



**HAL**  
open science

# Local approximation of the Maximum Cut in regular graphs

Étienne Bamas, Louis Esperet

► **To cite this version:**

Étienne Bamas, Louis Esperet. Local approximation of the Maximum Cut in regular graphs. 45th International Workshop on Graph-Theoretic Concepts in Computer Science (WG 2019), Jun 2019, Vall de Núria, Spain. pp.45-59, 10.1016/j.tcs.2020.03.008 . hal-02147410

**HAL Id: hal-02147410**

**<https://hal.science/hal-02147410>**

Submitted on 22 Aug 2022

**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.



Distributed under a Creative Commons Attribution - NonCommercial 4.0 International License

## LOCAL APPROXIMATION OF THE MAXIMUM CUT IN REGULAR GRAPHS

ÉTIENNE BAMAS AND LOUIS ESPERET

**ABSTRACT.** This paper is devoted to the distributed complexity of finding an approximation of the maximum cut (MAXCUT) in graphs. A classical algorithm consists in letting each vertex choose its side of the cut uniformly at random. This does not require any communication and achieves an approximation ratio of at least  $\frac{1}{2}$  in expectation. When the graph is  $d$ -regular and triangle-free, a slightly better approximation ratio can be achieved with a randomized algorithm running in a single round. Here, we investigate the round complexity of *deterministic* distributed algorithms for MAXCUT in regular graphs. We first prove that if  $G$  is  $d$ -regular, with  $d$  even and fixed, no deterministic algorithm running in a constant number of rounds can achieve a constant approximation ratio. We then give a simple one-round deterministic algorithm achieving an approximation ratio of  $\frac{1}{d}$  for  $d$ -regular graphs when  $d$  is odd. We show that this is best possible in several ways, and in particular no deterministic algorithm with approximation ratio  $\frac{1}{d} + \epsilon$  (with  $\epsilon > 0$ ) can run in a constant number of rounds. We also prove results of a similar flavour for the MAXDICCUT problem in regular oriented graphs, where we want to maximize the number of arcs oriented from the left part to the right part of the cut.

**Keywords.** Maximum Cut, Approximation algorithm, Distributed algorithm, Regular graphs.

### 1. INTRODUCTION

Although the maximum cut problem (MAXCUT) is fundamental in combinatorial optimization, it has not been intensively studied from the perspective of distributed algorithms. The folklore algorithm consisting in choosing uniformly at random one side of the cut for each vertex of a graph  $G$  can however be seen as a distributed randomized algorithm with no rounds of communication. By the linearity of expectation, this algorithm gives a cut (a bipartition of the vertex set) of size at least  $m/2$  in expectation, where  $m$  is the number of edges of  $G$ . Here, by the *size* of the cut, we mean the number of edges connecting the two parts of the bipartition. Since every cut in  $G$  contains at most  $m$  edges, this algorithm has *approximation ratio* at least  $\frac{1}{2}$  in expectation, which means that the size of the cut given by the algorithm is at least  $\frac{1}{2}$  of the size of the maximum cut in expectation.

A natural question is whether a better approximation ratio can be obtained if more rounds of communications are allowed. This question was answered positively by Shearer [28] in the case of triangle-free  $d$ -regular graphs. A  *$d$ -regular* graph is a graph in which every vertex has degree  $d$ . In the case of triangle-free  $d$ -regular graphs, Shearer gave a simple randomized algorithm finding a cut of size at least  $m \cdot (\frac{1}{2} + \frac{0.177}{\sqrt{d}})$  in expectation, and thus achieving an approximation ratio of  $(\frac{1}{2} + \frac{0.177}{\sqrt{d}})$  in expectation. Shearer's algorithm uses a single round of communication, messages consisting of a single bit, and at most 3 random bits per vertex. This was recently improved by Hirvonen, Rybicki, Schmid and Suomela [16], who obtained a simpler algorithm finding a cut of size

---

An extended abstract of this work appeared in the proceedings of the 45th International Workshop on Graph-Theoretic Concepts in Computer Science (WG), 2019.

Partially supported by ANR Projects GATO (ANR-16-CE40-0009-01) and GrR (ANR-18-CE40-0032), and LabEx PERSYVAL-Lab (ANR-11-LABX-0025).

at least  $m \cdot \left(\frac{1}{2} + \frac{0.28125}{\sqrt{d}}\right)$  in expectation. Their algorithm uses a single round of communication, messages consisting of a single bit, and a single random bit per vertex.

The case where  $d$  is small and the *girth* (length of a shortest cycle) is large has also been considered: for 3-regular graphs, Kardoš, Král' and Volec [17] showed that when the girth is at least 637789, there exists a randomized distributed algorithm that outputs a cut of average size at least  $0.88672m$  in at most 318894 rounds (the important value here is the size of the cut). This was improved by Lyons [20], who proved a lower bound of  $0.89m$  for cubic graphs of girth at least 655. The best known lower bound for cubic graphs of large girth,  $0.90m$ , was proved by Gamarnik and Li [10], using a result of Csóka, Gerencsér, Harangi, and Virág [5]. The bound of Lyons [20] holds for any  $d$ -regular graphs of large enough (but constant) girth: such graphs have a cut of size at least  $m \cdot \left(\frac{1}{2} + \frac{2}{\pi\sqrt{d}}\right) \approx m \cdot \left(\frac{1}{2} + \frac{0.637}{\sqrt{d}}\right)$ . On the other hand, Dembo, Montanari and Sen [7] showed that in random  $d$ -regular graphs, the maximum cut has size  $m \cdot \left(\frac{1}{2} + \frac{0.763+o(1)}{\sqrt{d}}\right) + o(m)$  with high probability, proving a conjecture of [10]. The existence of this constant  $\approx 0.763$  is also connected to a conjecture of Hatami, Lovász and Szegedy [15] on limits of sparse graphs (see also the conclusion of [26] where the conjecture is strongly disproved for maximum independent sets, improving on an earlier result of [11]).

All the results mentioned above (except the result of Gamarnik and Li [10])<sup>1</sup> can be translated into algorithms working in the CONGEST model in a constant number of rounds. In this model, each node of the graph corresponds to a processor with infinite computational power and has a unique ID (each ID is an integer between 1 and  $\text{poly}(n)$ , where  $n$  denotes the number of vertices in the graph). Nodes can communicate with their neighbors in the graph in synchronous rounds until each node outputs 0 or 1, corresponding to its side in the cut. In the CONGEST model, each message sent by a node to a neighbor has size  $O(\log n)$ , while in some of the algorithms above, the messages have size at most 1. Let us call  $\text{CONGEST}(B)$  the variant of the CONGEST model in which messages are restricted to have size at most  $B$  (instead of  $O(\log n)$ ), and let us say that an algorithm is *local* in a model if it runs in a constant number of rounds in this model. In particular the results of [16, 17, 28] mentioned above can be translated into local algorithms in the  $\text{CONGEST}(O(1))$  model, while the results of [5, 20] can be translated into local algorithms in the CONGEST model.

Note that some of our lower bounds are also valid in the less restricted LOCAL model where the size of each message is not limited. In the following, we will make it clear if this applies. On the other hand, all our algorithms can be implemented in the PO model (anonymous network with port numbering and orientations), whose assumptions are significantly weaker than the CONGEST model (see [14] for some results on local algorithms in PO and CONGEST).

We now review recent results on distributed approximation of MAXCUT. On the deterministic side, Censor-Hillel, Levy, and Shachnai [3] designed a deterministic  $\frac{1}{2}$ -approximation that runs in  $\tilde{O}(\Delta + \log^* n)$  rounds in the CONGEST model on any graph of maximum degree at most  $\Delta$ . More recently, Kawarabayashi and Schwartzman [18] improved the complexity for constant factor approximation by providing a deterministic  $(\frac{1}{2} - \epsilon)$ -approximation that runs in  $O(\log^* n)$  rounds (for any  $\epsilon > 0$ ), in the CONGEST model. However, no deterministic *local* approximation for MAXCUT (i.e. running in a constant number of rounds) in the CONGEST model is known.

---

<sup>1</sup>Their algorithm has two phases: The first consists in finding a large bipartite induced subgraph (this can be done in a constant number of rounds), and the second phase greedily assigns each remaining vertex to the side of the bipartition where it has fewer neighbors (it is unlikely that this greedy procedure can be performed in a constant number of rounds, and at least Theorem 2.2 shows that it cannot be performed in a constant number of rounds when there are no restrictions on the girth).

There is a similar gap between randomized and deterministic approximations for the maximum directed cut problem. Censor-Hillel, Levy, and Shachnai [3] provided a deterministic algorithm running in  $O(\Delta + \log^* n)$  rounds that guarantees a  $\frac{1}{3}$ -approximation as well as a randomized  $\frac{1}{2}$ -approximation with the same round complexity. The round complexities were improved by Kawarabayashi and Schwartzman [18] who provided a deterministic  $(\frac{1}{3} - \epsilon)$ -approximation running in  $O(\log^* n)$  rounds as well as a randomized  $(\frac{1}{2} - \epsilon)$ -approximation in  $O(\epsilon^{-1})$  rounds. All these results are stated in the CONGEST model. Similarly, no deterministic local algorithm is known to achieve a constant factor approximation for this problem.

**1.1. Our results.** Our work focuses on bridging the gap between extremely efficient randomized local algorithms and slower deterministic algorithms for MAXCUT. It should be noted that there are generic tools to derandomize distributed algorithms (see [4, 12] for recent results in this direction) but existing techniques mainly apply to *locally checkable* problems (problem for which a solution can be checked locally), which is not the case for (approximations of) MAXCUT.

In Section 2 we show that any deterministic algorithm that guarantees a constant factor approximation for MAXCUT on the class of bipartite  $d$ -regular graphs when  $d$  is a (constant) even integer requires  $\Omega(\log^* n)$  rounds, which matches the complexity of the algorithm of Kawarabayashi and Schwartzman [18] mentioned above. When  $d$  is odd, we show that one cannot achieve an approximation ratio better than  $\frac{1}{d}$  in a constant number of rounds. Our proofs use an elementary graph construction and then apply Ramsey's theorem [27]. Both of these arguments are not new in distributed algorithms: our construction is inspired by Linial's seminal paper [19] that provides a lower bound on the round complexity of coloring cycles and by a more recent paper by Åstrand, Polishchuk, Rybicki, Suomela, and Uitto [1] which applies Ramsey's theorem in a similar setting to prove that there is no deterministic and local constant factor approximation for the maximum matching problem. Note that similar arguments were also used by Czygrinow, Hanckowiak, and Wawrzyniak [6] to prove lower bounds for the approximation of maximum independent sets in cycles. Our results hold for any  $d$ -regular graph ( $d$  is not necessarily equal to 2), so some additional work needs to be done compared to the simple case of cycles.

In Section 3, we show that this barrier of  $\frac{1}{d}$  when  $d$  is odd is sharp: we first remark that a result of Naor and Stockmeyer [24] on weak 2-coloring of graphs directly gives a deterministic local algorithm that guarantees a  $\frac{1}{d}$ -approximation. We then provide a much simpler and faster deterministic local algorithm achieving the same approximation ratio. It runs in a single round with messages of size  $O(\log n)$  and we show that this cannot be improved.

For the Maximum Directed Cut problem in  $d$ -regular graphs, we prove that a similar situation occurs. If  $d$  is even, a constant factor approximation cannot be achieved in  $o(\log^* n)$  rounds, and if  $d$  is odd, no  $(\frac{2}{d} + \epsilon)$ -approximation can be achieved in  $o(\log^* n)$  rounds (for any  $\epsilon > 0$ ). On the other hand, if  $d$  is odd, a  $(\frac{2}{d+1/d})$ -factor approximation can be achieved in 0 rounds, and a  $(\frac{2}{d+1/d-3/d^2+O(d^{-3})})$ -factor approximation can be achieved in 2 rounds. Note that there is a small gap between the lower bounds and the upper bound of  $\frac{2}{d}$ , and we explain some obstacles towards closing the gap.

Our results imply that while finding a constant factor approximation for the (directed) maximum cut problem in regular graphs of even degree does not require any communication for randomized distributed algorithm (i.e. it can be solved in 0 rounds), for deterministic algorithms an unbounded number of rounds is needed in this case. Note that this separation is not possible for locally checkable problems (see Theorem 3 in [4]). The (perhaps) surprising aspect is that in the case of regular graphs

of *odd* degree, the problem can be solved by a *deterministic* algorithm without communication (if some orientation is given).

Note that another example of a non locally checkable problem with such a separation between the randomized and deterministic complexities was given in [12]. Their problem consists in marking  $(1 + o(1))\sqrt{n}$  vertices of an  $n$ -cycle; the randomized version can also be solved in 0 rounds, while the deterministic version needs  $\Omega(\sqrt{n})$  rounds.

**1.2. Definitions.** A *cut* in a graph  $G$  is a bipartition  $(A, B)$  of its vertex set  $V(G)$ . We usually refer to  $A$  and  $B$  as the left side and the right side of the cut, respectively. The *size* of a cut  $(A, B)$  is the number of edges with one end in  $A$  and the other in  $B$ . The MAXCUT problem in a graph  $G$  consists in finding a cut in  $G$  whose size is maximum.

Given an oriented graph  $G$ , a *directed cut* is again a bipartition  $(A, B)$  of the vertex set of  $G$ , and the size of the directed cut  $(A, B)$  is the number of arcs with their tail in  $A$  and their head in  $B$ . The MAXDICCUT problem in an oriented graph  $G$  consists in finding a directed cut in  $G$  whose size is maximum.

Our results in this paper mainly concern  $d$ -regular graph, i.e. graphs in which each vertex has degree  $d$ . When we refer to an oriented  $d$ -regular graph  $G$ , we mean that the underlying unoriented graph is  $d$ -regular (the out-degrees can be arbitrary).

For an integer  $k \geq 1$ , the *tower function*  $\text{twr}_k$  is the function defined as  $\text{twr}_1(x) = x$  and  $\text{twr}_k(x) = 2^{\text{twr}_{k-1}(x)}$  for  $k \geq 2$ . The *iterated logarithm* of an integer  $n$ , denoted by  $\log^* n$  is defined as 0 if  $n \leq 1$ , and as  $1 + \log^*(\log n)$  otherwise (here and everywhere else in the paper,  $\log$  denotes the base 2 logarithm). The following can be easily derived by induction on  $k$ .

**Claim 1.1.** *For any  $k, n \geq 1$ :*

$$\log^*(\text{twr}_k(n)) = k - 1 + \log^*(n)$$

## 2. DETERMINISTIC CONSTANT FACTOR APPROXIMATION IN REGULAR GRAPHS

As mentioned in the introduction of this paper, Kawarabayashi and Schwartzman [18] provided a deterministic approximation algorithm running in  $O(\log^* n)$  rounds for both problems studied here. In this section, we show using simple arguments based on bounds on Ramsey numbers that their bound is best possible.



Let  $[N] = \{1, \dots, N\}$ . The  $q$ -color Ramsey number  $r_k(n; q)$  is the minimum  $N$  such that in any  $q$ -coloring of the  $k$ -element subsets of  $[N]$ , there is an  $n$ -element subset  $S$  of  $[N]$  such that all  $k$ -element subsets of  $S$  have the same color (see [23] for a recent survey on Ramsey numbers).

**Theorem 2.1** ([9, 8]). *There exists  $c > 0$  such that for any positive integers  $q, k$ , and  $n$ , we have  $r_k(n; q) \leq \text{twr}_k(c \cdot n \cdot q \log q)$ .*

We will also need two simple constructions of  $d$ -regular bipartite graphs.

First we assume that  $d$  is even. We consider a cycle  $C$  of size  $n \geq 2d$ , with  $n$  even, and then add an edge between each pair of vertices that are at distance exactly  $i$  in  $C$  for every  $i \in \{3, 5, 7, \dots, d-1\}$ . This graph, which we denote by  $C_n^d$ , is certainly bipartite (the bipartition corresponds to the vertices at even distance from some arbitrary vertex in  $C$ , and the vertices at odd distance from this vertex). See figure 1 for an example of this graph. By a slight abuse of notation, we say that two (or more) vertices of  $C_n^d$  are *consecutive* if they are consecutive in  $C$ . Similarly, when we refer to the *clockwise order* around  $C_n^d$ , we indeed refer to the clockwise order around  $C$ .

Assume now that  $d$  is odd. We take two disjoint copies of  $C_n^{d-1}$  and assume that the vertices of the cycle  $C$  in the first copy are  $u_1, u_2, \dots, u_n$ , in clockwise order, and the vertices of the cycle  $C$  in the second copy are  $v_1, v_2, \dots, v_n$  in clockwise order. We then connect  $u_i$  and  $v_i$  by an edge, for any  $1 \leq i \leq n$  (and we say it what follows that  $u_i$  and  $v_i$  are *matched*). This graph, which we denote by  $D_{2n}^d$ , is clearly bipartite and  $d$ -regular, see Figure 2 for an example.


 FIGURE 1.  $C_{12}^4$ 

 FIGURE 2.  $D_{24}^5$ 

We are now ready to state the main result of this section.

**Theorem 2.2.** *Let  $d \geq 2$  be a fixed integer.*

- *If  $d$  is even, then any deterministic algorithm in the LOCAL model that guarantees a constant factor approximation for MAXCUT on the class of bipartite  $d$ -regular  $n$ -vertex graphs runs in  $\Omega(\log^* n)$  rounds.*
- *If  $d$  is odd, then for any  $\epsilon > 0$ , any deterministic  $(\frac{1}{d} + \epsilon)$ -approximation algorithm in the LOCAL model for MAXCUT on the class of bipartite  $d$ -regular  $n$ -vertex graphs runs in  $\Omega(\log^* n)$  rounds.*

Note that since the LOCAL model is less restrictive than the CONGEST model, this theorem is also valid in the CONGEST model.

*Proof.* We prove the two cases of the theorem separately starting with  $d$  even:

Let  $d$  be an even integer and assume that there exists a local deterministic algorithm **A** running in  $T$  rounds and guaranteeing a  $\frac{1}{a}$ -approximation for some fixed  $a \geq 1$ , with  $T$  to be defined later. Note that since **A** runs in  $T$  rounds, the output of a vertex  $v$  in  $C_n^d$  only depends on the IDs of the vertices at distance at most  $T$  from  $v$  in  $C_n^d$ , and thus at distance at most  $(d-1)T$  of  $v$  in  $C$  (more precisely, since the subgraph induced by each ball of a given radius is the same, the output of a vertex only depends on the sequence of the IDs of its neighbors at distance  $(d-1)T$  in  $C$ , in clockwise order).

Consider a subset  $U = \{u_1, \dots, u_\ell\}$  of  $[n]$ , with  $u_1 < \dots < u_\ell$ , and assume that  $\ell$  consecutive vertices of  $C_n^d$  (in clockwise order) have IDs  $u_1, u_2, \dots, u_\ell$ , in this order. In what follows, we identify each vertex of  $C_n^d$  with its ID. We now set  $r = 2(d-1)T + 1$ , and start by proving the following claim (recall that  $1/a$  is the approximation ratio of **A**):

**Claim 2.3.** *If  $\ell \geq 4adT$ , at least  $\frac{\ell d}{2} (1 - \frac{1}{2a})$  edges of  $C_n^d$  have both endpoints in  $\tilde{U} = \{u_{(r-1)/2+1}, \dots, u_{\ell-(r-1)/2}\}$ .*



*Proof.* Count the edges by the length of the jump they make around the cycle  $C$ . For a jump of size  $k$ , there are more than  $\ell - (r - 1) - 2k$  such edges. Summing for  $k$  odd from 1 to  $d - 1$  we obtain at least

$$\frac{\ell d}{2} - \frac{d(r-1)}{2} - \frac{d^2}{2}$$

edges that have both endpoints in  $\tilde{U}$ . When  $\ell \geq 4adT \geq 2a(r-1-d)$ , this is at least  $\frac{\ell d}{2} \left(1 - \frac{1}{2a}\right)$ , which proves the claim.  $\square$

Consider some  $r$ -element subset  $S = \{a_1, a_2, \dots, a_r\}$  of  $[n]$ , with  $a_1 < a_2 < \dots < a_r$ . We assign the IDs  $a_1, a_2, \dots, a_r$  (in this order) to  $r$  consecutive vertices of  $C_n^d$ , in clockwise order, and look at the output of the vertex with ID  $a_{(r+1)/2}$  (call it  $v$ ) given by the algorithm **A**. Note that this output only depends on  $(a_1, a_2, \dots, a_r)$ . If  $v$  joins the left side of the cut (according to **A**), we color the set  $S$  with color 0 and otherwise with color 1.

Consider again a subset  $U = \{u_1, \dots, u_\ell\}$  of  $[n]$ , with  $u_1 < \dots < u_\ell$ , and assume that  $\ell$  consecutive vertices of  $C_n^d$  (in clockwise order) have IDs  $u_1, u_2, \dots, u_\ell$ . It follows from the definition of  $\tilde{U}$  that if  $v \in \tilde{U}$ , all the vertices at distance at most  $T$  from  $v$  in  $C_n^d$  are in  $U$ . This implies that if all  $r$ -elements subsets of  $U$  are assigned the same color in the coloring defined above, all the vertices of  $\tilde{U}$  choose the same side of the cut.

We can now apply Theorem 2.1 with  $r$  as defined above and  $\ell = \lceil 4adT \rceil$  satisfying the condition of Claim 2.3. Let  $N = r_r(\ell; 2) \leq \text{twr}_r(c\ell)$  be given by Theorem 2.1. We now let  $n$  be the smallest even integer which is greater than  $2aN$  and consider  $C_n^d$ . Observe that by Claim 1.1,  $\log^* n = O(\log^* N) = O(r + \log^*(c\ell)) = O(dT + \log^*(dT)) = O(dT)$ . Since  $d$  is a fixed constant, it follows that we have  $T = \Omega(\log^* n)$ , as desired.

By Theorem 2.1, there is an  $\ell$ -element subset  $U_1$  of  $[n]$  such that all  $r$ -elements subsets of  $U_1$  have the same color. As long as there are more than  $N$  remaining labels, we repeatedly apply Theorem 2.1 and thus find disjoint  $\ell$ -element subsets  $U_1, U_2, \dots, U_k$  of  $[n]$ , with the same property (for each  $U_i$ , all  $r$ -elements subsets of  $U_i$  have the same color), until  $[n] - \bigcup_{i=1}^k U_i$  contains fewer than  $N$  elements. We write each set  $U_i$  as  $\{u_1^i, \dots, u_\ell^i\}$ , with  $u_1^i < \dots < u_\ell^i$ .

Finally, we assign these IDs to consecutive vertices in clockwise order around  $C_n^d$ :

$$u_1^1, \dots, u_\ell^1, u_1^2, \dots, u_\ell^2, \dots, u_1^k, \dots, u_\ell^k$$

By Claim 2.3, the subgraph induced by each  $\tilde{U}_i$  contains at least  $\frac{nd}{2} \left(1 - \frac{1}{2a}\right)$  edges for all  $1 \leq i \leq k$ , and it follows from Theorem 2.1 and our coloring of the  $r$ -elements sets that for each  $1 \leq i \leq k$ , all vertices in  $\tilde{U}_i$  choose the same side of the cut.

Since  $k \geq \frac{n-N}{\ell}$ , by running algorithm **A** on this particular labelling of  $C_n^d$ , at least

$$\frac{n-N}{\ell} \frac{nd}{2} \left(1 - \frac{1}{2a}\right) = \left(1 - \frac{1}{2a}\right) \frac{d}{2} (n-N) > \frac{nd}{2} \left(1 - \frac{1}{2a}\right)^2 > \frac{nd}{2} \left(1 - \frac{1}{a}\right)$$

edges are not in the cut (the last inequality uses the fact that  $1/a \leq 1$ ). Thus, there are strictly less than  $\frac{nd}{2a}$  edges in the cut. Since  $C_n^d$  is bipartite and  $d$ -regular, the optimal cut contains  $\frac{nd}{2}$  edges (i.e. all the edges are in the cut). This proves that **A** cannot be a  $\frac{1}{a}$ -approximation, yielding a contradiction.

Assume now that  $d$  is odd and that an algorithm **B** achieves a  $(\frac{1}{d} + \epsilon)$ -approximation for some  $\epsilon > 0$  in  $T$  rounds. We proceed as before except that instead of considering vertices one by one, we consider pairs of vertices  $u_i, v_i$  that are matched in the construction of  $D_{2n}^d$ , the graph that will be used here. Similarly as before, a matched pair  $(u_i, v_i)$  cannot see more than  $T(d-1)$  labels away.

Set  $r = 4T(d-1) + 4$ , and consider  $r/2$  consecutive vertices (in clockwise order)  $u_1, \dots, u_{r/2}$  on the outer cycle of  $D_{2n}^d$ . For each  $1 \leq i \leq r/2$ , let  $v_i$  be the neighbor of  $u_i$  on the inner cycle. This implies that  $v_1, \dots, v_{r/2}$  are consecutive (in clockwise order) on the inner cycle.

Fix an arbitrary  $r$ -element subset  $\{a_1, \dots, a_r\}$ , with  $a_1 < a_2 < \dots < a_r$ . For each  $1 \leq i \leq r/2$  assign the ID  $a_i$  to  $u_i$  and the ID  $a_{r/2+i}$  to  $v_i$ . Since the sides of the cut chosen by  $u_{r/4}$  and its neighbor  $v_{r/4}$  are entirely determined by the set  $\{a_1, \dots, a_r\}$ , we can color each set  $\{a_1, \dots, a_r\}$  with the pair  $(x, y) \in \{0, 1\}^2$  such that  $u_{r/4}$  chooses side  $x$  and  $v_{r/4}$  chooses side  $y$  (again we associate the left side of the cut with 0, and the right side of the cut with 1).

By exactly the same argument as in Claim 2.3, we can take  $\ell$  even and large enough so that at least  $\frac{\ell(d-1)}{2} (1 - \frac{\alpha}{2})$  edges are not in the cut in both copies of  $C_n^{d-1}$ , for any fixed  $\alpha$ . By Theorem 2.1 (with  $q = 4$ ), there is an integer  $N = r_r(2\ell; 2)$  such that for all  $n \geq N$ , there is a  $2\ell$ -element subset  $U_1 = \{a_1, \dots, a_{2\ell}\}$  of  $[2n]$  with  $a_1 < \dots < a_{2\ell}$  such that all  $r$ -element subsets  $X$  of  $U_1$  have the same color. We take  $n > \frac{N}{\alpha}$  even and repeatedly apply Theorem 2.1 as before, obtaining  $2\ell$ -element subsets  $U_1, U_2, \dots, U_k$ . We then assign the elements of each  $U_i$  to consecutive vertices in clockwise order in  $D_{2n}^d$  (the  $\ell$  smaller elements of  $U_i$  are assigned to the vertices of the outer cycle, and the  $\ell$  larger elements are assigned to their corresponding neighbors in the inner cycle). As before, the fact that all  $r$ -element subsets of  $U_i$  have the same color implies that on the portion of  $D_{2n}^d$  corresponding to  $U_i$ , all the vertices of the outer cycle (except at the boundary) choose the same side  $x$  of the cut, and all the vertices of the inner cycle (except at the boundary) choose the same side  $y$  of the cut (but  $x$  and  $y$  might be different).

This ensures that on each copy of  $C_n^{d-1}$ , at least

$$\frac{2n - N}{2\ell} \frac{m(d-1)}{2} \left(1 - \frac{\alpha}{2}\right) > \frac{n(d-1)}{2} \left(1 - \frac{\alpha}{2}\right) - \frac{n(d-1)}{2} \left(1 - \frac{\alpha}{2}\right) \frac{\alpha}{2} > \frac{n(d-1)}{2} (1 - \alpha)$$

edges are not in the cut after running algorithm B. It follows that in  $D_{2n}^d$ , at least  $n(d-1)(1-\alpha)$  edges are not in the cut. Since  $D_{2n}^d$  is bipartite and contains exactly  $nd$  edges, this shows there are less than  $n + n(d-1)\alpha$  edges in the cut, which is a  $(\frac{1}{d} + (d-1)\alpha)$ -fraction. Setting  $\alpha = \frac{\epsilon}{d-1}$ , this fraction is less than  $\frac{1}{d} + \epsilon$ , which is a contradiction.  $\square$

A direct consequence of our theorem is the following corollary that matches the round complexity obtained by Kawarabayashi and Schwartzman [18]:

**Corollary 2.4.** *Deterministic constant factor approximation on general graphs for MAXCUT in the LOCAL model requires  $\Omega(\log^* n)$  rounds.*

**2.1. Directed cut.** In this section, we consider the similar problem MAXDICUT where edges are oriented and we only count the edges going from the left side of the cut to the right side. We can prove similar bounds on the quality of the solution one can hope to achieve by simply orienting our lower bound graphs  $C_n^d$  and  $D_{2n}^d$ : we will define  $\overrightarrow{C}_n^d$  as the same graph as  $C_n^d$  where we orient all the edges in clockwise order. Similarly,  $\overrightarrow{D}_{2n}^d$  is obtained from  $D_{2n}^d$  by orienting all the edges in clockwise order on both the inner and outer cycle, and all the edges in the remaining perfect matching from the outer cycle to the inner cycle. We can again apply Ramsey's theorem as in the proof of Theorem 2.2 to obtain the following result :

**Theorem 2.5.** *Let  $d > 0$  be a fixed integer.*

- *If  $d$  is even, any deterministic algorithm that guarantees a constant factor approximation for MAXDICUT on the class of  $d$ -regular bipartite  $n$ -vertex oriented graphs requires  $\Omega(\log^* n)$  rounds in the LOCAL model.*



- If  $d$  is odd, then for any  $\epsilon > 0$ , any deterministic  $(\frac{2}{d} + \epsilon)$ -approximation of MAXDicut on the class of  $d$ -regular bipartite  $n$ -vertex oriented graphs requires  $\Omega(\log^* n)$  rounds in the LOCAL model.

We note a slight difference with Theorem 2.2 in the case where  $d$  is odd. In Theorem 2.5 the approximation ratio is only bounded by  $\frac{2}{d}$ , instead  $\frac{1}{d}$ . This happens because with our definition of  $\overrightarrow{D}_{2n}^d$ , one can check that the optimal directed cut is of size  $\frac{nd}{4} = \frac{m}{2}$  instead of  $m$  in the undirected case.

### 3. MATCHING THE APPROXIMATION RATIO WHEN $d$ IS ODD

**3.1. Weak-coloring.** In a landmark paper, Naor and Stockmeyer [24] addressed the issue of what can or cannot be computed locally. In particular, they proved one result that is relevant in our case.

A *weak coloring* of a graph is a coloring of its vertices such that each vertex has at least one neighbor with a different color. Observe that a weak coloring using only 2 colors is a  $\frac{1}{d}$ -approximation of the MAXCUT problem when the graph is  $d$ -regular. Let  $O_d$  be the class of graphs of maximum degree  $d$  where the degree of every vertex is odd. Naor and Stockmeyer proved the following theorem.

**Theorem 3.1** ([24]). *There is a constant  $b$  such that, for every  $d$ , there is a deterministic algorithm with round complexity  $\log^* d + b$  in the CONGEST model that solves the weak 2-coloring problem in the class  $O_d$ .*

As discussed above, this result directly implies that one can produce a local deterministic  $\frac{1}{d}$ -approximation of the MAXCUT problem on  $d$ -regular graphs. However, the result given here is much stronger than what we are looking for as in this case *every* vertex has at least one incident edge in the cut. A natural question is whether a faster algorithm (of round complexity that does not depend on  $d$ ) exists for the MAXCUT problem on  $d$ -regular graphs with  $d$  odd. In the next section, we prove that such an algorithm exists.

**3.2. A simpler and faster algorithm.** Consider the following algorithm: every vertex  $v$  collects the list of IDs of its neighbors, then  $v$  chooses its side of the cut depending on whether the median value of this list is higher or lower than its own ID. We call this algorithm the *median algorithm*. It runs in a single round and we prove the following theorem:

**Theorem 3.2.** *When the input is a  $d$ -regular graph on  $n$  vertices, with  $d$  odd, the median algorithm finds in 1 round a  $\frac{1}{d}$ -approximation for the MAXCUT problem in the CONGEST model.*

We will actually give two different proofs of this result (i.e. Theorem 3.2 will be a direct consequence of Theorem 3.3, which we proved next, but also of Theorem 3.8, which will be proved in Section 3.3).

**Theorem 3.3.** *When the input is a  $d$ -regular graph on  $n$  vertices, with  $d$  odd, the median algorithm outputs in 1 round (in the CONGEST model) a cut of size at least  $\frac{n}{2} + \frac{(d-1)(d+1)}{4}$ .*

*Proof.* Let  $G = (V, E)$  be a  $d$ -regular graph. In the proof, we say that a vertex is colored 0 or 1 by the median algorithm if it is assigned to the left side or the right side of the cut (respectively). We define the boundary of a subgraph of  $G$  as the set of edges of  $G$  that have exactly one endpoint in this subgraph. For every subset  $A \subset V$  we denote by  $G[A]$  the subgraph induced by  $A$  and  $\partial A$  the boundary of  $G[A]$ .

We now orient each edge of  $G$  from the vertex of lower ID to the vertex of higher ID: it can be observed that the median algorithm assigns color 0 to vertices that have more outgoing edges than

ingoing edges and side 1 when it is the opposite. We also note that this orientation is acyclic, which is the key property that will be used in this proof.

A *monochromatic component* is a connected component of the subgraph of  $G$  induced by one of two sides of the cut. We now prove the following two simple claims.

**Claim 3.4.** *After running the median algorithm, any subset  $A$  of a monochromatic component contains a vertex with at most  $\frac{d-1}{2}$  neighbors in  $A$ .*

Assume for the sake of contradiction that any vertex of  $A$  has at least  $\frac{d+1}{2}$  neighbors in  $A$ . Since the orientation of  $G[A]$  is acyclic, there must be a sink and a source. It follows that one has outdegree at least  $\frac{d+1}{2}$ , and the other has indegree at least  $\frac{d+1}{2}$ . By definition of the median algorithm, the source and the sink must be on different sides of the cut, which contradicts the fact that  $A$  is monochromatic. This concludes the proof of Claim 3.4.

**Claim 3.5.** *After running the median algorithm, for every monochromatic component  $A$  of size  $k$ ,  $\partial A$  contains at least  $k + \frac{(d-1)(d+1)}{4}$  edges if  $k \geq \frac{d+1}{2}$ , and at least  $\frac{d+1}{2} \cdot k$  edges otherwise.*

Let  $A$  be a monochromatic component of  $V$  of size  $k$ . Following Claim 3.4, one can order the vertices  $v_1, \dots, v_k$  of  $A$  such that for all  $1 \leq i \leq k$ ,  $v_i$  has at most  $\frac{d-1}{2}$  neighbors in  $\{v_{i+1}, \dots, v_k\}$ . If  $k \geq \frac{d+1}{2}$ , it follows that there are at most  $\frac{d-1}{2} \cdot k - \frac{(d-1)(d+1)}{8}$  edges in  $G[A]$ . But  $G$  is  $d$ -regular, therefore we have in this case:

$$|\partial A| + 2 \left( \frac{d-1}{2} \cdot k - \frac{(d-1)(d+1)}{8} \right) \geq d \cdot k$$

which implies  $|\partial A| \geq k + \frac{(d-1)(d+1)}{4}$ .

If  $k \leq \frac{d+1}{2}$ ,  $G[A]$  contains at most  $\binom{k}{2} \leq k \cdot \frac{d-1}{4}$  edges, and a similar computation shows that  $|\partial A| \geq k(d - \frac{d-1}{2}) \geq k \cdot \frac{d+1}{2}$ . This concludes the proof of Claim 3.5.

Finally, let  $X$  be the larger side of the cut output by the median algorithm, i.e.  $|X| \geq \frac{n}{2}$ . Observe that the boundary  $\partial X$  is the union of the boundaries of the connected components of  $G[X]$  (since there are no edges between two such components). If at least one of these components has size at least  $\frac{d+1}{2}$ , then it follows from Claim 3.5 that  $|\partial X| \geq \frac{n}{2} + \frac{(d-1)(d+1)}{4}$ , as desired. Otherwise all connected components of  $G[X]$  have size at most  $\frac{d-1}{2}$ , and it follows from Claim 3.5 that  $|\partial X| \geq \frac{n}{2} \cdot \frac{d+1}{2} = \frac{n}{2} + \frac{d-1}{2} \cdot \frac{n}{2} \geq \frac{n}{2} + \frac{(d-1)(d+1)}{4}$ , since  $n \geq d+1$  (recall that  $G$  is  $d$ -regular, and thus contains at least  $d+1$  vertices). This concludes the proof of Theorem 3.3.  $\square$

Since a  $d$ -regular graph has  $\frac{dn}{2}$  edges, we conclude that the median algorithm gives a  $\frac{1}{d}$ -approximation for the maximum cut when  $d$  is odd, which proves Theorem 3.2.

Figure 3 gives an example of labelling of  $D_{2n}^d$  for which the median algorithm gives a cut of size  $\frac{n}{2} + (d-2)^2 + 1$ . This shows that our analysis of the median algorithm in Theorem 3.2 is close to being best possible.

Another interesting aspect of Theorem 3.3 is that it shows that in (the second item of) Theorem 2.2, it is crucial that  $d$  is a fixed constant (independent of  $n$ ). Indeed, if  $d = \Omega(\sqrt{n})$ , then  $\frac{n}{2} + \frac{(d-1)(d+1)}{4} \geq (1 + \Omega(1))\frac{n}{2}$  and thus the median algorithm achieves a  $\frac{1+\epsilon}{d}$ -approximation, for some  $\epsilon > 0$ . This is impossible when  $d$  is a constant, as shown by Theorem 2.2.

The median algorithm is based on finding an (acyclic) orientation of the input graph. Here, we do it by simply orienting the edges from the end with lower ID to the end with higher ID. This costs a single round of communication, with messages of size  $\log n$  (since vertices have to send their




FIGURE 3. Extremal labelling of  $D_{24}^5$  for the *median* algorithm

ID to their neighbors). It follows that in the more restricted  $\text{CONGEST}(b)$  model, where messages have size at most  $b$ , our algorithm takes  $\frac{\log n}{b}$  rounds (here and in the remainder, we omit floors and ceilings whenever they are not necessary in the discussion). In particular, if only messages of size 1 are allowed, our algorithm takes  $\log n$  rounds.

A natural question is whether this can be improved. We now argue that it cannot be improved in general if the algorithm is based on some orientation in the graph. Consider the case where  $G$  contains an isolated edge  $uv$  (two adjacent vertices  $u, v$  of degree 1 in  $G$ ) and we want to construct some orientation of  $G$  (and in particular of the edge  $uv$ ) in the  $\text{CONGEST}(1)$  model (that is with messages of size 1). It seems that the argument below is not original, but we have not been able to find a written source.

Observe that at each round of communication, the message sent by each of  $u, v$  only depends on its ID and the bits received from its neighbor at previous steps. At the first round, at least half of the IDs (call this set  $S_1$ ) would send the same bit, say  $b_1$ , to their neighbor. At the next round, at least half of the IDs of  $S_1$ , upon receiving  $b_1$ , would send the same bit to their neighbor, say  $b_2$ . We continue this process by constructing sets  $S_i$  and bits  $b_i$  for any  $1 \leq i < \log n$  as above (except that for the final round, we define  $b_i$  as the bit output by the vertices, instead of the bit sent to the neighbor). If we use less than  $\log n$  rounds of communication, we can find two distinct IDs in the last set  $S_i$  such that if we assign these IDs to  $u$  and  $v$ , these two vertices will output the same bit  $b_i$ , and therefore they will not be able to deterministically agree on an orientation of the edge  $uv$ . Actually the result holds even if randomization is allowed and we want  $u$  and  $v$  to agree on some orientation of the edge  $uv$  with high probability.

This remark leads us to a similar result for approximating  $\text{MAXCUT}$  in regular graphs. We prove the following:

**Theorem 3.6.** *Let  $D^d = \{D_{2n}^d, n > 0\}$  for  $d$  odd. Any deterministic constant factor approximation of  $\text{MAXCUT}$  on the class  $D^d$  requires at least  $(1 - o(1)) \log n$  rounds in the  $\text{CONGEST}(1)$  model.*

*Proof.* Assume that an algorithm **A** achieves a  $\frac{1}{a}$ -approximation ( $a > 0$ ) on the class of  $d$ -regular graphs in at most  $(1 - \alpha) \log n$  rounds (with  $\alpha > 0$ ).

Assume first that we have a bipartite  $d$ -regular with  $m$  edges and consisting of  $\ell = n^{1-\alpha}$  connected components of size  $k = n^\alpha$ , labelled  $C_1, C_2, \dots, C_\ell$ , and in which each vertex has the same “view” at distance  $\log n$  (i.e. the balls of radius  $\log n$  centered in each of the vertices are isomorphic). We argue, in the same way as above, that we can choose  $\frac{n}{2^{(1-\alpha) \log n}} = n^\alpha = k$  labels such that during  $(1 - \alpha) \log n$  rounds, every vertex in  $C_1$  outputs the same bit. More precisely, for each round  $1 \leq i \leq (1 - \alpha) \log n$ , each vertex of  $C_1$  outputs the same bit  $b_i$ . We proceed similarly for  $C_2, \dots, C_{\ell'}$  with  $\ell' = \ell(1 - \frac{1}{2a})$ : this is possible since before labelling  $C_i$  ( $i \leq \ell'$ ) there are at least  $n - k\ell' = \frac{n}{2a} = \Omega(n)$  available labels.

Hence, we give a labelling of  $C_1, \dots, C_{\ell'}$  such that during  $(1 - \alpha) \log n$  rounds (and in particular, at the end of the algorithm), all the vertices in each given connected component output the same bit. In particular no edge of  $C_1, \dots, C_{\ell'}$  appears in the cut, which implies that at most  $\frac{m}{2a}$  edges appear in the cut. Since every component is bipartite, the maximum cut contains  $m$  edges, which contradicts the hypothesis that **A** achieves a  $\frac{1}{a}$ -approximation.

We then show that this proof can be adapted to the case of  $C_n^d$  or  $D_{2n}^d$  (depending on the parity of  $d$ ). Partition the graph into sets of vertices that appear consecutively in the clockwise order around the cycle(s):  $C_1, \dots, C_\ell$  (each of size  $n^\alpha$ ). As before, we can label  $C_1, \dots, C_{\ell'}$  such that during  $(1 - \alpha) \log n$  rounds, the output bit of every vertex of  $C_i$  is the same (for every  $1 \leq i \leq \ell'$ ). When **A** runs on the graph, a small perturbation of the output bit may appear on the boundary of each  $C_i$ , therefore the output bit of vertices near the boundary is not guaranteed to be the same anymore. However, this does not hurt us since we run the algorithm for  $O(\log n)$  rounds therefore the perturbation can reach at most  $d \log n$  vertices in every  $C_i$ . Therefore, in each  $C_i$  at least  $n^\alpha - O(\log n)$  vertices have the same output during  $(1 - \alpha) \log n$  rounds (and in particular, the same output at the end of the algorithm). This implies there are at least  $\ell'(n^\alpha - O(\log n)) = n(1 - \frac{1}{2a}) - o(n)$  vertices that have the same output as their neighbors at the end of the algorithm. It follows that the cut output by the algorithm has size at most  $\frac{m}{2a} + o(m)$ , while the maximum cut has size  $m$ , which contradicts the definition our initial assumption that that **A** achieves a  $\frac{1}{a}$ -approximation.  $\square$

Note that the randomized, high probability version of Theorem 3.6 does not hold. More precisely, we can show that the trivial randomized  $\frac{1}{2}$ -approximation indeed produces a constant factor approximation with high probability on the class of regular graphs in the CONGEST(0) model. This is certainly a classic result but we have not been able to find it in the literature.

**Theorem 3.7** (folklore). *For any  $\epsilon > 0$ , the folklore algorithm produces a  $\frac{1-\epsilon}{2}$ -approximation with high probability on the class of  $n$ -vertex graphs with degrees bounded by a constant, and  $m = \Omega(n) = \Omega(1/\epsilon^2)$  edges (and in particular in the class of  $d$ -regular graphs of sufficiently large size).*

*Proof.* Let  $G = (V, E)$  be a graph of maximum degree  $d$  containing  $m$  edges. By Vizing’s Theorem,  $G$  has a  $(d + 1)$ -edge-coloring, i.e. a partition of its edge-set into  $d + 1$  matchings  $M_1, \dots, M_{d+1}$ . Assume that  $|M_1| \geq \dots \geq |M_{d+1}|$ , and discard all the matchings  $M_i$  such that  $|M_i| \leq \epsilon' m/d$ , for some  $\epsilon' > 0$  whose value will be fixed later in the proof. Note that the remaining matchings  $M_1, \dots, M_k$  satisfy  $|\bigcup_{i=1}^k M_i| \geq (1 - \epsilon')m$ . Recall that the folklore algorithm assigns each vertex to one of the two sides of the cut, uniformly at random. Note that for each matching  $M_i$ , and for any two edges  $e, f \in M_i$ , the events that  $e$  and  $f$  are in the cut are independent.

We now recall the following Chernoff bound (see e.g. Chapter 5 in [21]): For any  $0 \leq t \leq np$ ,  $\mathbb{P}(|\text{BIN}(n, p) - np| > t) \leq 2 \exp(-t^2/3np)$ , where  $\text{BIN}(n, p)$  denotes the binomial distribution with parameters  $n$  and  $p$ . Thus, for each  $1 \leq i \leq k$ , it follows that with probability at least  $1 - n^{-2}$ ,

at least  $\frac{1}{2}|M_i| - c\sqrt{|M_i|}$  edges of  $M_i$  are in the cut output by the algorithm, for some absolute constant  $c > 0$ . Using the Union bound, with probability at least  $1 - 1/n$  the cut output by the algorithm contains at least

$$\sum_{i=1}^k (\frac{1}{2}|M_i| - c\sqrt{|M_i|}) \geq \frac{m}{2}(1 - \epsilon') - c \cdot \sqrt{k} \cdot \sqrt{\sum_{i=1}^k |M_i|} \geq \frac{m}{2}(1 - \epsilon') - c \cdot d \cdot \sqrt{m} = \frac{m}{2}(1 - \epsilon' - \frac{2cd}{\sqrt{m}})$$

edges. Setting  $\epsilon' = \epsilon - \frac{2cd}{\sqrt{m}}$  yields the desired result (recall that  $m = \Omega(1/\epsilon^2)$  and thus such an  $\epsilon' > 0$  exists).  $\square$

**3.3. Directed cuts.** Given a bipartition  $(V_1, V_2)$  of an oriented graph  $G$ , the set of arcs oriented from  $V_1$  to  $V_2$  (the *directed cut* from  $V_1$  to  $V_2$ ) is denoted by  $\vec{E}(V_1, V_2)$ . The maximum cardinality of a directed cut in  $G$  is denoted by  $\text{maxdicut}(G)$ .

Let  $G$  be an oriented graph. For each vertex  $v$ , we define the *deficit* of  $v$  as  $\delta(v) = d^+(v) - d^-(v)$ , where  $d^+(v)$  and  $d^-(v)$  denote the out-degree and in-degree of  $v$ , respectively. We define the *sign* of a vertex  $v$  as the sign of  $\delta(v)$ , and we say that a vertex is *positive* or *negative* accordingly. The set of positive vertices is denoted by  $V^+$  and the set of negative vertices is denoted by  $V^-$ . Note that if all the vertices of  $G$  have odd degree (in particular if  $G$  is  $d$ -regular with  $d$  odd), then every vertex is positive or negative and this case  $V^+, V^-$  form a bipartition of the vertex set  $V$  of  $G$ .

Note that the median algorithm described in the previous subsection can be rephrased as: find an acyclic orientation of  $G$  and then choose the cut  $(V^+, V^-)$  with respect to this orientation. Our second proof of Theorem 3.2 will be a direct consequence of the following general result (which proves that not only the cut, but also the *directed* cut between  $V^+$  and  $V^-$  has size at least  $n/2$ , and that the original orientation does not need to be acyclic).

**Theorem 3.8.** *Let  $G$  be an  $n$ -vertex oriented  $d$ -regular graph with  $d$  odd, and let  $V^+$  and  $V^-$  be defined as above. Then the directed cut  $\vec{E}(V^+, V^-)$  contains at least  $\max\{\frac{n}{2}, \frac{2}{d+1/d} \cdot \text{maxdicut}(G)\}$  arcs.*

*Proof.* We write  $\text{CUT} = |\vec{E}(V^+, V^-)|$  and  $\text{OPT} = \text{maxdicut}(G)$ , and set  $D = \sum_{v \in V^+} d^+(v) + \sum_{v \in V^-} d^-(v) \geq n \cdot \frac{d+1}{2}$ . Observe that  $D$  counts the number of arcs of  $G[V^+]$  and  $G[V^-]$  once, while the arcs of  $\vec{E}(V^+, V^-)$  are counted twice. Since  $G$  contains  $dn/2$  arcs, it follows that


$$(1) \quad \text{CUT} \geq D - \frac{dn}{2} \geq \frac{n}{2}.$$

This proves that the directed cut output by the algorithm has size at least  $\frac{n}{2}$ , which readily implies Theorem 3.2.

We now consider an optimal directed cut  $\vec{E}(V_1, V_2)$  of  $G$  (i.e. a directed cut of cardinality  $\text{OPT} = \text{maxdicut}(G)$ ), and define  $M$  as the set of vertices  $(V_1 \cap V^-) \cup (V_2 \cap V^+)$  (see Figure 4 for an illustration). Note that each arc of  $\vec{E}(V_1, V_2)$  which is not incident to a vertex of  $M$  is also an arc of  $\vec{E}(V^+, V^-)$ . Since the vertices of  $V_1 \cap V^-$  have out-degree at most  $\frac{d-1}{2}$  and the vertices of  $V_2 \cap V^+$  have in-degree at most  $\frac{d-1}{2}$ , we have

$$(2) \quad \text{CUT} \geq \text{OPT} - \frac{d-1}{2} \cdot |M|.$$

Now observe that

FIGURE 4. Sets  $V^+$ ,  $V^-$ ,  $V_1$ ,  $V_2$ , and  $M$ .

$$\begin{aligned}
2|\vec{E}(V_1, V_2)| &\leq \sum_{v \in V^+ \setminus M} d^+(v) + \sum_{v \in V^- \setminus M} d^-(v) + \sum_{v \in V^- \cap M} d^+(v) + \sum_{v \in V^+ \cap M} d^-(v) \\
&= D + \sum_{v \in V^- \cap M} (d^+(v) - d^-(v)) + \sum_{v \in V^+ \cap M} (d^-(v) - d^+(v)) \leq D - |M|.
\end{aligned}$$

This implies

$$(3) \quad 2 \cdot \text{OPT} \leq D - |M|.$$

It follows from (2) that  $|M| \geq \frac{2}{d-1}(\text{OPT} - \text{CUT})$ , which we can plug into (3) to obtain:

$$\frac{2d}{d-1} \text{OPT} \leq D + \frac{2}{d-1} \text{CUT}.$$

It directly follows from (1) that  $D \leq \text{CUT} + \frac{dn}{2} \leq (d+1) \text{CUT}$  and plugging it into the previous inequality, we obtain:

$$\frac{2d}{d-1} \text{OPT} \leq (d+1 + \frac{2}{d-1}) \text{CUT} = \frac{d^2+1}{d-1} \text{CUT},$$

and finally:

$$\frac{\text{CUT}}{\text{OPT}} \geq \frac{2d}{d^2+1} = \frac{2}{d+1/d},$$

as desired.  $\square$

From now on, we call the 0-round algorithm resulting from Theorem 3.8 the *oriented median algorithm*. The factor  $\frac{2}{d+1/d}$  might seem a little surprising, but it turns out to be sharp, in the following sense: there are  $d$ -regular oriented graphs for which the oriented median algorithm outputs a cut of size precisely  $\frac{2}{d+1/d} \text{OPT}$ . To see this, take  $n$  to be a multiple of  $4d$ , and take 4 sets of vertices  $A, B, C, D$  as in Figure 5. Each set is an independent set, and its size is  $n$  times the fraction indicated in the figure (for instance  $A$  and  $B$  both contain  $\frac{d+1}{4d} \cdot n$  vertices). The arc labelled  $\frac{1}{2}$  between  $A$  and  $B$  indicates that we add  $\frac{1}{2} \cdot n$  arcs joining  $A$  to  $B$ , and similarly for the arcs joining  $A$  and  $D$ , and the arcs joining  $B$  and  $C$ ). It can be checked that the number of arcs incident to each set is precisely  $d$  times the size of the set, so the graph can be made  $d$ -regular (and we can make sure that the out-degree of each vertex is equal to the average out-degree of its part, for instance the vertices of  $A$  have out-degree  $\frac{4d}{d+1} \cdot (\frac{(d-1)^2}{8d} + \frac{1}{2}) = \frac{d+1}{2}$ , so they lie in  $V^+$ ). It can also be checked that the directed cut output by the algorithm,  $\vec{E}(A \cup D, B \cup C)$ , has cardinality  $n/2$  (the arcs joining  $A$  to  $B$ ), while the optimal directed cut  $\vec{E}(A \cup C, B \cup D)$  contains  $\frac{d^2+1}{2d} \cdot \frac{n}{2} = \frac{d+1/d}{2} \cdot \frac{n}{2}$  arcs.






FIGURE 5. An example showing the sharpness of the analysis of Theorem 3.8

So the problem does not come from the analysis of the algorithm, but rather from the algorithm itself, which is unable to detect the kind of structure depicted in Figure 5.

To overcome this issue and close the gap with the  $\frac{2}{d}$  bound, one might be tempted to consider local improvements. In the following, a vertex will be *stable* if it has at least one neighbor on the other side of the cut. Otherwise it will be *unstable*. We now consider the following simple algorithm: at every round, every unstable vertex changes side. The algorithm stops when all vertices are stable. As we can see, the running time (and even the termination) of this algorithm is highly dependent on the starting point: for instance if all vertices start on the same side of the cut, then the algorithm never ends. When we perform one round of this algorithm, we say that we perform a *flip* (as this algorithm can be seen as a variant of the well known FLIP algorithm that is further discussed in the conclusion).

Even though this algorithm may never end, we will prove shortly that if we take as starting point the cut given by the oriented median algorithm (that gives a  $\frac{2}{d+1/d}$ -approximation in 0 rounds) and perform 2 flips, we then improve slightly on the approximation ratio of  $\frac{2}{d+1/d}$ .

We first explain some useful properties of stability, as defined above.

**Claim 3.9.** *Once a vertex is stable, it remains stable after any number of flips.*

*Proof.* Simply notice that if a vertex  $u$  is stable, then it has a neighbor  $v$  on the other side of the cut and  $v$  must be stable too. If  $u$  or  $v$  become unstable then one of them becomes unstable while the second one is on the other side of the cut, which is impossible.  $\square$

**Claim 3.10.** *Once an edge is in the directed cut, it remains in the directed cut after any number of flips.*

*Proof.* Any edge in the directed cut joins two stable vertices, and thus remains in the cut after any number of flips.  $\square$

We define  $\text{CUT}_i$  to be the number of edges in the directed cut after  $i$  flips. We take the notation defined in the proof of Theorem 3.8:  $V^+$  and  $V^-$  are the sets of vertices of positive and negative deficit, respectively,  $\text{CUT}_0 = \text{CUT} = |\vec{E}(V^+, V^-)|$  is the size of the dicut given by the oriented median algorithm (running in 0 rounds),  $\text{OPT} = \text{maxdicut}(G)$  is the size of the maximum dicut,

$D = \sum_{v \in V^+} d^+(v) + \sum_{v \in V^-} d^-(v)$ , and  $M$  is the set of vertices whose side differ in  $\vec{E}(V^+, V^-)$  and in some fixed maximum dicut  $(V_1, V_2)$ .

By Claim 3.10, we have that  $\text{CUT}_j \geq \text{CUT}_i$  for any  $j \geq i$ . Using this, a simple modification of the proof of Theorem 3.8 shows the following:

**Claim 3.11.** *Assume we have the following inequalities for some  $\alpha, \beta \in [0, 1]$ :*

$$\begin{aligned} \text{CUT}_j &\geq \text{OPT} - \frac{d-1}{2} \cdot |M| + \alpha \cdot |M| \\ 2 \cdot \text{OPT} &\leq D - |M| - \beta \cdot |M| \end{aligned}$$

then

$$\frac{\text{CUT}_j}{\text{OPT}} \geq f_d(\alpha, \beta) := \frac{d - 2\alpha + \beta}{d^2/2 - \alpha \cdot (d+1) + \beta + 1/2}$$

The following claim immediately holds as well:

**Claim 3.12.** *For any,  $d \geq 3$ ,  $\alpha, \beta \in [0, 1]$  such that  $\alpha + \beta > 0$ ,  $f_d(\alpha, \beta) > f_d(0, 0) = \frac{2}{d+1/d}$ .*

*More precisely, for any  $y \in (0, 1)$ ,*

$$\inf_{\alpha, \beta \in [0, 1], \alpha + \beta \geq y} f_d(\alpha, \beta) = f_d(0, y) = \frac{3 + y}{5 + y}$$

and for any  $d \geq 5$ ,

$$\inf_{\alpha, \beta \in [0, 1], \alpha + \beta \geq y} f_d(\alpha, \beta) = f_d(y, 0) = \frac{d - 2y}{d^2/2 - y(d+1) + 1/2}$$

*Proof.* Notice that, for any  $d \geq 3$ ,  $\alpha, \beta \in (0, 1)$ ,  $\frac{\partial}{\partial \alpha} f_d(\alpha, \beta) = \frac{4(d-1)(\beta+1)}{(d^2 - 2\alpha \cdot (d+1) + 2\beta + 1)^2} > 0$  and  $\frac{\partial}{\partial \beta} f_d(\alpha, \beta) = \frac{2(d-1)(d-2\alpha+1)}{(d^2 - 2\alpha \cdot (d+1) + 2\beta + 1)^2} > 0$ . This immediately proves the first claim.

For the second claim, by previous calculations we can set  $\beta = y - \alpha$  and compute:

$$\frac{d}{d\alpha} f_d(\alpha, y - \alpha) = -\frac{2(d-1)(d-2y-3)}{(d^2 - 2\alpha(d+2) + 2y + 1)^2}$$

Clearly, if  $d > 3$ , then the minimum is reached for  $\alpha = y$  and  $\beta = 0$ . And if  $d = 3$  then the minimum is reached for  $\alpha = 0$  and  $\beta = y$ .  $\square$

Knowing these claims, we now prove that  $\text{CUT}_2/\text{OPT}$  is greater than some  $f_d(\alpha, \beta)$  with  $\alpha + \beta > 0$ . To show this, we need to prove refined versions of inequalities (2) and (3).

Recall the proof of inequality (2) ( $\text{CUT}_0 \geq \text{OPT} - \frac{d-1}{2} \cdot |M|$ ): Start with some optimum cut and remove the edges of the cut leaving  $M \cap V^-$  and the edges entering  $M \cap V^+$ . By definition we remove at most  $\frac{d-1}{2} \cdot |M|$  edges, which implies (2).

Let  $E_0$  be the set of edges with one end in  $V^-$  and the other in  $M \cap V^+$ , or with one end in  $M \cap V^-$  and the other in  $V^+$ , or between two vertices of  $M \cap V^+$ , or between two vertices of  $M \cap V^-$ . We claim that

$$(4) \quad \text{CUT}_0 \geq \text{OPT} - \frac{d-1}{2} \cdot |M| + |E_0|.$$

To see this, observe that the only edges of  $E_0$  that are in the optimum cut are the edges going from  $M \cap V^-$  to  $M \cap V^+$ , and these are counted twice on our computation. The remaining edges of  $E_0$  are counted once in  $\frac{d-1}{2} \cdot |M|$  or  $\text{CUT}_0$  but not in  $\text{OPT}$ .

We denote the stable and unstable vertices (with respect to  $\text{CUT}_0$ ) by  $S_0$  and  $U_0$ , respectively.

Let  $E_1$  be the set of edges directed from  $V^+ \cap S_0$  to  $V^+ \cap U_0$ , or from  $V^- \cap U_0$  to  $V^- \cap S_0$ . Observe that each such edge is added to the cut after one flip and thus

$$(5) \quad \text{CUT}_1 \geq \text{OPT} - \frac{d-1}{2} \cdot |M| + |E_0| + |E_1|.$$

Let  $U_1$  be the set of vertices that are unstable after 1 flip. Consider by symmetry the subset  $U_1^r$  of  $U_1$  that lie on the right side of the cut. These vertices were in  $V^+$  (i.e. on the left side of the cut) before the first flip. Since each vertex of  $V^+$  has deficit at least 1, the sum of the deficits of the vertices of  $U_1^r$  is at least  $|U_1^r|$  and thus the number of edges leaving  $U_1^r$  is at least  $|U_1^r|$ . After one flip, by definition of  $U_1$ , all these edges are directed toward stable vertices on the right side of the cut. It follows that after a second flip, all these edges join the cut, and thus

$$(6) \quad \text{CUT}_2 \geq \text{OPT} - \frac{d-1}{2} \cdot |M| + |E_0| + |E_1| + |U_1|.$$

We now focus on finding a refined version of inequality (3) (which states that  $2 \cdot \text{OPT} \leq D - |M|$ ). Let  $F_0$  be the set of edges between two vertices of  $V^+ \setminus M$ , or between two vertices of  $V^- \setminus M$ . Note that each edge of  $F_0$  is counted in  $D$  but does not appear in  $\text{OPT}$ , thus we obtain

$$(7) \quad 2 \cdot \text{OPT} \leq D - |M| - |F_0|.$$

If we denote by  $M^*$  the set of vertices of  $M$  with deficit larger than 1 in absolute value (in other words, with deficit at least 3 or at most  $-3$ ), the two main inequalities can be slightly refined as

$$(8) \quad 2 \cdot \text{OPT} \leq D - |M| - |F_0| - |M^*|.$$

$$(9) \quad \text{CUT}_2 \geq \text{OPT} - \frac{d-1}{2} \cdot |M| + |E_0| + |E_1| + |U_1| + |M^*|.$$

We are now ready to prove the following.

**Theorem 3.13.** *Assume that  $d \geq 3$  is odd. Then the 2-round algorithm consisting of the oriented median algorithm followed by two flips provides a  $\frac{2}{\frac{d+1}{d-3} + \frac{2}{d^2 + O(d-3)}}$ -approximation for the MAXDICUT problem in  $d$ -regular graphs.*

*Proof.* We use the notation defined in this section. Let  $M_1$  be the set of vertices of  $M \setminus M^*$  (i.e. the subset of vertices of  $M$  of deficit 1 in absolute value) that are not incident to any edge of  $E_0$  or  $E_1$ . In particular all the vertices of  $M_1$  are unstable, and their in-degrees and out-degrees are  $\frac{d-1}{2}$  or  $\frac{d+1}{2}$ .

Assume first that  $|M_1| \geq x \cdot |M|$ , with  $x = \frac{d^2+d}{d^2+4d+1}$ . Note that for every  $v \in M_1$ , all the neighbors of  $v$  are in  $V \setminus M$ , on the same side of the cut as  $v$  (with respect to  $\text{CUT}_0$ ). Since  $v \in M_1$  is unstable and not incident to any edge of  $E_1$ , all the in-neighbors of  $v$  are unstable as well. It follows that  $|U_0 \setminus M| \geq \frac{d-1}{2d} |M_1|$ . Let  $W$  be the subset of vertices  $w \in U_0 \setminus M$  that have a stable out-neighbor in  $M$  (if  $w \in V^+$ ) or a stable in-neighbor in  $M$  (if  $w \in V^-$ ). Since stable vertices of  $M$  are in  $M \setminus M_1$ , we have  $|W| \leq \frac{d-1}{2} (|M| - |M_1|)$ , and thus

$$\begin{aligned} |U_0 \setminus M| - |W| &\geq \frac{d-1}{2d} |M_1| - \frac{d-1}{2} |M| + \frac{d-1}{2} |M_1| \\ &\geq \frac{d^2-1}{2d} |M_1| - \frac{d-1}{2} |M| \\ &\geq \frac{1}{2d} (x(d^2-1) - d(d-1)) |M| \\ &= (1-x) |M| \end{aligned}$$

by definition of  $x$ . Note that each vertex  $u \in (U_0 \setminus M) \setminus W$  is unstable, so all its neighbors are on the same side as  $u$ . Consider by symmetry the case  $u \in V^+$ . By definition,  $u$  has no stable out-neighbor in  $M$ . If  $u$  has a stable in-neighbor  $v$ , then the edge  $uv$  is in  $E_1$ . If  $u$  has a neighbor

$v \notin M$ , then the edge  $uv$  is in  $F_0$ . If none of these cases occur, then all neighbors of  $u$  are in  $M$ , and are unstable. It follows that  $u$  is in  $U_1$ , the set of vertices that are still unstable after 1 flip. It follows that

$$(10) \quad |E_1| + |F_0| + |U_1| \geq \frac{1}{2}(|U_0 \setminus M| - |W|) \geq \frac{1}{2}(1-x)|M|.$$

Using inequalities (6) and (8) together with Claims 3.11 and 3.12, we obtain that  $\text{CUT}_2/\text{OPT} \geq f_d(\alpha, \beta)$  with  $\alpha, \beta \geq 0$  and  $\alpha + \beta \geq \frac{1}{2}(1-x)$ .

Assume now that  $|M_1| < x \cdot |M|$ . In this case, it means that at least  $(1-x) \cdot |M|$  vertices  $v$  are in  $M^*$  (i.e. have deficit at least 3 in absolute value) or are incident to an edge of  $E_0$  or  $E_1$ . It follows that

$$(11) \quad |M^*| + |E_0| + |E_1| \geq \frac{1}{2}(1-x)|M|,$$

and thus  $\text{CUT}_2/\text{OPT} \geq f_d(\alpha, \beta)$  with  $\alpha, \beta \geq 0$  and  $\alpha + \beta \geq \frac{1}{2}(1-x)$ .

In both cases we obtain  $\text{CUT}_2/\text{OPT} \geq f_d(\alpha, \beta)$  with  $\alpha, \beta \geq 0$  and  $\alpha + \beta \geq \frac{1}{2}(1-x) = \frac{3d+1}{2d^2+8d+2}$ . It follows from Claim 3.12, that  $\text{CUT}_2/\text{OPT} \geq \frac{2}{d+1/d-3/d^2+O(d^{-3})}$ .  $\square$

Theorem 3.13 proves that after 2 flips, we can slightly improve on the approximation ratio of Theorem 3.8. A natural question is whether the same can be achieved after a single flip. The construction of Figure 6, which is a refinement of the construction of Figure 5, shows that it is not the case: if we apply the oriented median algorithm and then perform a single flip, the size of the cut does not change (it remains  $\frac{2}{d+1/d}$  OPT).




FIGURE 6. A refinement of the construction of Figure 5. The subscripts  $u$  and  $s$  stand for *unstable* and *stable*, respectively.

#### 4. CONCLUSION

**4.1. FLIP.** In Section 3, we have designed a very simple one-round algorithm approximating MAX-CUT in regular graphs (with odd degrees). Once a solution has been obtained, it might be tempting to run a few more rounds of computation to see if the solution can be improved locally.

We have already seen a simple way to improve the quality of a solution (by moving the so-called unstable vertices to the other side of the cut), but the notion of stability we used was specifically designed to improve the approximation ratio in a small number of rounds. Another simple way to locally improve a cut (in the sequential setting this time) is to take a vertex with more neighbors

in its own part than in the other part, and change its side. If this is done until no such vertex exists, the resulting cut is *maximal*, and in this case is a  $\frac{1}{2}$ -approximation of the maximum cut. This operation, called *FLIP*, has been studied for a long time. When the edges are weighted, it was proved by Poljak [25] that any sequence of FLIPs takes only polynomially many steps before reaching a maximal cut in cubic graphs, while Monien and Tscheuschner [22] proved that there are graphs of maximum degree 4 for which a sequence of FLIPs can take exponentially many steps to reach a maximal cut. In the unweighted case however, since each flip improves the cut by at least one, the maximum number of flips before reaching a maximal cut is bounded by the number edges (which is linear in  $n$  in bounded degree graphs). In the distributed framework, it might be tempting to consider running some rounds of the *distributed FLIP* dynamics: at each round, each vertex with more neighbors in its own part than in the other part changes side. The graph  $D_{2n}^d$  constructed in the previous sections shows that it might not be helpful at all: if all the vertices of the outer cycle are in one side of the cut, and all the vertices of the inner cycle are on the other side of the cut, then at each round, all the vertices of the graph would change side, not improving the solution.

It might be worth noting that in our application of the median algorithm, not all vertices of the outer cycle of  $D_{2n}^d$  are on the same side of the cut (given the bad labelling of Figure 3): due to some side-effects, roughly  $d$  vertices in the outer cycle are not on the same side of the cut as the others, and similarly for the inner cycle. It can then be checked that if we run the distributed FLIP dynamics in this instance, the solution does improve over time, but improving the approximation ratio from  $\frac{1}{d}$  to  $\frac{1}{d} + \epsilon$  requires  $\Omega(\epsilon n)$  rounds, which is extremely unpractical. This has to be compared with the lower bound of Theorem 2.2, which says that in order to achieve an approximation ratio of  $\frac{1}{d} + \epsilon$  in general, one needs a number of rounds of the order of  $\Omega(\log^* n)$ .

**4.2. Maximal cut.** An interesting aspect of the Maximal cut problem defined in the previous subsection is that it is an LCL problem (of locality 1): each vertex only needs to check that at least half of its neighbors lie on the other side of the cut. This is in stark contrast with MAXCUT, as we have seen already. It was recently proved by Balliu, Hirvonen, Lenzen, Olivetti, and Suomela [2] that any deterministic algorithm finding a maximal cut in  $d$ -regular graphs ( $d \geq 3$ ) takes  $\Omega(\log n)$  rounds and any randomized algorithm takes  $\Omega(\log \log n)$  rounds in the LOCAL model. It would be interesting to find algorithms matching these round complexities. Note that if we merely require that each vertex has at least  $d/2 - O(\sqrt{d})$  neighbors on the other side of the cut, the problem can be easily reduced to the distributed Lovász Local Lemma, and therefore solved efficiently.

*Acknowledgments.* We would like to thank Jérémie Chalopin and Keren Censor-Hillel for their remarks on the complexity of finding an orientation using very small messages in the CONGEST model. We also thank Michal Dory for calling reference [6] to our attention, and David Gamarnik for pointing out references [5, 7, 10, 20] to us.

## REFERENCES

- [1] M. Åstrand, V. Polishchuk, J. Rybicki, J. Suomela and J. Uitto. *Local algorithms in (weakly) coloured graphs*, CoRR abs/1002.0125, 2010.
- [2] A. Balliu, J. Hirvonen, C. Lenzen, D. Olivetti, and J. Suomela, *Locality of not-so-weak coloring*, In Proc. of the 26th International Colloquium on Structural Information and Communication Complexity (SIROCCO), 2019.
- [3] K. Censor-Hillel, R. Levy and H. Shachnai, *Fast Distributed Approximation for Max-Cut*. In Proc. of the 13th International Symp. on Algorithms and Experiments for Wireless Sensor Networks (ALGOSENSORS), 2017.
- [4] Y.-J. Chang, T. Kopelowitz, and S. Pettie, *An Exponential Separation between Randomized and Deterministic Complexity in the LOCAL Model*, In Proc. of the IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS), 2016.
- [5] E. Csóka, B. Gerencsér, V. Harangi, and B. Virág, *Invariant gaussian processes and independent sets on regular graphs of large girth*, Random Structures Algorithms **47** (2015), 284–303.

- [6] A. Czygrinow, M. Hanckowiak, and W. Wawrzyniak, *Fast Distributed Approximations in Planar Graphs*, In Proc. of the 22nd International Symposium on Distributed Computing (DISC), 2008.
- [7] A. Dembo, A. Montanari, and S. Sen, *Extremal cuts of sparse random graphs*, Ann. Probab. **45(2)** (2017), 1190–1217.
- [8] P. Erdős and R. Rado, *Combinatorial theorems on classifications of subsets of a given set*, Proc. London Math. Soc. **3** (1952), 417–439.
- [9] P. Erdős and G. Szekeres, A combinatorial problem in geometry, Compos. Math. **2** (1935), 463–470.
- [10] D. Gamarnik and Q. Li, *On the max-cut of sparse random graphs*, Random Structure Algorithms **52(2)** (2018), 219–262.
- [11] D. Gamarnik and M. Sudan, *Limits of local algorithms over sparse random graphs*, In Proc. of Innovations in Theoretical Computer Science (ITCS), 369–376, 2014.
- [12] M. Ghaffari, D.G. Harris, and F. Kuhn, *On Derandomizing Local Distributed Algorithms*, In Proc. of the IEEE Symposium on Foundations of Computer Science (FOCS), 2018.
- [13] M. Ghaffari, F. Kuhn and Y. Maus, *On the complexity of local distributed graph problems*, In Proc. of the 49th Annual ACM Symp. on Theory of Computing (STOC), 2017, 784–797.
- [14] M. Göös, J. Hirvonen, and J. Suomela, *Lower bounds for local approximation*, J. ACM **60** (2013), #39.
- [15] H. Hatami, L. Lovász, and B. Szegedy, *Limits of local-global convergent graph sequences*, Geom. Funct. Anal. **24(1)** (2014), 269–296.
- [16] J. Hirvonen, J. Rybicki, S. Schmid and J. Suomela, *Large Cuts with Local Algorithms on Triangle-Free Graphs*, Electron. J. Combin. **24(4)** (2017), #P4.21.
- [17] F. Kardoš, D. Král’ and J. Volec. *Maximum edge-cuts in Cubic Graphs With Large Girth and in Random Cubic Graphs*, Random Structures Algorithms **41(4)** (2012), 506–520.
- [18] K.-i. Kawarabayashi and G. Schwartzman, *Adapting Local Sequential Algorithms to the Distributed Setting*, In Proc. 32nd International Symposium on Distributed Computing (DISC), 2018, 35:1—35:17.
- [19] N. Linial, *Locality in distributed graph algorithms*, SIAM J. Comput. **21(1)** (1992), 193–201.
- [20] R. Lyons, *Factors of IID on trees*, Combin. Probab. Comput. **26(2)** (2017), 285–300.
- [21] M. Molloy and B. Reed, *Graph Colouring and the Probabilistic Method*, Springer, 2002.
- [22] B. Monien and T. Tscheuschner, *On the Power of Nodes of Degree Four in the Local Max-Cut Problem*, In: Calamoneri T., Diaz J. (eds) Algorithms and Complexity. CIAC 2010. Lecture Notes in Computer Science, vol 6078. Springer, Berlin, Heidelberg.
- [23] D. Mubayi and A. Suk, *A survey of hypergraph Ramsey problems*, ArXiv e-prints, 2017.
- [24] M. Naor and L. Stockmeyer, *What can be computed locally?*, In Proc. of the 25th Annual ACM Symp. on Theory of Computing (STOC), 1993, 184–193.
- [25] S. Poljak, *Integer linear programs and local search for max-cut*, SIAM J. Comput. **21(3)** (1995), 450—465.
- [26] M. Rahman and B. Virág, *Local algorithms for independent sets are half-optimal*, Ann. Probab. **45(3)** (2017), 1543–1577.
- [27] F. P. Ramsey, *On a problem of formal logic*, Proc. Lond. Math. Soc. **30** (1930), 264–286.
- [28] J. B. Shearer, *A note on bipartite subgraphs of triangle-free graphs*, Random Structures Algorithms **3(2)** (1992), 223–226.

SCHOOL OF COMPUTER AND COMMUNICATION SCIENCES, ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE,  
SWITZERLAND

*Email address:* `etienne.bamas@epfl.ch`

LABORATOIRE G-SCOP (CNRS, UNIV. GRENOBLE ALPES), GRENOBLE, FRANCE

*Email address:* `louis.esperet@grenoble-inp.fr`