for simplicity of the analysis that if a tree  $T_s(u)$  is incomplete, the tree has exactly  $s + 1$  vertices (in general, such a tree could have either  $s$  or  $s + 1$  vertices). Thus the condition  $|T_s(u)| < s + 1$  denotes the condition that  $T_s(u)$  is complete.

■ **Algorithm 1** Construct Algorithm.

---

```

1 Function Construct( $u, s$ ):
2   Compute  $T_s(u)$  using at most  $O(s)$  queries.
3   if  $b \in T_s(u)$  then
4     //  $b$  can be reached from  $u$  within  $s$  queries
4      $T'_s(u) \leftarrow (\{b\}, \emptyset)$ 
5   else if  $|T_s(u)| < s + 1$  then
6     //  $b$  cannot be reached from  $u$ 
6      $T'_s(u) \leftarrow (\{u\}, \emptyset)$ 
7   else
8     //  $b$  cannot be reached from  $u$  within  $s$  queries, need to explore
8     Compute  $I_s(u)$  using at most  $O(s)$  queries
9      $T'_s(u) \leftarrow$  A complete binary tree whose leaves are exactly  $I_s(u)$ 
10  Write  $T'_s(u)$  to the DDS

```

---

Let  $T'_s(u)$  be the output of *Construct*( $u, s$ ) in Algorithm 1.  $T'_s(u)$  is a contracted version of  $T_s(u)$ . If  $b \in T_s(u)$  or  $|T_s(u)| < s + 1$ , the search from  $u$  is completed, and we can contract the tree to a single node. Otherwise, the tree is contracted to a complete binary tree whose leaves are  $I_s(u)$ , which are precisely the candidates that can lead to  $b$ . Locally, it is possible that  $T'_s(u)$  does not contract. Claim 6 shows that globally the contraction will occur.

Define the tree  $T'$  generated by starting with  $T'_s(a)$ , and recursively substituting every leaf  $l \in T'$  with  $T'_s(l)$ . Continuing the process until substituting leaves does not change the tree. This graph has the property that  $\langle T', a, b \rangle \in \text{REACH-UNAMBIGUOUS} \iff \langle G, a, b \rangle \in \text{REACH-UNAMBIGUOUS}$  since the only vertices that remain are those vertices that have the potential to reach  $b$ . For an AMPC model,  $T'$  need not be explicitly constructed since each

tree is locally computed and is then updated in the DDS. We shall now show that the graph size reduces by a factor of  $s/2$ , which is sufficient to get the algorithm to halt in constant rounds by setting  $s = O(N^\epsilon)$ .

▷ **Claim 6.**  $|T'| \leq 2N/s$

Given  $u \in V$ , for analysis' sake construct  $H_s(u)$  by making every descendant in  $T_s(u)$  a child of  $u$ , i.e.  $H_s(u)$  is a re-arranged version of  $T_s(u)$  such that edges go from  $u$  to the descendants of  $u$  in  $T_s(u)$ .

We now construct an  $s$ -ary tree,  $H$ , such that it is always full (vertices either have out degree 0 or  $s$ ). Start with  $H_s(a)$ , then for every leaf  $l$  whose parent is  $p$  such that  $l \in I_s(p)$  substitute  $l$  with  $H_s(l)$  if  $T_s(p) = s + 1$ . If the BFS search was incomplete, substitute it with  $b$  if  $b \in H_s(l)$ , otherwise, do nothing. Repeat the process until no more substitutions can be done. This construction leads to a full  $s$ -ary tree  $H$ , such that  $|H| \leq N$ .  $H$  represents a BFS traversal done in “batches” of size  $s$ , where only intermediate nodes are substituted with another  $s$ -ary tree. Exploring non-intermediate nodes would be redundant. Let  $i$  denote the number of internal nodes of this  $s$ -ary tree.

Proof of Claim 6. Since  $H$  is a full  $s$ -ary tree,  $i = |H|/s \leq N/s$ . And  $H$  is essentially a rearrangement of  $T'$  such that internal vertices in  $H$  correspond to intermediate vertices in  $T'$ . However, due to line 9 of Algorithm 1, we may be adding more vertices, but however, it is no more than twice. i.e. we have  $|T'| \leq 2i$ . Therefore we have  $|T'| \leq 2N/s$  ◁

► **Theorem 7.** REACH-UNAMBIGUOUS  $\in$  AMPC<sup>0</sup>

**Proof.** Let  $\langle G, a, b \rangle$  be a problem instance of REACH-UNAMBIGUOUS with  $G = (V, E)$  such that  $V = \{v_1, \dots, v_N\}$ . Fix  $\epsilon \in (0, 1)$ . Define the AMPC algorithm with  $s = O(N^\epsilon)$  local memory. Assign  $G_1 \leftarrow G$  and  $R \leftarrow O(1/\epsilon)$ , for  $i = 1, \dots, R$  rounds do the following, for each  $v \in G_i$ , a machine executes  $Construct(v, s)$  for  $G_i$  to get a new graph  $T'$  as described earlier. Assign  $G_{i+1} \leftarrow T'$ . After the rounds are complete,  $|G_R| = \frac{N}{(s/2)^R} = O(1)$ . Then a single machine can perform normal reachability on  $G_R$ , accepting the input if and only if there is a path from  $a$  to  $b$ . ◀

► **Theorem 8.** For every  $c \in \mathbb{N}$ . There exists a problem in AMPC<sup>0</sup> that is not in DSPACE( $\log^c N$ )

**Proof.** Let  $A \in \text{DSPACE}(N^2)$  but  $A \notin \text{DSPACE}(N)$ . We know such a problem exists due to the space hierarchy theorem. Consider a padded language  $B$  defined as  $B = \{\langle x, y \rangle \mid x \in A, |x| = M, |y| = 2^{M^{1/c}} - M\}$ .

We claim that  $B \notin \text{DSPACE}(\log^c N)$ . Assume by contradiction,  $B \in \text{DSPACE}(\log^c N)$ . We show that in that case,  $A \in \text{DSPACE}(N)$ . Let  $x$  be an input instance of  $A$  of length  $M$ ; we shall solve it by reducing it to an instance of  $B$ , by generating  $z = \langle x, y \rangle$ , where  $y = 0^{2^{M^{1/c}} - M}$ , we have  $|z| = N = 2^{M^{1/c}}$ . The instance  $z$  need not be explicitly stored; every bit can be computed on the fly. Thus the space used is  $O(M)$ . Then use algorithm for  $B$  to solve  $z$  using space  $O(\log^c N) = O(\log^c 2^{M^{1/c}}) = O(M)$ . Thus  $A \in \text{DSPACE}(N)$  a contradiction. Therefore,  $B \notin \text{DSPACE}(\log^c N)$ .

Now we show that  $B$  is in AMPC<sup>0</sup>. Let  $z = \langle x, y \rangle$  be a problem instance of  $B$  of length  $N$ . The crucial point to note is that the membership of  $z$  in  $B$  depends only on  $x$ , which has length  $O(\log^c N)$ . If  $x$  does not have this length, we can safely reject it. Fix an arbitrary  $\epsilon \in (0, 1)$ . Consider the AMPC algorithm with just one machine and one round. Let  $z = \langle x, y \rangle$  be an input of length  $N$ . The machine  $\mathcal{M}$  reads  $x$  which has length  $M = O(\log^c N) \subseteq O(N^\epsilon)$ ,

then checks if  $x \in A$  using space  $O(M^2) = O(\log^{2c} N) \subseteq O(N^\epsilon)$ , via our assumption that  $A \in \text{DSPACE}(N^2)$ . If  $x \in A$  then  $\mathcal{M}$  accepts otherwise rejects. Thus we have exhibited a language  $B$  such that  $B \in \text{AMPC}^0$  but  $B \notin \text{DSPACE}(\log^c N)$ . ◀

► **Theorem 9.**  $\text{ReachUL} \subsetneq \text{AMPC}^0$ . That is,  $\text{ReachUL}$  is a proper subset of  $\text{AMPC}^0$ .

**Proof.** The containment follows since  $\text{REACH-UNAMBIGUOUS}$  is complete for  $\text{ReachUL}$  under logspace reductions and by Corollary 4,  $\text{AMPC}^0$  is closed under logspace reductions. The strict containment follows from the fact that  $\text{ReachUL} \subseteq \text{NLOG} \subseteq \text{DSPACE}(\log^2 N)$ . Hence by Theorem 8, there is a language in  $\text{AMPC}^0$  that is not in  $\text{ReachUL}$ . ◀

### 3.2 Limitations

This section discusses the limitations of the AMPC model in relation to well-known complexity classes.

**Uniform Model.** Since each processor in each round has a memory bound of  $O(N^\epsilon)$ , the number of configurations of each processor is  $\text{poly}(2^{N^\epsilon})$  and hence runs in  $O(2^{N^{\epsilon'}})$  for some  $0 < \epsilon' < 1$  (since the processors are halting). Thus it is clear that uniform  $\text{AMPC}^{\text{poly}}$  is in  $\text{SubEXP}$ . Thus by time-hierarchy theorem, there is a language in  $\text{EXP}$  that is not in  $\text{AMPC}^{\text{poly}}$ . This also establishes that it is unlikely that  $\text{PSPACE}$  is in  $\text{AMPC}^{\text{poly}}$  as this will imply  $\text{PSPACE}$  is a subset of  $\text{SubEXP}$ . Moreover, no NP-complete problem (under logspace reduction) is in  $\text{AMPC}^{\text{poly}}$  unless  $\text{NP} \subseteq \text{SubEXP}$ .

**Non-uniform Model.** In the case of non-uniform AMPC computations, we can argue that any language accepted by a polynomial round AMPC algorithm can be simulated by a Boolean circuit of size  $\text{poly}(2^{N^\epsilon})$ . This is because every bit computed by a processor is a decision tree of size  $O(2^{N^\epsilon})$  and hence has a Boolean circuit of size  $O(2^{N^\epsilon})$ . Since the number of bits written by all machines overall (polynomial) rounds is bounded by a polynomial, the size of the Boolean circuit simulating the whole computation is  $\text{poly}(2^{N^\epsilon})$ . It is known that there is a language  $L$  in  $\text{E}^{\Sigma_2^P}$  that has maximum circuit complexity [14], it follows that  $L$  is not in non-uniform  $\text{AMPC}^{\text{poly}}$ . This lower bound establishes that showing  $\text{PSPACE}$  is in non-uniform  $\text{AMPC}^{\text{poly}}$  is difficult as this will imply an unknown complexity theory separation that  $\text{PSPACE}$  is a proper subset of  $\text{EXP}^{\Sigma_2^P}$ .

---

### References

- 1 E. Allender and K.-J. Lange.  $\text{RSPACE}(\log n) \subseteq \text{DSPACE}(\log^2 n / \log \log n)$ . *Theory of Computing Systems*, 31(5):539–550, October 1998. doi:10.1007/s002240000102.
- 2 Alexandr Andoni, Zhao Song, Clifford Stein, Zhengyu Wang, and Peilin Zhong. Parallel graph connectivity in log diameter rounds. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, 2018.
- 3 Alexandr Andoni, Clifford Stein, and Peilin Zhong. Log Diameter Rounds Algorithms for 2-Vertex and 2-Edge Connectivity. In *46th International Colloquium on Automata, Languages, and Programming (ICALP 2019)*, Leibniz International Proceedings in Informatics (LIPIcs), pages 14:1–14:16, 2019.
- 4 Sepehr Assadi, Xiaorui Sun, and Omri Weinstein. Massively parallel algorithms for finding well-connected components in sparse graphs. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*, PODC '19, 2019. doi:10.1145/3293611.3331596.
- 5 Paul Beame, Paraschos Koutris, and Dan Suciu. Communication steps for parallel query processing. *J. ACM*, 64(6), October 2017. doi:10.1145/3125644.

