Oblivious tight compaction in $O(n)$ time with smaller constant

Samuel Dittmer[*]   Rafail Ostrovsky[†]

June 5, 2020

Abstract

Oblivious compaction is a crucial building block for hash-based oblivious RAM. Asharov et al. recently gave a $O(n)$ algorithm for oblivious tight compaction. Their algorithm is deterministic and asymptotically optimal, but it is not practical to implement because the implied constant is $\gg 2^{228}$. We give a new algorithm for oblivious tight compaction that runs in time $< 9405.73n$. As part of our construction, we give a new result in the bootstrap percolation of random regular graphs.

1   Introduction

1.1   Oblivious RAM

Oblivious RAM has been a critical object of study in cryptography since its introduction by Goldreich and Ostrovksy [13]. ORAM now arises in a number of contexts such as improving security on SGX chips [9], garbled RAM programs [19] and cloud computing [32]. Many implementations and variations of ORAM exist, including distributed ORAM, where multiple clients jointly access a shared array (see e.g. [6, 18]) and parallel ORAM (see e.g. [23]). Many variations have explored optimizations with larger block sizes or client memory, such as the well-known Path ORAM [29], which has blocks of size $\Omega(\log^2 n)$ and $O(\log n) \cdot \omega(1)$ blocks of client storage.

In this paper, we consider the classic ORAM model, with a single server, a single client, blocks of size $O(\log n)$, and $O(1)$ blocks of client storage. Recently Asharov et al. [1] gave the first ever asymptotically optimal ORAM construction under these conditions, with $O(\log n)$ amortized overhead, which they called OptORAMa. A key building block in their construction is an oblivious tight compaction algorithm with $O(n)$ runtime.

OptORAMa is noteworthy for being asymptotically optimal and deterministic, but it is not intended for practical applications. The implied constant in the big-O notation is enormous, $\gg 2^{228}$.

We construct a new oblivious tight compaction algorithm that matches the asymptotics of OptORAMa and makes a substantial improvement to the constant. We do this in part by removing the requirement that the algorithm be deterministic, and by analyzing the behavior of expander graphs under majority bootstrap percolation. Our algorithm requires $9405.73n$ steps to compact an array of length n with negligible failure probability.

---

[*]Stealth Software Technologies Inc. `samuel.dittmer@gmail.com`

[†]University of California, Los Angeles. Department of Computer Science and Mathematics. `rafail@cs.ucla.edu`

1.2 Main results

Our main theorem is our oblivious tight compaction algorithm:

Theorem 1. *There is an algorithm to obliviously tightly compact an array in* $9405.73n$ *time.*

To demonstrate the correctness of our algorithm, we need the following new result on bootstrap percolation, which we prove in Appendix C. The proof follows quickly from the expander mixing lemma and is similar to the proof of Lemma 10 in [8].

Theorem 2. *Let* G *be a d-regular Ramanujan graph. For* $d \geq 13$ *odd and* $d \geq 16$ *even there is some constant* $c < 1$ *such that every subset* $S \subseteq G$ *with* $|S|/|G| \geq c$ *is a percolating set for majority bootstrap percolation.*

1.3 Bootstrap percolation

Bootstrap percolation of a graph is type of cellular automaton where, at each time step, a vertex's behavior is determined by the behavior of its neighbors. It is an area of extensive research in computer science and combinatorics, and also arises in areas from physics studies of the Bethe lattice [7] and Ising model [28] to the study of social networks [17, 10]. One prominent area of research is on the size of percolating sets, which are sets that, if activated initially, will cause the entire graph to be activated.

One common way to phrase this question is to ask when choosing vertices from a graph with probability p gives a percolating set with high probability (see e.g. [2], [5], [31]). In the case of random regular graphs, Balogh and Pittel compute the p^* at which a transition from probability almost zero to probability almost one occurs [4]. We use their result to determine the density of bins needed for majority bootstrap percolation to occur with high probability However, we also desire bounds on initial sets above which percolation to the entire graph is guaranteed, rather than occurring with high probability. We give such bounds in Theorem 2 and prove them in Appendix C.

There has been broad interest in the question of the smallest possible percolating set for many different families of graphs, including expander graphs [8], regular graphs [14], d-dimensional hypercubes [21, 22], and powers of complete graphs [3]. Coja-Oghlan, Feige, Krivelevich and Reichman studied the size of the *largest* minimal percolating set for expander graphs [8] in the case of a constant size activation threshold. This question has also been studied for the grid [21], hypercubes [25] and trees [26].

We are interested in majority bootstrap percolation, where a vertex becomes active if a majority of its neighbors are active. Gärtner and Zehmakan gave a result equivalent to establishing an upper bound on the size of the largest minimal percolating set for expander graphs of sufficiently large degree in [12], (see also [33]) but their result is for two-way majority bootstrap percolation, a different cellular automaton.

1.4 Comparison to OptORAMa

To clearly present both the similarities to and the improvements over OptORAMa, we describe below several of the key features of OptORAMa that facilitate its asymptotic optimality and drive its large runtime. We also describe the approach our algorithm takes to address the same problems.

1. In OptoRAMa, large arrays are divided into smaller bins, where all but a constant fraction of bins have their density of marked elements bounded above by some constant. For large enough arrays (compared to W, the word size), this subdivision step happens twice. Bins are

first chosen to have size less than $(W/\log W)^2$, and then those bins are further subdivided into bins of size less than $W/\log W$. These bins are further divided into constant size buckets. This double division causes the initial density requirement to be $\rho < 1/2^{38}$.

In our algorithm, we likewise divide large arrays into much smaller bins. However, we avoid making subdivisions of already subdivided bins. We use bins of size $O(W/\log W)$, with an explicit constant we give later.

This improvement is made possible by various bin distribution schemes that only produces overly dense bins with probability $n^{-\frac{a}{\log\log n}}$ or n^{-c}. We then collect the elements from the dense bins into another large runoff array, which we must also compact.

2. In OptORAMa, private data about small bins are packed into $O(1)$ memory words, so that this data can be accessed obliviously in $O(1)$ time. We use the same bitpacking schemes in a similar way.

3. In OptORAMa, expander graphs are used to simulate random access. This is used in particular for their LooseSwapMisplaced algorithm, which controls density of marked elements. Our LooseSwap algorithm is essentially the "re-randomised" version of OptORAMa's "de-randomised" LooseSwapMisplaced algorithm.

4. In OptORAMa, expander graphs are also used as a building block for linear superconcentrators on small bins. OptoRAMa adapts the non-oblivious Pippenger construction [24] that builds from a bipartite expander graph.

   We replace linear superconcentrators with bootstrap percolation on expander graphs. This allows to distribute balls evenly between buckets without forcing us to designate a certain set of vertices as "input" or "output" vertices.

5. Estimate of OptORAMa runtime: Using the Jimbo-Maruoka expander [16], an improvement on the Margulis-Gabber-Galil in OptORAMa, requires a bipartite expander graph with degree at least d^k, where $d = 8$, and k is an integer such that $(8/(5\sqrt{2}))^k > L = 2\sqrt{(2^{25})}$. This forces $k \geq 76$, so the degree is at least 2^{228}, and so the runtime at least $2^{228}n$.

2   Structure of paper

Our algorithm is made up of three rounds of compaction, and three levels of compaction. We give an overview of the construction, including a summary of the accounting of the overall cost of the algorithm, in Section 3. We describe the top level of the construction in Section 4 and the next two levels of construction in Section 5. The descriptions of some sub-algorithms are moved to the appendices.

In the first round, we extract most marked elements, leaving behind a shortened array of size $O(n/\log^3 n)$, which can be Batcher sorted, and a sparse array of density $n^{-a/\log\log n}$, for some absolute constant $a > 0$. In the second round, we reduce the density from $n^{-a/\log\log n}$ to n^{-c}, for some constant $c > 0$. In the third round, we complete the compaction.

At the top level, we control density through LooseSwap. At the middle level we shuffle elements into bins via LooseCompact or SparseLooseCompact. At the bottom level, we perform loose compaction on bins of size $W/\log W$ using majority bootstrap percolation via BinCompact or BinCompactMargulis.

We follow [1] in using *tight compaction* to refer to moving all marked elements to the left of all unmarked elements, and *loose compaction* to refer to moving all marked elements into an array

whose size is decreased by some constant fraction. As in [1], we note that tight compaction can be reduced to loose compaction.

Our loose compaction algorithm uses an involved construction to control the density of marked elements on part of the array. It then shuffles elements from that part of the array into bins of size $\Theta(\log n)$, i.e. proportional to $W/\log W$. We then run BINCOMPACT to loosely compact each bin. This produces a sparse runoff array and a discard array populated with unmarked and dummy elements. The sparse runoff array we compact with a modified bin compaction algorithm BINCOMPACTMARGULIS. The discard array we compact by running TIGHTCOMPACT again, since the discard array produced in this iteration is made up entirely of dummy elements. We give a more detailed description of how these algorithms interact in Section 4.

In the algorithm BINCOMPACT, we shuffle each bin into buckets of constant size, and identify each bucket with the vertex of an expander graph G. Via bitpacking, we can store the private data of G in a single memory word.

By treating empty buckets as active in the sense of majority bootstrap percolation, and applying Theorem 2, we assign weights to edges of G, so that every non-empty vertex has a majority of edges pointing away from it. This algorithm is essentially nonoblivious, since we can analyze the private data of G locally. Once the edge weights are set, we perform oblivious swaps in the directions indicated by the edge weights. We give a more detailed description of BINCOMPACT in Section 5.

For sparse arrays with density bounded above by $\exp\left(-O(\log n(\log\log n)^t)\right)$, for $t \in \{-1, 0\}$, we use a variant of BINCOMPACT built on Margulis graphs that we call BINCOMPACTMARGULIS. We use here a tree traversal algorithm that is similar to bootstrap percolation in its effects. We give the details in Appendix A.

We defer proofs of a number of lemmas to Appendix B. The proof of Theorem 2 is given in Appendix C. A careful accounting of runtime cost is given in Appendix D.

3 Main algorithm

3.1 Description of algorithm

In the first round of compaction, we use TIGHTCOMPACTBYTWO to produce arrays with density $\gamma < 0.005307$. Then we shuffle into bins of size $O(\log n/\log\log n)$ in LOOSECOMPACT. By Lemma B.6, we find that the probability of density greater than γ is

$$e^{-O(\log n/\log\log n)}.$$

For each bin, we run BINCOMPACT, using a random expander graph of degree $d = 13$. By Corollary C.1, all elements will be reached by majority bootstrap percolation unless the density is greater than γ. When the density is greater than γ, we transfer the entire bin to the runoff array.

When we complete this process, we have a sorted array with $ne^{-O(\log n/\log\log n)}$ dummy elements, and a runoff array of size $14n$ and density $e^{-O(\log n/\log\log n)}$. Running SPARSELOOSECOMPACT repeatedly gives a new runoff array of size $56n$ and density n^{-c}, for some constant $c < 1$. Running SPARSELOOSECOMPACT repeatedly again gives no runoff, since by Lemma B.7 we now have negligible failure probability in BINCOMPACTMARGULIS.

We run the LOOSECOMPACT array in all three rounds to extract most of the marked elements, leaving behind three sub-arrays of size $O(n/\log^3 n)$, which can be compacted in $o(n)$ time by Batcher sort.

3.2   Runtime of main algorithm

Here we give the cost of our main algorithm by analyzing the previous section, summarizing the argument in Appendix D. The first round of the algorithm, achieving $O(e^{-O(\log n/\log\log n)})$ compaction, by § D.2, has a cost of

$$
\left[\left(5+\lg(d)+\frac{(1-(1-\gamma)^{d+1}+\varepsilon)(5d+1)-2\lg(d)}{2d+2}\right)\cdot\left((d+1)+(d^2-1)\right)+\frac{\alpha}{\beta}(d+1)\right]n
$$

$$
+\left[\left(1-\frac{\beta}{1-\alpha}\right)^{-1}(\log(1+1/d)d+(1-\alpha)\log(\beta/\gamma))+2(d^2-1)\log(1+1/d)\right]n
$$

$$
+\left(\frac{1}{\beta}-2\right)n+(d-1)(1-\frac{2}{d})\log(d/\gamma)n+\frac{2}{d+1}U(n).
$$

Taking $d=13, \alpha=0.01, \beta=0.2801, \gamma=0.005307$, gives a run time of

$$
1525.21n.
$$

We then have to run the remainder of our algorithm on the array of size $14n$ with density $e^{-O(\log n/\log\log n)}$. However, the guarantee of low density eliminates the costs from α and β. The cost for the second round is:

$$
14\left[\log(1+1/d)+\left(5+\lg(d)+\frac{o(1)}{2d+2}\right)\cdot\left((d+1)+(d^2-1)\right)\right]n
$$

$$
+\frac{2}{d+1}U(n).
$$

Similarly, the cost for the third round is:

$$
56\left[\log(1+1/d)+\left(5+\lg(d)+\frac{o(1)}{2d+2}\right)\cdot\left((d+1)+(d^2-1)\right)\right]n
$$

$$
+\frac{2}{d+1}U(n).
$$

Which gives a cost of

$$
16359.37n.
$$

Adding back in the cost of the original algorithm gives a cost of

$$
U(n)=17884.58n,
$$

for TIGHTCOMPACTBYTWO and a cost of

$$
T(n)=35769.15n
$$

for general tight compaction.

Algorithm 1 Oblivious tight compaction

Require: Array A of size n with some subset of elements marked
Ensure: Obliviously shuffle A so that all marked elements are at the front of the array
 1: **procedure** TIGHTCOMPACT(A)
 2:     $s \leftarrow 0$
 3:     **if** More than half of A is marked **then**
 4:         $s \leftarrow 1$
 5:         Switch all markings
 6:     **end if**
 7:     Temporarily mark additional elements until exactly half of A is marked
 8:     TIGHTCOMPACTBYTWO(A)
 9:     Unmark the temporarily marked elements
10:     TIGHTCOMPACT($A[1 : \frac{n}{2}]$)
11:     **if** $s = 1$ **then**
12:         Swap element in position i with position $n + 1 - i$
13:         Unswitch all markings
14:     **else**
15:         Perform dummy access to positions i and $n + 1 - i$
16:     **end if**
17: **end procedure**
---

4 Compaction algorithms

4.1 Analysis of Algorithm 1

The purpose of this algorithm is to reduce tight compaction to the case where exactly half of the elements are marked. First, we observe that compressing marked elements to the left side of an array is equivalent to compressing unmarked elements to the right half. The formal implementation of this observation is given in lines 1-5 and lines 10-15.

Next, assume less than or equal to half of the elements are marked. If we can loosely compact those elements into the left half of the array, then, by induction we can tightly compact the left half of the array. Equivalently, we mark additional elements until exactly half of the elements are marked, and then compact those elements tightly into the left half. The formal implementation is given in lines 6-9.

The correctness of this algorithm is an immediate consequence of induction and the correctness of TIGHTCOMPACTBYTWO, that is, Algorithm 2 shown in § 4.2.

4.2 Analysis of Algorithm 2

In TIGHTCOMPACTBYTWO we introduce a number of constants, to be optimized later. The constant α represents the fraction of the array on which we will perform LOOSECOMPACT, while β, γ, and γ' are bounds on density of marked elements in subarrays. The degree of the expander graph is given by d, and the word size is $W = \Theta(\log n)$.

The purpose of TIGHTCOMPACTBYTWO is to isolate a subarray $A_1 \subseteq I$ whose density of marked elements is very low. We then call LOOSECOMPACT on A_1, which in turn calls BINCOMPACT. We require A_1 to have density $\leq \gamma'$, which ensures that all but an $n^{-c/\log\log n}$ fraction of the bins created in LOOSECOMPACT will have density $\leq \gamma$, for some constant $c > 0$ (where c is a function

Algorithm 2 Tight compaction by two
***
Require: Array I of size n with exactly half of I marked, constants α, β, γ, γ', W, and d
Ensure: Obliviously shuffle I so that all marked elements are at the front of the array
 1: **procedure** TIGHTCOMPACTBYTWO($A, \alpha, \beta, \gamma, W, d$)
 2:     Divide I into its left and right half, $I = A \cup \tilde{A}$
 3:     LOOSESWAP($A, \tilde{A}, \frac{1}{2}, \frac{1}{2}, \beta$)
 4:     For every statement in lines 5-23 other than line 12, execute the statement again on $\tilde{A}$, reversing the role of marked and unmarked elements
 5:     Divide $A = A_1 \cup A_2$ with $|A_1| = \alpha n/2$ and $|A_2| = (1 - \alpha)n/2$
 6:     LOOSESWAP($A_1, A_2, \beta, \beta, \gamma'$)
 7:     $B, R, S \leftarrow \{\}$
 8:     LOOSECOMPACT($A_1, \gamma', \gamma, W, d$)
 9:     Append RUNOFF(A_1) to B
 10:     Append DISCARD(A_1) to R
 11:     **while** $|A_1| + |A_2| > n/(6 \log^3 n)$ **do**
 12:         Divide $A_2 = A_1' \cup A_2'$ so that $(1 - \alpha)(|A_1| + |A_1'|) = \alpha|A_2'|$
 13:         LOOSESWAP($A_2, \tilde{A}_2, \frac{\beta}{1-\alpha}, 1 - \frac{\beta}{1-\alpha}, \beta$)
 14:         LOOSESWAP($A_1, A_2', \gamma'(1 + \frac{1}{d}), \beta, \gamma'$)
 15:         LOOSESWAP($A_1', A_2', \beta, \beta, \gamma'$)
 16:         $A_1 \leftarrow A_1 \cup A_1'$
 17:         $A_2 \leftarrow A_2'$
 18:         LOOSECOMPACT($A_1, \gamma', \gamma, W, d$).
 19:         Append RUNOFF(A_1) to B
 20:         Append DISCARD(A_1) to R
 21:     **end while**
 22:     Divide $R = R_1 \cup R_2$ with $|R_1| = \frac{d-1}{d+1}|R|$ and $|R_2| = \frac{2}{d+1}|R|$
 23:     LOOSESWAP($R_1, R_2, \frac{1}{d}, \frac{1}{d}, \gamma'$)
 24:     **while** $|R_1| > \frac{1}{\gamma'(d+1)}|R|$ **do**
 25:         LOOSESWAP($R_1, R_2, \gamma'(1 + \frac{1}{d}), \frac{1}{2}, \gamma'$)
 26:         LOOSECOMPACT($R_1, \gamma', \gamma, W, d$)
 27:     **end while**
 28:     $R \leftarrow$ TIGHTCOMPACT($R_1 \cup R_2, \alpha, \beta, \gamma, W, d$)
 29:     $R \leftarrow$ leading $\frac{1}{d+1}|R|$ elements of R
 30:     **while** $|B| > n/(8 \log^3 n)$ **do**
 31:         SPARSELOOSECOMPACT on B
 32:         Append DISCARD(B) to S
 33:     **end while**
 34:     **while** $|S| > n/(8 \log^3 n)$ **do**
 35:         SPARSELOOSECOMPACT on S
 36:     **end while**
 37:     $B \leftarrow B \cup \tilde{S}$
 38:     $I \leftarrow$ Batcher sort of $(A \cup B) \cup (\tilde{A} \cup \tilde{B})$, sorted with marked elements first, then unmarked elements, then dummy elements
 39:     $I \leftarrow$ leading $|A| + |\tilde{A}|$ elements of I
 40:     $I \leftarrow \tilde{R} \cup I \cup R$
 41: **end procedure**
***

of γ, γ' and W; see Lemma B.6).

The algorithm relies on *exactly* half of the elements being marked. This allows us to divide the array in half and control the density in both halves. We write $I = A \cup \tilde{A}$, and treat $\tilde{A}$ as the mirror of A, reversing the role of marked and unmarked balls. To begin, we use LooseSwap to force the density of A to be less than β, which pushes the density of $\tilde{A}$ above $1 - \beta$.

We divide A into $A = A_1 \cup A_2$, with $|A_1| = \alpha|A|$, and likewise divide $\tilde{A} = \tilde{A}_1 \cup \tilde{A}_2$. We push the density of A_2 below γ' by runing LooseSwap on A_1 and A_2, and likewise push the density of $\tilde{A}_2$ above γ'. We have now increased the density of A_1 and decreased the density of $\tilde{A}_1$, so we use LooseSwap to push their densities back to β and $1 - \beta$. Note that, by Lemma B.5, the densities at every stage of the algorithm can, in principle, be computed exactly up to a term of order $1 + O(n^{-1/4})$.

We can now run LooseCompact on A_1 and $\tilde{A}_1$. This compacts A_1 by a factor of $\frac{d}{d+1}$, and so increases the density of A_1 by a factor of $(1 + \frac{1}{d})$. It compacts A by a factor of $(1 - \frac{\alpha}{d+1})$. We adjust the borders of A_1 and A_2 so that A_1 again covers α of A. We now need to run LooseSwap twice, both to push down the density of the former A_1 array from $(1 + \frac{1}{d})\gamma'$ to γ', and to push down the density of what used to be A_2 from β to γ'. We can then run LooseCompact again on the new A_1. This is the loop described in Lines 11-21.

After line 21, we must deal with two additional arrays, B, the union of RunOff(A_1) and R the union of Discard(A_1).

The array R is a mix of unmarked and dummy elements. We show in § 4.3 that the density of unmarked elements is exactly $\frac{1}{d+1}$. We no longer can exchange with the mirror $\tilde{R}$, since $\tilde{R}$ also has mostly dummy elements. However, using a slightly modified version of our loop in lines 23-26, we can compact the non-dummy elements of R into an array of size $\frac{1}{\gamma'(d+1)}|R|$, and then tightly compact the result.

The array B is sparse, with density bounded above by n^{c-1}. We run SparseLooseCompact repeatedly on B until we have compressed it by a factor of $8\log^3 n$. Note that after compressing it by this factor, B remains sparse.

This step also produces a Discard array S that is a mix of unmarked and dummy elements. We compress this array in the same way. Then $B \cup \tilde{S}$ is a sparse array of unmarked elements among dummy elements, and $\tilde{B} \cup S$ is a sparse array of marked elements among dummy elements. The number of non-dummy elements in these two arrays matches exactly the number of dummy elements in $A \cup \tilde{A}$. The Batcher sort in line 37 therefore moves all non-dummy elements to the leading $|A| + |\tilde{A}|$ positions of I. Appending $\tilde{R}$ to the front of the result and R to the end completes the compaction.

We show in § D.3, § 4.3 and § 4.4 that the probability of failure for LooseSwap, LooseCompact and SparseLooseCompact, respectively, is negligible in n for $|A| = \Omega(n/\log^3 n)$. The loops in this subsection will run at most $O(\log \log n)$ times, so by a union bound the probability of failure of this algorithm is also negligible.

4.3   Analysis of Algorithm 4

This algorithm shuffles elements into bins where we can then apply BinCompact. By Lemma B.6, the probability that a bin D_i has more than $\gamma|D_i|$ marked balls is $n^{-c/\log \log n}$, for some constant $c < 1$.

The proof of correctness is an immediate consequence of the proof of correctness of BinCompact

Algorithm 3 Loose swap misplaced

Require: Arrays A_1 and A_2 with densities of marked elements bounded above by ρ_1 and ρ_2 respectively and $|A_i| = \Omega(n/\log^3 n)$.

Ensure: Density of A_1 bounded above by $\rho_1' < \rho_1$ with all but negligible probability

1: **procedure** LOOSESWAP(A_1,A_2,ρ_1,ρ_2,ρ_1')
2:      $t \leftarrow 0$
3:      $T \leftarrow \epsilon n + |A_1| \int_{\rho_1'}^{\rho_1} \left(1 - \frac{|A_1|(\rho_1-\rho)+|A_2|\rho_2}{|A_2|}\right)^{-1} \frac{d\rho}{\rho}$
4:      **for** $t < T$ **do**
5:          $t \leftarrow t + 1$
6:          Choose random indices i_1 and i_2 from A_1 and A_2
7:          **if** $A_1[i_1]$ is marked and $A_2[i_2]$ is unmarked **then**
8:              Obliviously swap the contents of i_1 and i_2
9:          **else**
10:            Perform fake accesses
11:          **end if**
12:      **end for**
13: **end procedure**

Algorithm 4 Loose compaction

Require: Array A with $|A| = \Omega(n/\log^3 n)$ and at most $\gamma'|A|$ marked elements, for some constant $\gamma' < \gamma$.

Ensure: Compact A by a factor of $(1 - \frac{1}{d+1})$. The algorithm also returns arrays RUNOFF(A) and DISCARD(A). The array RUNOFF(A) satisfies $|A| = |$RUNOFF(A)$|$ and has $O(n^{-c})$ marked elements. The array DISCARD(A) satisfies $|$DISCARD(A)$| = |A|$. It is made up of a mix of dummy and unmarked elements.

1: **procedure** LOOSECOMPACT(A, γ, W, d)
2:      Shuffle A in the clear with Fischer-Yates.
3:      Divide A into bins $\{D_i\}$ of size $(d+1)W/(1 + \lg(d) + \frac{d}{2})$.
4:      $A \leftarrow \{\}$
5:      RUNOFF(A) $\leftarrow \{\}$
6:      DISCARD(A) $\leftarrow \{\}$
7:      **for** bin D_i **do**
8:          BINCOMPACT(D_i, W, d)
9:          Append RUNOFF(D_i) to RUNOFF(A)
10:          Append DISCARD(D_i) to DISCARD(A)
11:          Append D_i to A
12:      **end for**
13: **end procedure**

Algorithm 5 Sparse loose compaction

Require: Array A with $|A| = \Omega(n/\log^3 n)$ and at most $|A| \exp\left((c-1)\log n (\log\log n)^t\right)$ marked elements, for $c < 1$ and $t \in \{-1, 0\}$.

Ensure: Compact A by a factor of $\frac{2}{3}$. The algorithm also returns an array $\text{DISCARD}(A)$ of size $\frac{2}{3}|A|$, made up of a mix of unmarked and dummy elements, and an array $\text{RUNOFF}(A)$, with density $\exp\left((c-1)\log n (\log\log n)^{k+1}\right)$.

1: **procedure** SPARSELOOSECOMPACT(A, c, W)
2:     Shuffle A in the clear with Fischer-Yates
3:     Divide A into bins $\{D_i\}$ of size $3W/2$.
4:     $A \leftarrow \{\}$
5:     $\text{DISCARD}(A) \leftarrow \{\}$
6:     $\text{RUNOFF}(A) \leftarrow \{\}$
7:     **for** bin D_i **do**
8:         BINCOMPACTMARGULIS(D_i, W)
9:         Append D_i to A
10:       Append $\text{DISCARD}(D_i)$ to $\text{DISCARD}(A)$
11:       Append $\text{RUNOFF}(D_i)$ to $\text{RUNOFF}(A)$
12:     **end for**
13: **end procedure**
---

given in § 5.1.

4.4 Analysis of Algorithm 5

This is the sparse analogue of LOOSECOMPACT given in the previous section.

By Lemma B.7, the probability that a fixed bin has more than $\log\log n$ elements is negligible in n. By taking a union bound over the bins, the probability that at least one bin has more than $\log\log n$ marked elements is also negligible in n.

The proof of correctness relies on the proof of correctness of BINCOMPACTMARGULIS given in § A.1.

5 Bin compaction with bootstrap percolation

In this section, we give the algorithm BINCOMPACT, and related algorithms needed to implement it. The algorithm BINCOMPACTMARGULIS used for SPARSECOMPACT is similar, and we describe it in Appendix A. In both algorithms, the idea is the same: we assign elements to buckets of constant size, place each bucket at the vertex of an expander graph, and use some graph algorithm to assign directions to the edges of the graph.

5.1 Analysis of Algorithm 6

The algorithm BINCOMPACT(A, W, d) replaces A with a new array that is $(1 - \frac{1}{d+1})$ times smaller, with all marked elements from the old A moved into the new A. In this process, some unmarked elements will be removed from A. They are placed into a new array $\text{DISCARD}(A)$, which also contains a number of dummy elements.

Algorithm 6 Loose bin compaction with degree d Rucinski-Wormald graphs

---

Require: Array A with $|A| = (d+1)W/(d\log W + 1 + \lg d + \frac{d}{2})$, and some constant γ
Ensure: Compact A by a factor of $(1 - \frac{1}{d+1})$. Also returns arrays $\textsc{RunOff}(A)$, which has only dummy elements with probability $1 - n^{c-1}$, and $\textsc{Discard}(A)$, which has no marked elements and dummy elements distributed independently of A.

1: **procedure** $\textsc{BinCompact}(A, W, d)$
2:      $B \leftarrow |A|$ dummy elements
3:      **if** more than $\gamma|A|$ marked elements **then**
4:          $\textsc{RunOff}(A) \leftarrow A$
5:          $\textsc{Return } B$
6:      **end if**
7:      $k \leftarrow |A|/(d+1)$
8:      $G \leftarrow$ a d-regular expander graph with k vertices $\{v_1, \dots, v_k\}$
9:      Divide A into k bins $D_1, \dots, D_k$ of size $(d+1)$
10:     Let w be a single memory word of W bits
11:     $w \leftarrow wt(v_i) \in \{\textsc{Empty}, \textsc{Nonempty}\} \cup wt(e(v_i, v_j)) \in \{\textsc{Incoming}, \textsc{Outgoing}\} \cup \textsc{Count}(v_i) \in \{0, \dots, d\}$
12:     **for** $i \in \{1, \dots, k\}$ **do**
13:         $\textsc{Count}(v_i) = 0$
14:         **if** D_i contains no marked elements **then**
15:             $wt(v_i) \leftarrow \textsc{Empty}$
16:         **else**
17:             $wt(v_i) \leftarrow \textsc{Nonempty}$
18:         **end if**
19:     **end for**
20:     **if** proportion of vertices marked $\textsc{Empty}$ is less than $(1 - \gamma)^{d+1} - \varepsilon$ **then**
21:         $\textsc{RunOff}(A) \leftarrow A$
22:         $\textsc{Return } B$
23:     **end if**
24:     **for** $v \in V$ **do**
25:         $\textsc{Explore}(v)$
26:     **end for**
27:     **for** $i \in \{1, \dots, k\}$ **do**
28:         Add d dummy elements to D_i
29:     **end for**
30:     **for** $v_i \in V$ **do**
31:         **for** neighbors $\{v_{i_1}, \dots, v_{i_j}, \dots, v_{i_d}\}$ of v_i **do**
32:             **if** $wt(e(v_i, v_{i_j})) = \textsc{Outgoing}$ **then**
33:                 $j' \leftarrow$ index of v_i in list of neighbors of v_{i_j}
34:                 Swap element j in D_i with element $d + 1 + j'$ in D_{i_j}
35:             **else**
36:                 Perform dummy memory accesses
37:             **end if**
38:         **end for**
39:     **end for**
40:     $A, \textsc{Discard}(A) \leftarrow \{\}$
41:     **for** $i \in \{1, \dots, k\}$ **do**
42:         $\textsc{SlowCompact}(D_i)$
43:         Append leading d elements of D_i to A
44:         Append the remainder of D_i to $\textsc{Discard}(A)$
45:     **end for**
46: **end procedure**

The first time we call BINCOMPACT, we generate a d-regular expander graph G of size $|G| = W/(\log W + 1 + \lg d + \frac{d}{2})$. This takes polylogarithmic time, as shown in Lemmas B.2, B.3 and B.4.

Each time we call BINCOMPACT, we also assign to G a set of private data so that the edge set of G and its private data fit into a single memory word. The edge set and private data of G can therefore be accessed and altered without revealing the access pattern.

The array A is required to have size $|A| = (d+1)|G|$ so that we can assign $d+1$ elements of A to each vertex of G. We consider a vertex to be EMPTY if it contains no marked elements, and NONEMPTY otherwise. As part of the private data of G, for each vertex we store a single bit representing EMPTY or NONEMPTY, and another $\lg(d)$ bits to store a variable COUNT that counts the number of EMPTY neighbors. COUNT is used only in the EXPLORE subroutine.

The rest of the private data of G is made up of edge weights. We use $kd/2$ bits to assign a direction to each edge of the graph. Formally, we replace each edge e between vertices v_i and v_j with a pair of directed edges $e(v_i, v_j)$ and $e(v_j, v_i)$. To indicate an edge pointing from v_i to v_j, we assign OUTGOING to $wt(e(v_i, v_j))$ and INCOMING to $wt(e(v_j, v_i))$. To indicate an edge pointing in the opposite direction, we reverse the labels. Of course, since the weights of $e(v_i, v_j)$ and $e(v_j, v_i)$ are always opposite, it only requires one bit to store both weights. The total storage cost for the private data of G is thus

$$|G|(1 + \lg d + \tfrac{d}{2})$$

bits, which explains our choice of the size of $|G|$.

We can now describe the algorithm BINCOMPACT in terms of the private data of G. In lines 2-25, we mark vertices as EMPTY or NONEMPTY and check against certain failure conditions. In lines 26-28, we mark every edge as INCOMING or OUTGOING. Then in lines 29-47 we swap elements between vertices in the directions indicated by the edge weights, and compact each vertex's elements.

The edge weights are assigned by the subroutine EXPLORE, which guarantees that every vertex with marked elements has more OUTGOING edges than INCOMING edges. We show this in § 5.2.

The loop in lines 32-41 ensures that every bucket D_i holds at most d marked elements, and at most $d+1$ dummy elements. Indeed, vertices marked initially as EMPTY will receive at most d marked elements from their d neighbors. Vertices marked initially as NONEMPTY will have more edges marked OUTGOING than INCOMING. Note that the swap in line 37 does not necessarily transfer a marked element, but it does guarantee that position j in D_i holds an unmarked element. Let the number of OUTGOING edges be o and the number of INCOMING edges be i. After lines 33-42 the number of marked elements will be at most $d + 1 - o + i = d$. Likewise the number of dummy elements for an initially NONEMPTY vertex will be exactly $d + 1 - o + i = d$.

Now D_i has $2d+1$ elements, with at most d elements marked, so we can run the algorithm SLOWCOMPACT(D_i) described in § 5.3. We obtain an output array A of size dk holding all the marked elements from the original A along with an array DISCARD(A) of size $(d+1)k$, as desired.

BINCOMPACT fails when A begins with more than $\gamma|A|$ elements marked, or when fewer than $(1-\gamma)^{d+1} - \varepsilon$ bins contain no marked elements. Since $|A| = \Omega(W/\log W) = \Omega(\log n)$, The first event occurs with probability n^{c-1}, by Lemma B.6. The second event occurs with probability $n^{-c'}$, by Lemma B.8. Thus the failure probability of BINCOMPACT is $n^{-c''}$, for some constant $c'' > 0$. In case of failure, the contents of A are emptied into RUNOFF(A), and BINCOMPACT(A) returns an array of dummy elements.

5.2 Analysis of Explore

We desire to implement majority bootstrap percolation, where EMPTY vertices are treated as active, and NONEMPTY vertices are inactive. Then vertices that are initially NONEMPTY will be activated

Algorithm 7 Explore subroutine for Rucinski-Wormald graphs

```
 1: procedure EXPLORE(vertex v)
 2:     if wt(v) = NONEMPTY then
 3:         if COUNT = 0 then
 4:             for neighbors w of v do
 5:                 if wt(w) = EMPTY then
 6:                     COUNT ← COUNT + 1
 7:                 end if
 8:             end for
 9:         else
10:             COUNT = COUNT + 1
11:         end if
12:         if COUNT > d/2 then
13:             wt(v) ← EMPTY
14:             for neighbors w of v do
15:                 if wt(w) = EMPTY then
16:                     wt(e(v, w)) ← OUTGOING
17:                 else
18:                     wt(e(v, w)) ← INCOMING
19:                     EXPLORE(w)
20:                 end if
21:             end for
22:         end if
23:     end if
24: end procedure
```

when a majority of their neighbors are active, so that a majority of its edges will be marked as OUTGOING.

Let S be the collection of vertices marked EMPTY. We show in Corollary C.1 that, for $\frac{|S|}{|G|} \geq (1-\gamma)^{d+1}$, $|S|$ is a percolating set, i.e. majority bootstrap percolation activates every vertex.

However, running majority bootstrap percolation would take time superlinear in $|A|$. The EXPLORE subroutine is similar to majority bootstrap percolation, but activates neighbors of active vertices as they become available, rather than looping through the entire array at each time step.

Claim 5.1. *The set of vertices activated after the* EXPLORE *subroutine has been run on every vertex is the same as when we run majority bootstrap percolation on G.*

Proof. A vertex v activated after t calls to EXPLORE will certainly have been activated after t time steps of majority bootstrap percolation. Conversely, suppose that t_0 is the first t for which some vertex v activated by majority bootstrap percolation is never activated by EXPLORE. Then after time step $t_0 - 1$, a majority of v's neighbors are activated, so a majority of v's neighbors are eventually activated by EXPLORE. But this is a contradiction, since each time we activate a vertex with EXPLORE, we check each of its neighbors to see if they can be activated. This completes the proof. □

Algorithm 8 Slow loose compaction

Require: Array A of size t with at most $t/2$ marked elements
Ensure: compact A so that all marked elements are in the first $\lfloor t/2 \rfloor$ positions
 1: **procedure** SLOWCOMPACT(A)
 2:    $u \leftarrow \lfloor t/2 \rfloor$, LCOUNT $\leftarrow 0$ and RCOUNT $\leftarrow t - 2u$
 3:    **for** $i \in \{1, \ldots, u\}$ **do**
 4:        **if** $A[i]$ is marked **then**
 5:            LCOUNT $\leftarrow$ LCOUNT $+ 1$
 6:        **end if**
 7:        **if** $A[u + i]$ is marked **then**
 8:            RCOUNT $\leftarrow$ RCOUNT $+ 1$
 9:        **end if**
10:        **if** $|\text{LCOUNT} - \text{RCOUNT}| = 2$ **then**
11:            Swap $A[i]$ and $A[u + i]$
12:            LCOUNT $\leftarrow$ LCOUNT ± 1
13:            RCOUNT $\leftarrow$ RCOUNT ∓ 1
14:        **end if**
15:    **end for**
16:    SLOWCOMPACT($A[1 \ldots u]$) and SLOWCOMPACT($A[u + 1 \ldots t]$)
17:    **for** $i \in 1, \ldots, \lfloor u/2 \rfloor$ **do**
18:        Swap $A[u + i]$ and $A[u + 1 - i]$
19:    **end for**
20: **end procedure**

5.3 Analysis of Algorithm 8

The loop in lines 3-15 ensures that both the left and right half of A have at most half of their elements marked. The correctness of this algorithm follows by induction. The runtime analysis

showing $O(n \log n)$ cost is given in § D.8.

References

[1] Gilad Asharov, Ilan Komargodski, Wei-Kai Lin, Kartik Nayak, Enoch Peserico, and Elaine Shi. Optorama: Optimal oblivious ram. Cryptology ePrint Archive, Report 2018/892, 2018. https://eprint.iacr.org/2018/892.

[2] József Balogh, Béla Bollobás, and Robert Morris. Graph bootstrap percolation. *Random Structures & Algorithms*, 41(4):413–440, 2012.

[3] József Balogh, Béla Bollobás, Robert Morris, and Oliver Riordan. Linear algebra and bootstrap percolation. *Journal of Combinatorial Theory, Series A*, 119(6):1328–1335, 2012.

[4] József Balogh and Boris G Pittel. Bootstrap percolation on the random regular graph. *Random Structures & Algorithms*, 30(1-2):257–286, 2007.

[5] Milan Bradonjić and Iraj Saniee. Bootstrap percolation on random geometric graphs. *Probability in the Engineering and Informational Sciences*, 28(2):169–181, 2014.

[6] Paul Bunn, Jonathan Katz, Eyal Kushilevitz, and Rafail Ostrovsky. Efficient 3-party distributed oram. *IACR Cryptology ePrint Archive*, 2018:706, 2018.

[7] John Chalupa, Paul L Leath, and Gary R Reich. Bootstrap percolation on a bethe lattice. *Journal of Physics C: Solid State Physics*, 12(1):L31, 1979.

[8] Amin Coja-Oghlan, Uriel Feige, Michael Krivelevich, and Daniel Reichman. Contagious sets in expanders. In *Proceedings of the twenty-sixth annual ACM-SIAM symposium on Discrete algorithms*, pages 1953–1987. SIAM, 2014.

[9] Shujie Cui, Sana Belguith, Ming Zhang, Muhammad Rizwan Asghar, and Giovanni Russello. Preserving access pattern privacy in sgx-assisted encrypted search. In *2018 27th International Conference on Computer Communication and Networks (ICCCN)*, pages 1–9. IEEE, 2018.

[10] Matıas A Di Muro, Sergey V Buldyrev, and Lidia A Braunstein. Reversible bootstrap percolation: fake news and fact-checking. *arXiv preprint arXiv:1910.09516*, 2019.

[11] Joel Friedman et al. Relative expanders or weakly relatively ramanujan graphs. *Duke Mathematical Journal*, 118(1):19–35, 2003.

[12] Bernd Gärtner and Ahad N Zehmakan. Majority model on random regular graphs. In *Latin American Symposium on Theoretical Informatics*, pages 572–583. Springer, 2018.

[13] Oded Goldreich and Rafail Ostrovsky. Software protection and simulation on oblivious rams. *Journal of the ACM (JACM)*, 43(3):431–473, 1996.

[14] Alberto Guggiola and Guilhem Semerjian. Minimal contagious sets in random regular graphs. *Journal of Statistical Physics*, 158(2):300–358, 2015.

[15] Willem H Haemers. Interlacing eigenvalues and graphs. *Linear Algebra and its applications*, 226(228):593–616, 1995.

[16] Shuji Jimbo and Akira Maruoka. Expanders obtained from affine transformations. In *Proceedings of the seventeenth annual ACM symposium on Theory of computing*, pages 88–97, 1985.

[17] Marc Lelarge. Diffusion and cascading behavior in random networks. *ACM SIGMETRICS Performance Evaluation Review*, 39(3):34–36, 2011.

[18] Steve Lu and Rafail Ostrovsky. Distributed oblivious ram for secure two-party computation. In *Theory of Cryptography Conference*, pages 377–396. Springer, 2013.

[19] Steve Lu and Rafail Ostrovsky. How to garble ram programs? In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 719–734. Springer, 2013.

[20] Grigorii A Margulis. Explicit constructions of graphs without short cycles and low density codes. *Combinatorica*, 2(1):71–78, 1982.

[21] Robert Morris. Minimal percolating sets in bootstrap percolation. *the electronic journal of combinatorics*, 16(1):R2, 2009.

[22] Natasha Morrison and Jonathan A Noel. Extremal bounds for bootstrap percolation in the hypercube. *Journal of Combinatorial Theory, Series A*, 156:61–84, 2018.

[23] Kartik Nayak and Jonathan Katz. An oblivious parallel ram with o (log2 n) parallel runtime blowup. *IACR Cryptology ePrint Archive*, 2016:1141, 2016.

[24] Nicholas Pippenger. Self-routing superconcentrators. *Journal of Computer and System Sciences*, 52(1):53–60, 1996.

[25] Eric Riedl. Largest minimal percolating sets in hypercubes under 2-bootstrap percolation. *the electronic journal of combinatorics*, 17(R80):1, 2010.

[26] Eric Riedl. Largest and smallest minimal percolating sets in trees. *the electronic journal of combinatorics*, pages P64–P64, 2012.

[27] Andrzej Ruciński and Nicholas C Wormald. Random graph processes with degree restrictions. *Combinatorics, Probability and Computing*, 1(2):169–180, 1992.

[28] Sanjib Sabhapandit, Deepak Dhar, and Prabodh Shukla. Hysteresis in the random-field ising model and bootstrap percolation. *Physical Review Letters*, 88(19):197202, 2002.

[29] Emil Stefanov, Marten Van Dijk, Elaine Shi, Christopher Fletcher, Ling Ren, Xiangyao Yu, and Srinivas Devadas. Path oram: an extremely simple oblivious ram protocol. In *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*, pages 299–310, 2013.

[30] Angelika Steger and Nicholas C Wormald. Generating random regular graphs quickly. *Combinatorics, Probability and Computing*, 8(4):377–396, 1999.

[31] Tatyana S Turova and Thomas Vallier. Bootstrap percolation on a graph with random and local connections. *Journal of Statistical Physics*, 160(5):1249–1276, 2015.

[32] Peter Williams and Radu Sion. Single round access privacy on outsourced storage. In *Proceedings of the 2012 ACM conference on Computer and communications security*, pages 293–304, 2012.

[33] Ahad N Zehmakan. Opinion forming in erdős–rényi random graph and expanders. *Discrete Applied Mathematics*, 2019.

A Margulis graph compaction

A.1 Analysis of Algorithm 9

The argument here is similar to the argument in § 5.1. We explain the modifications from BINCOMPACT we make due to the sparsity of the array and the use of a Margulis graph.

The private data of G here is made up of two choices for the weight of a vertex $wt(v)$, either EMPTY or NONEMPTY. We also have three choices for the edge weights: INCOMING, OUTGOING or OFF. There are $2|G|$ edges, and each of the above variables require 2 bits, so we require $5|G|$ bits to store the private data of G. This gives $|G| = W/(4 \log W + 5)$. We use a bucket size of 3 elements from A per vertex, giving $A = 3W/(4 \log W + 5)$.

We show in § A.2 that the algorithm EXPLOREMARGULIS fails with negligible probability. The guarantee from EXPLOREMARGULIS is that every vertex initially marked NONEMPTY will have at least 3 edges marked OUTGOING, and every vertex will have at most 1 edge marked INCOMING. Therefore running lines 22-37 moves all marked elements out of $D_i[j]$, for all i and $j = 1, 2, 3$, and at most one element is moved into $D_i[4]$, for all i. Note that OUTGOING edges from vertices initially marked as EMPTY will never carry marked elements; these edges are marked only to guarantee that no vertex marked EMPTY has two or more INCOMING edges.

A.2 Analysis of Algorithm 10

Claim A.1. *Let H be the set of vertices visited when we call* SUBEXPLOREMARGULIS(v), *for some vertex v. Then H is a tree with all but negligible probability.*

Proof. Suppose by way of contradiction that $C \subseteq H$ is a cycle. By Lemma B.1, the shortest cycle of G has length at least $0.756 \log \log n$. Every vertex in C initially marked as EMPTY is adjacent to a vertex initially marked NONEMPTY. Therefore at least $1/3$ of the elements of C must be initially

Algorithm 9 Loose compaction for bins of Margulis graphs

Require: Array A with $|A| \leq 3W/(4 \log W + 5)$

Ensure: Compact A so that none of the final $\frac{2|A|}{3}$ elements are marked.

1: **procedure** BINCOMPACTMARGULIS($A, W, b = 3, \ell = 3, d = 4$)
2:     $p \leftarrow \min\{\text{prime } i \mid i^3 - i \geq |A|/3\}$
3:     $k \leftarrow p^3 - p$
4:     $G \leftarrow$ a 4-regular Margulis expander graph with k vertices, i.e. a Cayley graph for $SL_2(\mathbb{F}_p)$
5:     Label the vertices V of G as $V = \{v_1, \ldots, v_k\}$
6:     Divide A into k bins $D_1, \ldots, D_k$ of size 3, adding dummy elements if necessary
7:     Let w be a single memory word of W bits
8:     $w \quad \leftarrow \quad wt(v_i) \quad \in \quad \{\text{EMPTY}, \text{NONEMPTY}, \text{TERMINAL}\} \quad \cup \quad wt(e(v_i, v_j)) \quad \in$
   $\{\text{INCOMING}, \text{OUTGOING}, \text{OFF}\}$
9:     **for** $i \in \{1, \ldots, k\}$ **do**
10:         **if** D_i contains no marked elements **then**
11:             $wt(v_i) \leftarrow$ EMPTY
12:         **else**
13:             $wt(v_i) \leftarrow$ NONEMPTY
14:         **end if**
15:     **end for**
16:     $B \leftarrow |A|$ dummy elements
17:     **if** more than $\ell \log |G|$ NONEMPTY vertices **then**
18:         RUNOFF(A) $\leftarrow A$
19:         RETURN B
20:     **end if**
21:     **for** $v \in V$ **do**
22:         EXPLOREMARGULIS(v)
23:     **end for**
24:     **for** $i \in \{1, \ldots, k\}$ **do**
25:         Add 1 blank space to D_i
26:     **end for**
27:     **for** $v_i \in V$ **do**
28:         **for** First three neighbors $\{v_{i_1}, v_{i_2}, v_{i_3}\}$ of v_i **do**
29:             **if** $wt(e(v_i, v_{i_j})) =$ OUTGOING **then**
30:                 Swap element j in D_i with element 4 in D_{i_j}
31:             **else**
32:                 Perform dummy memory accesses
33:             **end if**
34:         **end for**
35:         **for** $j \in \{1, 2, 3\}$ **do**
36:             **if** $D_i[j]$ is marked and $wt(e(v_i, v_{i_4})) =$ OUTGOING **then**
37:                 Swap element j in D_i with element 4 in D_{i_4}
38:             **else**
39:                 Perform dummy memory accesses
40:             **end if**
41:         **end for**
42:     **end for**
43:     **for** $i \in \{1, \ldots, k\}$ **do**
44:         $A[i] \leftarrow D_i[4]$
45:     **end for**
46: **end procedure**

Algorithm 10 Explore subroutine for Margulis graphs

1: **procedure** EXPLOREMARGULIS(vertex v)
2:     **if** $wt(v) = $ NONEMPTY **then**
3:         SUBEXPLOREMARGULIS(v)
4:     **end if**
5: **end procedure**
6: **procedure** SUBEXPLOREMARGULIS(v)
7:     **if** $wt(v) = $ NONEMPTY **then**
8:         $wt(v) \leftarrow$ EMPTY
9:         **for** neighbors w of v **do**
10:             **if** $wt(e(v, w)) = $ OFF **then**
11:                 $wt(e(v, w)) \leftarrow$ OUTGOING
12:                 SUBEXPLOREMARGULIS(w)
13:             **end if**
14:         **end for**
15:     **else**
16:         **for** neighbors w of v **do**
17:             **if** $wt(e(v, w)) = $ OFF **then**
18:                 **if** $wt(w) = $ NONEMPTY **then**
19:                     $wt(e(v, w)) \leftarrow$ OUTGOING
20:                     SUBEXPLOREMARGULIS(w)
21:                 **end if**
22:             **end if**
23:         **end for**
24:     **end if**
25: **end procedure**

marked. But by Lemma B.7 the probability that there are more than $0.25 \log \log n$ marked elements in A is negligible in n. Since

$$0.25 \log \log n < \frac{0.756 \log \log n}{3}$$

this is impossible. $\qquad \square$

Since H is a tree, all edges are oriented away from the root v, and every vertex has at most one INCOMING edge. Additionally, all initially NONEMPTY vertices have all four neighbors in H, and so three OUTGOING edges. Note that all vertices in H with fewer than four edges in H were initially marked EMPTY and all neighbors of such vertices in GnH were also initially EMPTY. Since SUBEXPLOREMARGULIS never assigns a weight to an edge between two EMPTY vertices, no swaps will occur between H and GnH.

B   Combinatorial Lemmas

We state here a number of combinatorial lemmas we need throughout the paper, and give a citation if the result is well-known, or a proof otherwise.

Lemma B.1. *[Margulis [20]] There is a family of 4-regular expander graphs $\{G_k\}_{k \geq 0}$ with girth satisfying $|G_k| \geq 0.756 \log G_k$, with $G_k = (p_k^3 - p_k)$, for p_k the kth prime.*

The following result is due to Friedman [11] proving a conjecture of Alon.

Lemma B.2 (Friedman)**.** *The probability that a random d-regular graph G on n vertices, with dn even and $n > d$, has second eigenvalue $\leq 2\sqrt{d-1} + \epsilon$, is $1 - O(n^{-c})$.*

The following is due to Steger and Wormald [30], though the algorithm was first studied by Rucinski and Wormald [27].

Lemma B.3 (Steger, Wormald)**.** *There is an algorithm for generating a random d-regular graph, uniformly, in time $O(nd^2)$.*

Their algorithm is not difficult to describe. They take dn points, and, while possible, choose two points, and, if suitable, draw an edge between the corresponding vertices of G. If successful, the algorithm terminates. If unsuccessful, the algorithm resets. The following result is an immediate consequence of the well-known polynomial time algorithms for computing the characteristic polynomial and approximating roots.

Lemma B.4. *It is possible to test whether a d-regular graph G on n vertices is Ramanujan in time $poly(n)$.*

Combining these three results gives us an algorithm for generating a d-regular expander graph on n vertices in time $poly(n)$. Since our expander graphs will have $O(\log n)$ vertices, and we only need to generate $O(\log n)$ such graphs, this is sufficient for our needs.

Lemma B.5. *Consider a random walk along the integers $\mathbb{Z}$ where the probability of moving from i to $i+1$ is $f(\frac{i}{n})$ for some increasing function f. Let X be a random variable denoting the time it takes to walk from ρn to $\rho' n$. Then*

$$\Pr\left(|X - n\int_\rho^{\rho'} \frac{dx}{f(x)}| > \epsilon n^{3/4}\right) = e^{-8(\rho'-\rho)\epsilon^2\sqrt{n}}.$$

Proof. Let X_i be the time it takes to walk from ρn to $\rho' n$. From the definition of f, we have

$$\mathbb{E}[X_i] = \frac{1}{f(i/n)},$$

so that

$$\mathbb{E}[X] = \sum_{i=\rho n}^{\rho' n} \mathbb{E}[X_i] = n \sum_{i=\rho n}^{\rho' n} \frac{1}{nf(i/n)}.$$

The right hand side is a Riemann sum for $1/f(i)$ from ρ to ρ' including both the left and right endpoints, which gives

$$\frac{1}{f(\rho')} + n \int_\rho^{\rho'} \frac{dx}{f(x)} \leq n \sum_{i=\rho n}^{\rho' n} \frac{1}{nf(i/n)} \leq \frac{1}{f(\rho)} + n \int_\rho^{\rho'} \frac{dx}{f(x)}.$$

For sufficiently large n, we therefore have

$$\left| \mathbb{E}[X] - n \int_\rho^{\rho'} \frac{dx}{f(x)} \right| \leq \frac{\epsilon n^{3/4}}{2}.$$

By the Chernoff-Hoeffding Theorem, we have

$$\Pr\left(|X - \mathbb{E}[X]| \geq \frac{cn}{2} \right) \leq e^{-8(\rho'-\rho)cn},$$

for any $c \geq 0$. Setting $c = \epsilon n^{-1/4}$ and applying the triangle inequality completes the proof. $\square$

Lemma B.6. *Fix constants $\gamma' < \gamma$, and let A be an array of length n with $\gamma' n$ marked balls. Then if $k \log n / \log \log n$ balls are chosen at random from A, the probability there are more than $\gamma k \log n$ marked balls chosen is $n^{-\frac{(\gamma-\gamma')^2 k}{(\gamma+\gamma') \log \log n}}$.*

Proof. Let X be the number of marked balls chosen. The probability of $X = \ell$ is bounded above by

$$\binom{k \log n / \log \log n}{\ell} \left(\gamma' + \frac{\ell}{n} \right)^\ell \left(1 - \gamma' + \frac{\ell}{n} \right)^{k \log n - \ell}$$

$$\leq \binom{k \log n / \log \log n}{\ell} (\gamma')^\ell (1-\gamma')^{k \log n - \ell} \left(1 + \frac{k \log n}{n \gamma'} \right)^{k \log n} \left(1 + \frac{k \log n}{(1-\gamma')n} \right)^{k \log n}$$

$$\leq \binom{k \log n / \log \log n}{\ell} (\gamma')^\ell (1-\gamma')^{k \log n - \ell} \left(1 + O\left(\frac{\log^2 n}{n} \right) \right).$$

The probability of $X = \ell$ is similarly bounded below by

$$\binom{k \log n / \log \log n}{\ell} (\gamma')^\ell (1-\gamma')^{k \log n - \ell} \left(1 - O\left(\frac{\log^2 n}{n} \right) \right).$$

Therefore, up to a factor of $\left(1 \pm O\left(\frac{\log^2 n}{n} \right) \right)$, we can approximate X by a binomial random variable with $k \log n$ trials and probability γ'. Let

$$\delta = \frac{\gamma}{\gamma'} - 1.$$

21

By a Chernoff bound, we have

$$
\begin{aligned}
\Pr(X &\geq \gamma k \log n / \log \log n) \\
&= \Pr(X \geq (1 + \delta)\mathbb{E}[X]) \\
&\leq e^{-\frac{\delta^2}{2+\delta}\mathbb{E}[X]} \\
&= n^{-\frac{\delta^2 k \gamma'}{(2+\delta)\log \log n}} \\
&= n^{-\frac{(\gamma-\gamma')^2 k}{(\gamma+\gamma')\log \log n}},
\end{aligned}
$$

as desired. $\qquad\square$

Lemma B.7. *Fix a constant $c < 1$, and let A be an array of $\Omega(n/\log^3 n)$ balls with at most $|A| \exp\left((c-1)\log n (\log \log n)^t\right)$ marked balls, for $t > -2$. Then if $k \log n / \log \log n$ balls are chosen at random from A, the probability there are more than $\ell \log \log n$ marked balls chosen is $\exp\left(-O(\log n (\log \log n)^{t+1})\right)$, for every choice of $k, \ell > 0$. In particular, for $t = 0$, it is negligible in n.*

Proof. The probability a randomly selected ball is marked is $\exp\left((c-1)\log n (\log \log n)^t\right)$. Let X be the number of marked balls chosen out of $k \log n / \log \log n$. Then

$$
\begin{aligned}
\Pr(X > \log \log n) &\leq 2^{k \log n / \log \log n} \left(\exp\left((c-1)\log n (\log \log n)^t\right) + \ell \log \log n\right)^{\ell \log \log n} \\
&\leq n^{k \log 2 / \log \log n} \exp\left((c-1)\log n (\log \log n)^{t+1}\right)\left(1 + O\left(\frac{(\log \log n)^2}{n}\right)\right) \\
&= n^{k \log 2 / \log \log n + (c-1)\ell(\log \log n)^{t+1} + o(1/n)}.
\end{aligned}
$$

For $t > -2$, the $k \log 2 / \log \log n$ and $o(1/n)$ terms can be absorbed into the $-O((\log \log n)^{t+1})$ term, given the desired expression. In particular, for $t = 0$, this gives

$$
\Pr(X > \log \log n) \leq n^{-O(\log \log n)},
$$

which is negligible in n, as desired. $\qquad\square$

Lemma B.8. *Fix constants ε, γ and k. Let A be an array of kn balls, with γkn balls marked. Shuffle A and divide it into bins of size k. The probability that there are fewer than $\left((1-\gamma)^k - \varepsilon\right)n$ empty bins is e^{-cn}, for some constant c.*

Proof. We give two proofs.

In the first proof, we observe that the probability an individual bin is empty is

$$
\approx (1 - \gamma)^k.
$$

If Y_i is the event that bin i is empty, then the events $\{Y_i\}$ are negatively associated, and so we can apply a Chernoff bound to give the desired result.

In the second proof, we let X_i be random variables representing the number of bins with i marked balls in that bin, for $0 \leq i \leq k$.

Suppose that $X_i = nx_i$. We then have the following two equations:

$$
\sum x_i = 1 \tag{B.1}
$$

$$\sum i x_i = k\gamma \tag{B.2}$$

The number of ways to arrange marked balls to give $X_i = nx_i$ is

$$\binom{n}{nx_0, nx_1, \cdots, nx_k} \cdot \prod_{i=0}^{k} (\binom{k}{i})^{nx_i}.$$

We maximize this expression using Lagrange multipliers. Taking the logarithm, dividing by n, and taking the gradient gives

$$\nabla \mathbf{x} = \left(-1 - \log x_i + \log \binom{k}{i}\right)_i.$$

The method of Lagrange multipliers tells us

$$-1 - \log x_i + \log \binom{k}{i} = \lambda_A + i\lambda_B.$$

Taking $i = 0$, then $i = 1$, then i arbitrary allows us to cancel λ_A and λ_B to give

$$x_i = x_0 \frac{\binom{k}{i} x_1^i}{x_0^i}.$$

Returning to our original equations, we have

$$1 = x_0 \sum_{i=0}^{k} \binom{k}{i} (\frac{x_1}{x_0})^i = x_0 (1 + \frac{x_1}{x_0})^k$$

and

$$k\gamma = x_0 \sum_{i=0}^{k} i \binom{k}{i} (\frac{x_1}{x_0})^i = x_1 k (1 + \frac{x_1}{x_0})^{k-1}$$

. Solving gives

$$\frac{x_1}{x_0} = \frac{\gamma}{1-\gamma}$$

and

$$x_0 = (1-\gamma)^k.$$

In particular, this means for $x_0 < (1-\gamma)^k - \varepsilon$, the probability of this event occurring is exponentially smaller in n. $\qquad\square$

C   Proof of Theorem 2

Proof. Choose $c > \frac{1}{2}$ and suppose by way of contradiction there is some set U with $\frac{|U|}{|G|} > c$ that does not activate all of G under percolation. Let S be the set that is activated by U. Then we also have $\frac{|S|}{|G|} > c$, and all vertices in $T = G \backslash S$ have fewer than $\lceil \frac{d+1}{2} \rceil$ edges to S. By the expander mixing lemma (see e.g. [15]) we have

$$e(T, T) \le \frac{d|T|^2}{|G|} + \lambda |T|$$

$$\le \left(\frac{d|T|}{|G|} + 2\sqrt{d-1}\right) |T|,$$

when G is an expander graph. Taking $T = G \backslash S$ gives

$$e(T,T) \leq \left(d - \frac{dS}{G} + 2\sqrt{d-1}\right)|T|$$
$$\leq \left(d(1-c) + 2\sqrt{d-1}\right)|T|.$$

First consider the case when d is odd. We have

$$2\sqrt{d-1} < \frac{d+1}{2}$$

for $d \geq 13$, so that, for c sufficiently close to 1, we have

$$e(T,T) < \frac{d+1}{2}|T|$$

and

$$e(S,T) > \frac{d-1}{2}|T|.$$

Thus at least one vertex of T has at least $\frac{d+1}{2}$ edges to S, and can be activated, a contradiction.

The proof for d even is similar, and relies on the fact that $d = 16$ is the smallest value with

$$2\sqrt{d-1} < \frac{d}{2}.$$

$\square$

Corollary C.1. *For $d = 13$, there is some constant γ such that every set S with $|S|/|G| \geq (1-\gamma)^{d+1}$ is a percolating set for majority bootstrap percolation.*

Proof. This is an immediate consequence of Theorem 2 but we derive γ explicitly. We desire

$$\left(d - d(1-\gamma)^{d+1} + 2\sqrt{d-1}\right) < \frac{d+1}{2},$$

and so choose $\gamma < 0.005307$. $\square$

D Runtime analysis details

D.1 Runtime of Algorithm 1

To analyze the cost, suppose the cost of this algorithm is $T(n)$ and the cost of Algorithm 2 is $U(n)$. Then:

Lines 2-6 require n steps to scan and count markings, and n steps to scan again and switch markings.

Line 7 requires $U(n)$ steps.

Line 8 requires n steps.

Line 9 requires $T(n/2)$ steps.

Lines 10-14 require n steps.

We thus have

$$T(n) = 4n + U(n) + T(n/2).$$

Solving the recurrence gives

$$T(n) = 8n + \sum_{k \geq 0} U(n/2^k).$$

We show that $U(n)$ is linear in n in § D.2, so we have

$$T(n) = 8n + 2U(n).$$

D.2 Runtime of Algorithm 2

The cost of Line 3 is

$$(\frac{1}{\beta} - 2)n,$$

and the cost of Line 6 is

$$\log(\beta/\gamma)\left(1 - \frac{\beta}{1 - \alpha}\right)^{-1}\alpha n.$$

Both bounds come from the results in § D.3.

By the analysis in § D.4, the cost of Line 8 and Line 18 over the course of the loop is

$$\left(5 + \lg(d) + \frac{(1 - (1 - \gamma)^{d+1} + \varepsilon)(5d + 1) - 2\lg(d)}{2d + 2}\right)\alpha n \sum_{k \geq 0}\left(1 - \frac{\alpha}{d + 1}\right)^k$$

$$= \left(5 + \lg(d) + \frac{(1 - (1 - \gamma)^{d+1} + \varepsilon)(5d + 1) - 2\lg(d)}{2d + 2}\right)\cdot(d + 1)n$$

Lines 13-15 require again the estimates from § D.3. The cost of line 13 is

$$\sum_{k \geq 0}\left(\frac{1}{\beta} - \frac{1 - \alpha}{\beta}\right)\alpha\left(1 - \frac{\alpha}{d + 1}\right)^k n$$

$$= \frac{\alpha}{\beta}(d + 1)n.$$

The cost of line 14 is

$$\sum_{k \geq 0}\log(\gamma^{\frac{d+1}{d}}/\gamma)(1 - \frac{\beta}{1 - \alpha})^{-1}\alpha\left(1 - \frac{1}{d + 1}\right)\left(1 - \frac{\alpha}{d + 1}\right)^k$$

$$= \log(1 + 1/d)(1 - \frac{\beta}{1 - \alpha})^{-1}dn.$$

And the cost of line 15 is

$$(1 - \alpha)(1 - \frac{\beta}{1 - \alpha})^{-1}\log(\beta'/\gamma)n.$$

In lines 22-28, the size of DISCARD(A_1) is equal to the size of A_1 before running LOOSECOMPACT, so that the total size of R is

$$\sum_{k \geq 0}\left(1 - \frac{\alpha}{d + 1}\right)^k\alpha n.$$

Summing the geometric series gives $|R| = (d + 1)n$. The array R is a mix of unmarked elements and dummy elements, and we desire to compact out the unmarked elements. We show in § 5.1 that the density of unmarked elements in R is exactly $\frac{1}{d+1}$.

The cost of line 23 is

$$(d - 1)(1 - \frac{2}{d})\log(d/\gamma')n,$$

and the cost of line 25 is

$$\sum_{k \geq 0}2\log(1 + 1/d)(1 - \frac{1}{d + 1})^k(d - 1)n$$

$$= 2(d^2 - 1)\log(1 + 1/d).$$

25

The cost of line 26 is

$$\left(5 + \lg(d) + \frac{(1 - (1-\gamma)^{d+1} + \varepsilon)(5d+1) - 2\lg(d)}{2d+2}\right)(d-1)n\sum_{k \geq 0}\left(1 - \frac{1}{d+1}\right)^{k}$$

$$= \left(5 + \lg(d) + \frac{(1 - (1-\gamma)^{d+1} + \varepsilon)(5d+1) - 2\lg(d)}{2d+2}\right) \cdot (d^2 - 1)n$$

In addition, the cost of line 28 is $\frac{2}{d+1}U(n)$.

Line 30-33 requires $(3/2) \cdot 14|B|/3$ steps, by § 4.4. The size of $|B|$ is the sum of $|A_1|$ over all sizes of $|A_1|$ that arose in the previous loop. This is the same computation we did for the size of R above, so we have $|B| = (d+1)n$, and the total cost is $7 \cdot (d+1)n$.

The size of DISCARD(B) is equal to the input size of B, so that the size of S after line 33 is

$$\sum_{k \geq 0}(\frac{1}{3})^k(d+1)n = \frac{3}{2}|B|.$$

Therefore the cost from lines 34-36 is $3/2$ the cost of lines 30-33, i.e. $\frac{21}{2}(d+1)n$.

Summing all of these expressions gives, as the total cost:

$$\left[\left(5 + \lg(d) + \frac{(1 - (1-\gamma)^{d+1} + \varepsilon)(5d+1) - 2\lg(d)}{2d+2}\right) \cdot ((d+1) + (d^2 - 1)) + \frac{\alpha}{\beta}(d+1)\right]n$$

$$+ \left[\left(1 - \frac{\beta}{1-\alpha}\right)^{-1}(\log(1 + 1/d)d + (1-\alpha)\log(\beta/\gamma)) + 2(d^2 - 1)\log(1 + 1/d) + \frac{35}{2} \cdot (d+1)\right]n$$

$$+ \left(\frac{1}{\beta} - 2\right)n + (d-1)(1 - \frac{2}{d})\log(d/\gamma)n + \frac{2}{d+1}U(n).$$

The effect of the final $\frac{2}{d+1}$ term is equivalent to multiplying the rest of the runtime by $\frac{d+1}{d-1}$.

D.3   Runtime of Algorithm 3

When the density of A_1 is ρ, the number of marked balls in A_2 is $|A_1|(\rho_1 - \rho) + |A_2|(\rho_2)$. Thus the probability that a random ball from A_1 is marked is $1/\rho$ and the probability that a random ball from A_2 is unmarked is

$$\left(1 - \frac{|A_1|(\rho_1 - \rho) + |A_2|\rho_2}{|A_2|}\right)^{-1}.$$

Writing X for the number of random swaps required to force the density of A_1 to be less than ρ_1' and applying Lemma B.5 gives

$$\left|X - |A_1|\int_{\rho_1'}^{\rho_1}\left(1 - \frac{|A_1|(\rho_1 - \rho) + |A_2|\rho_2}{|A_2|}\right)^{-1}\frac{d\rho}{\rho}\right| < \epsilon n,$$

with all but negligible probability.

We bound the integral in two ways. When $\rho_2 = 1 - \rho_1$ and $|A_1| = |A_2|$, the integral simplifies to:

$$|A_1|\int_{\rho_1'}^{\rho_1}\frac{d\rho}{\rho^2} = \frac{1}{\rho_1} - \frac{1}{\rho_1'}.$$

This is the bound we use to evaluate the cost of lines 3 and 13 in Algorithm 2.

We can also bound the term

$$\left(1 - \frac{|A_1|(\rho_1 - \rho) + |A_2|\rho_2}{|A_2|}\right)^{-1} \le \left(1 - \frac{|A_1|\rho_1 + |A_2|\rho_2}{|A_2|}\right)^{-1}$$
$$= \frac{|A_2|}{|A_1|\rho_1 + |A_2|(1 - \rho_2)}.$$

This gives

$$X < \epsilon n + \frac{|A_2|}{|A_1|\rho_1 + |A_2|(1 - \rho_2)} \log\left(\frac{\rho_1}{\rho_1'}\right),$$

the estimate we use in our evaluation of lines 6, 14 and 15 of Algorithm 2.

D.4 Runtime of Algorithm 4

Line 2 requires $|A|$ steps.

Line 9 requires $\left(4 + \lg(d) + \frac{(1 - (1-\gamma)^{d+1} + \varepsilon)(5d+1) - 2\lg(d)}{2d+2}\right)|D_i|$ steps.

Thus the entire algorithm requires

$$\left(5 + \lg(d) + \frac{(1 - (1-\gamma)^{d+1} + \varepsilon)(5d+1) - 2\lg(d)}{2d+2}\right)|A|$$

steps.

D.5 Runtime of Algorithm 5

Line 2 requires $|A|$ steps.

Lines 6-10 require $11A/3$ steps, by § D.9.

This gives a total cost of $14A/3$ steps.

D.6 Runtime of Algorithm 6

The graph generation in line 8 only needs to occur once per call to LooseCompact, and by Lemmas B.2,B.3 and B.4 can be completed in polylogarithmic, i.e. $o(n)$ time.

Lines 3-23 can be accomplished with a single scan through A, followed by writing to Count(v_i) and $wt(v_i)$, which requires $|A| + 2k$ steps. If either of the runoff conditions are triggered, we must spend another $|A|$ steps obliviously swapping the contents of A into RunOff(A). This gives a total cost of $2|A| + 2k$ steps.

Lines 24-26 require $d + (1 - (1-\gamma)^{d+1} + \varepsilon)k \cdot (5d+1)/2$ steps, by § D.7.

Lines 30-39 require dk swap steps.

Line 42 requires $d\lg(d)$ steps, by § D.8, so that lines 41-44 require $dk\lg(d)$ steps.

Replacing k everywhere with $|A|/(d+1)$ gives

$$\frac{(4d+8) + 2d + (1 - (1-\gamma)^{d+1} + \varepsilon)(5d+1) + 2d + 2d\lg(d)}{2(d+1)}|A|$$
$$= \left(4 + \lg(d) + \frac{(1 - (1-\gamma)^{d+1} + \varepsilon)(5d+1) - 2\lg(d)}{2d+2}\right)|A|$$

total steps.

D.7 Runtime of Algorithm 7

The if statement on line 2 runs for every vertex, giving a cost of d steps. The rest of the algorithm only runs for the initially NONEMPTY vertices. Lines 4-8 and lines 13-21 will be run at most one time per vertex. Each loop requires one step, since we access and write over the single memory word containing the private data of G. Line 10 may be run up to $(d+1)/2$ times. This gives a total cost of $(5d+1)/2$ per NONEMPTY vertex. With all but negligible probability, there are less than $(1 - (1 - \gamma)^{d+1} + \varepsilon)$ NONEMPTY vertices, giving a total cost of

$$d + \left(1 - (1 - \gamma)^{d+1} + \varepsilon\right)\left(\frac{5d + 1}{2}\right)$$

D.8 Runtime of Algorithm 8

The cost of lines 3-15 is $t/2$ swaps. Lines 17-19 are not oblivious swaps - i.e., they could be represented by re-indexing, and cost nothing. Thus if the cost of the algorithm is $L(t)$, we have

$$L(t) = t/2 + 2L(t/2)$$

so that

$$L(t) = t/2 \lg(t).$$

D.9 Cost of Algorithm 9

The argument here is similar to the argument in § 5.1.

As in § D.6, we generate the required expander graph once, which contributes polylogarithmically to the total cost.

Lines 9-15 require $4k|D_i| = 4|A|/3$ steps, three steps to read and one step to write.

Lines 16-18 require $|G| = |A|/3$ steps.

Lines 22-37 require $6k|D_i| = 2|A|$ swap steps.

This gives a total of $11|A|/3$ steps.

D.10 Runtime of Algorithm 10

The cost of the subprocedure SUBEXPLOREMARGULIS is either 10 or 13 steps, depending on the branching statement in line 7. lines 3-16 by direct calculation, is 16 steps. Because SUBEXPLOREMARGULIS will be called at most $O(\log \log n)$ times, the total cost is $o(\log n)$. Therefore the total cost of EXPLOREMARGULIS is $(1 + o(1))|G|$.