New Algorithms and Lower Bounds for All-Pairs Max-Flow in Undirected Graphs^{*}

Amir Abboud[†]

Robert Krauthgamer[‡]

Ohad Trabelsi[§]

Abstract

We investigate the time-complexity of the All-Pairs Max-Flow problem: Given a graph with n nodes and m edges, compute for all pairs of nodes the maximum-flow value between them. If Max-Flow (the version with a given source-sink pair s, t) can be solved in time T(m), then an $O(n^2) \cdot T(m)$ is a trivial upper bound. But can we do better?

For directed graphs, recent results in fine-grained complexity suggest that this time bound is essentially optimal. In contrast, for undirected graphs with edge capacities, a seminal algorithm of Gomory and Hu (1961) runs in much faster time $O(n) \cdot T(m)$. Under the plausible assumption that Max-Flow can be solved in near-linear time $m^{1+o(1)}$, this half-century old algorithm yields an $nm^{1+o(1)}$ bound. Several other algorithms have been designed through the years, including $\tilde{O}(mn)$ time for unit-capacity edges (unconditionally), but none of them break the O(mn) barrier. Meanwhile, no super-linear lower bound was shown for undirected graphs.

We design the first hardness reductions for All-Pairs Max-Flow in undirected graphs, giving an essentially optimal lower bound for the *node-capacities* setting. For edge capacities, our efforts to prove similar lower bounds have failed, but we have discovered a surprising new algorithm that breaks the O(mn) barrier for graphs with unit-capacity edges! Assuming $T(m) = m^{1+o(1)}$, our algorithm runs in time $m^{3/2+o(1)}$ and outputs a cut-equivalent tree (similarly to the Gomory-Hu algorithm). Even with current Max-Flow algorithms we improve state-of-the-art as long as $m = O(n^{5/3-\varepsilon})$. Finally, we explain the lack of lower bounds by proving a *non-reducibility* result. This result is based on a new quasi-linear time $\tilde{O}(m)$ non-deterministic algorithm for constructing a cut-equivalent tree and may be of independent interest.

1 Introduction

In the maximum st-flow problem (abbreviated Max-Flow), the goal is to compute the maximum value of a feasible flow between a given pair of nodes s, t (sometimes called *terminals*) in an input graph.¹ Determining the time complexity of this problem is one of the most prominent open questions in fine-grained complexity and algorithms. The best running time known for directed (or undirected) graphs with nnodes, m edges, and largest integer capacity U is $\tilde{O}(\min\{m^{10/7}U^{1/7}, m\sqrt{n}\log U\})$ [Mad16, LS14], where throughout $\tilde{O}(f)$ hides logarithmic factors and stands for $O(f \log^{O(1)} f)$. To date, there is no $\Omega(m^{1+\varepsilon})$ lower bound for this problem, even when utilizing one of the popular conjectures of fine-grained complexity, such as the Strong Exponential-Time Hypothesis (SETH) of [IP01].² This gap is regularly debated among experts, and a common belief is that such a lower bound is not possible, since a near-linear-time algorithm exists but is not vet known. There is also a formal barrier for basing a lower bound for Max-Flow on SETH, as it would refute the so-called Non-deterministic SETH (NSETH) [CGI⁺16]. We will henceforth assume that Max-Flow can be solved in time $m^{1+o(1)}$, and investigate some of the most important questions that remain open under this favorable assumption. (None of our results need this assumption; it only serves for highlighting their significance.)

Perhaps the most natural next-step after the s, t version is the "all-pairs" version (abbreviated All-Pairs Max-Flow), where the goal is to solve Max-Flow for all pairs of nodes in the graph. This multi-terminal problem, dating back to 1960 [May60, Chi60], is the main focus of our work:

What is the time complexity of computing Max-Flow between all pairs of nodes?

Email:

^{*}A full version appears at arXiv:1901.01412

[†]IBM Almaden Research Center.

amir.abboud@ibm.com

[‡]Weizmann Institute of Science. Work partially supported by ONR Award N00014-18-1-2364, the Israel Science Foundation grant #1086/18, and a Minerva Foundation grant. Part of this work was done while the author was visiting the Simons Institute for the Theory of Computing. Email: robert.krauthgamer@weizmann.ac.il

[§]Weizmann Institute of Science. Work partly done at IBM Almaden. Email: ohad.trabelsi@weizmann.ac.il

¹Throughout, we focus on computing the value of the flow (rather than an actual flow), which is equal to the value of the minimum *st*-cut by the famous max-flow/min-cut theorem [FF56].

²SETH asserts that for every fixed $\varepsilon > 0$ there is an integer $k \geq 3$, such that kSAT on *n* variables and *m* clauses cannot be solved in time $2^{(1-\varepsilon)n}m^{O(1)}$.

We will discuss a few natural settings, e.g., directed vs. undirected, or node-capacities vs. edge-capacities, in which the answer to this question may vary. A trivial strategy for solving this problem (in any setting) is to invoke a T(m)-time algorithm for the s, t version $O(n^2)$ times, giving a total time bound of $O(n^2) \cdot T(m)$, which is $n^2 \cdot m^{1+o(1)}$ under our favorable assumption. But one would hope to do much better, as this all-pairs version arises in countless applications, such as a graphclustering approach for image segmentation [WL93].

In undirected edge-capacitated graphs, a seminal paper of Gomory and Hu [GH61] showed in 1961 how to solve All-Pairs Max-Flow using only n-1 calls to a Max-Flow algorithm, rather than $O(n^2)$ calls, yielding an upper bound $O(n) \cdot T(m)$. (See also [Gus90] for a different algorithm where all the n-1 calls can be executed on the original graph.) This time bound has improved over the years, following the improvements in algorithms for Max-Flow, and under our assumption it would ultimately be $n \cdot m^{1+o(1)}$. Even more surprisingly, Gomory and Hu showed that all the n^2 answers can be represented using a single tree, which can be constructed in the same time bound. Formally, A cutequivalent tree to a graph G is an edge-capacitated tree T on the same set of nodes, with the property that for every pair of nodes s, t, every minimum st-cut in T yields a bipartition of the nodes which is a minimum st-cut in G, and of the same value as in T^{3} . See also [GT01] for an experimental study, and the Encyclopedia of Algorithms [Pan16] for more background. The only algorithm that constructs a cut-equivalent tree without making $\Omega(n)$ calls to a Max-Flow algorithm was designed by Bhalgat, Hariharan, Kavitha, and Panigrahi [BHKP07]. It runs in time O(mn) in unit-capacity graphs (or equivalently, if all edges have the same capacity), and utilizes a tree-packing approach that was developed in [CH03, HKP07], inspired by classical results of [Gab95] and [Edm70]. However, if Max-Flow can indeed be computed in near-linear time, then none of the later algorithms beat by a polynomial factor the time bound $n \cdot m^{1+o(1)}$ of Gomory and Hu's half-century old algorithm.

The time complexity of All-Pairs Max-Flow becomes higher in settings where Gomory and Hu's "tree structure" [GH61] does not hold. For instance, in nodecapacitated graphs (where the flow is constrained at intermediate nodes, rather than edges) flow-equivalent trees are impossible, since there could actually exist $\Omega(n^2)$ different maximum-flow values in a single graph [HL07] (see therein also an interesting exposition of certain false claims made earlier).⁴ Directed edges make the all-pairs problem even harder; in fact, in this case node-capacities and edge-capacities are equivalent, and thus this setting does not admit flow-equivalent trees, see [May62, Jel63, HL07]. In the last decade, different algorithms were proposed to beat the trivial $O(n^2) \cdot T(m)$ time bound in these harder cases. The known bound for general graphs is $O(m^{\omega})$, due to Cheung, Lau, and Leung [CLL13], where $\omega < 2.38$ is the matrix multiplication exponent. A related version, which is obviously no harder than All-Pairs Max-Flow, is to ask (among all pairs of nodes) only for flow values that are at most k, assuming unit node capacities; for example, the case k = 1 is the transitive closure problem (reachability). For k = 2, an $O(n^{\omega})$ -time algorithm was shown in $[GGI^+17]$, and very recently a similar bound was achieved for all k = O(1) [AGI+19]. The aforementioned papers [CLL13, GGI⁺17, AGI⁺19] also present improved algorithms for acyclic graphs (DAGs). In addition, essentially optimal $\tilde{O}(n^2)$ -time algorithms were found for All-Pairs Max-Flow in certain graph families, including small treewidth [ACZ98], planar graphs [LNSW12], and surface-embedded graphs [BENW16].

The framework of fine-grained complexity has been applied to the all-pairs problem in a few recent papers, although its success has been limited to the directed case. Abboud, Vassilevska-Williams, and Yu [AVY15] proved SETH-based lower bounds for some multiterminal variants of Max-Flow, such as the single-source all-sinks version, but not all-pairs. Krauthgamer and Trabelsi [KT18] proved that All-Pairs Max-Flow cannot be solved in time $O(n^{3-\varepsilon})$, for any fixed $\varepsilon > 0$, unless SETH is false, even in the sparse regime $m = n^{1+o(1)}$. This holds also for unit-capacity graphs, and it essentially settles the complexity of the problem for directed sparse graphs, showing that the $O(n^2) \cdot T(m)$ upper bound is optimal if one assumes that $T(m) = m^{1+o(1)}$. Recently, Abboud et al. [AGI+19] proved a conditional lower bound that is even higher for dense graphs, showing that an $O(n^{\omega+1-\varepsilon})$ -time algorithm would refute the 4-Clique conjecture. However, no non-trivial lower bound is known for undirected graphs.

1.1 The Challenge of Lower Bounds in Undirected Graphs Let us briefly explain the difficulty in obtaining lower bounds for undirected graphs. Con-

³Notice that a minimum st-cut in T consists of a single edge that has minimum capacity along the unique st-path in T, and removing this edge disconnects T to two connected components. A flow-equivalent tree has the weaker property that for every pair of nodes s, t, the maximum st-flow value in T equals that in G. The key difference is that flow-equivalence maintains only the values of the flows (and thus also of the corresponding cuts).

 $^{^{4}}$ Granot and Hassin [GH86] considered a related but different notion of minimum *st*-cuts with node capacities, where an equivalent tree exists and can be computed.

sider the following folklore reduction from Boolean Matrix Multiplication (BMM) to All-Pairs Reachability in directed graphs (the aforementioned special case of All-Pairs Max-Flow with k = 1). In BMM the input is two $n \times n$ boolean matrices P and Q, and the goal is to compute the product matrix R given by

$$R(a,c) := \bigvee_{b=1}^{n} \left(P(a,b) \land Q(b,c) \right), \qquad \forall a,c \in [n].$$

Computing R can be reduced to All-Pairs Reachability as follows. Construct a graph with three layers A, B, Cwith n nodes each, where the edges are directed $A \rightarrow B \rightarrow C$ and represent the two matrices: $a \in A$ is connected to $b \in B$ iff P(a,b) = 1; and $b \in B$ is connected to $c \in C$ iff Q(b,c) = 1. It is easy to see that R(a,c) = 1 iff node $a \in A$ can reach node $c \in C$ (via a two-hop path).

This simple reduction shows an $n^{\omega-o(1)}$ lower bound for All-Pairs Reachability in dense directed graphs assuming the BMM conjecture. Higher lower bounds can be proved by more involved reductions that utilize the extra power of flow over reachability, e.g., an $n^{3-o(1)}$ lower bound in sparse directed graphs assuming SETH [KT18]. Nevertheless, this simple reduction illustrates the main difficulty in adapting such reductions to undirected graphs.

Consider the same construction but with undirected edges (i.e., without the edge orientations). The main issue is that paths from A to C can now have more than two hops – they can crisscross between two adjacent layers before moving on to the next one. Indeed, it is easy to construct examples in which the product R(a, c) = 0 but there is a path from a to c (with more than two hops). Even if we try to use the extra power of flow, giving us information about the number of paths rather than just the existence of a path, it is still unclear how to distinguish flow that uses a two-hop path (YES case) from flow that uses only longer paths (NO case).

A main technical novelty of this work is a trick to overcome this issue. The high-level idea is to design large gaps between the capacities of nodes in different layers in order to incentivize flow to move to the "next layer". Let us exhibit how this trick applies to the simple reduction above. Remove the edge orientations from our three-layer graph, and introduce node capacities, letting all nodes in B, the middle layer, have capacity 2n, and all nodes in $A \cup C$, the other two layers, have capacity 1. Now, consider the maximum flow from $a \in A$ to $c \in C$. If R(a, c) = 1 then there is a two-hop path through some $b \in B$, which can carry 2nunits of flow, hence the maximum-flow value is at least 2n. On the other hand, if R(a,c) = 0 then every path from a to c must have at least four hops, and a maximum flow must be composed of such paths. Any such path must pass through at least one node in $A \cup C \setminus \{a, c\}$, whose capacity is only 1, hence the maximum flow is bounded by $|A \cup C \setminus \{a, c\}| = 2n - 2$. This proves the same $n^{\omega - o(1)}$ lower bound as before, but now for undirected graphs with node capacities.⁵ In Section 4 we utilize this trick in a more elaborate way to prove stronger lower bounds.

1.2 Our Results Our main negative result is the first (conditional) lower bound for All-Pairs Max-Flow that holds in undirected graphs. For sparse, node-capacitated graphs we are able to match the lower bound $n^{3-o(1)}$ that was previously known only for directed graphs [KT18], and it also matches the hypothetical upper bound $n^{3+o(1)}$.

THEOREM 1.1. Assuming SETH, no algorithm can solve All-Pairs Max-Flow in undirected graphs on n nodes and O(n) edges with node capacities in $[O(n^2)]$ in time $O(n^{3-\varepsilon})$ for some fixed $\varepsilon > 0$.

Our lower bound holds even under assumptions that are weaker than SETH (see Section 4), as we reduce from the 3-Orthogonal-Vectors (3OV) problem. At a high level, it combines the trick described above for overcoming the challenge in undirected graphs, with the previous reduction of [KT18] from 3OV to the directed case. However, both of these ingredients have their own subtleties and fitting them together requires adapting and tweaking them very carefully.

Following our Theorem 1.1, the largest remaining gap in our understanding of All-Pairs Max-Flow concerns the most basic and fundamental setting: undirected graphs with edge capacities. What is the time complexity of computing a cut-equivalent tree? The upper bound has essentially been stuck at $n \cdot m^{1+o(1)}$ for more than half a century, while we cannot even rule out a near-linear $m^{1+o(1)}$ running time. To our great surprise, after a series of failed attempts at proving any lower bound, we have noticed a simple way to design a new algorithm for computing cut-equivalent trees for graphs with unit-capacities, breaking the longstanding mn barrier!

THEOREM 1.2. There is an algorithm that, given an undirected graph G with n nodes and m edges (and unit edge capacities) and parameter $1 \le d \le n$, constructs a cut-equivalent tree in time $\tilde{O}(md + \Phi(m, n, d))$, where

⁵The argument can be simplified a bit if we allow nodes of capacity 0. We also remark that restricting the flow to obey the capacities of the source and the sink makes the problem much easier; this is the version considered by Granot and Hassin [GH86] and mentioned in the previous footnote.

 $\Phi(m,n,d) = \max\left\{\sum_{i=1}^{m/d} T(m,n,F_i): F_1,\ldots,F_{m/d} \geq 0, \sum_{i=1}^{m/d} F_i \leq 2m\right\} \text{ and } T(m,n,F) \text{ is the time bound for Max-Flow on instances where whose flow value is at most a F.}$

Using the current bound on T(m, n, F) we achieve running time $\tilde{O}(m^{3/2}n^{1/6})$, and under the plausible hypothesis that $T(m, n) = m^{1+o(1)}$ our time bound becomes $m^{3/2+o(1)}$. In the regime of sparse graphs where $m = \tilde{O}(n)$ the previous best algorithm of Bhalgat et al. [BHKP07] had running time $\tilde{O}(n^2)$, whereas we achieve $\tilde{O}(n^{5/3})$, or conditionally $n^{3/2+o(1)}$. In fact, we improve on their upper bound as long as $m = O(n^{5/3-\varepsilon})$. Clearly, this also leads to improved bounds for All-Pairs Max-Flow (with unit edge capacities), for which the best strategy known is to compute the tree and then extract the answers in time $O(n^2)$.

The main open question remains: Can we prove any super-linear lower bounds for the edge capacitated case in undirected graphs? Is there an $m^{1+\varepsilon}$ lower bound under SETH for constructing a cut-equivalent tree? Perhaps surprisingly, we prove a strong barrier for the possibility of such a result.

We follow the non-reducibility framework of Carmosino et al. $[CGI^+16]$. Intuitively, if problem A is conjectured to remain hard for nondeterministic algorithms while problem B is known to become significantly easier for such algorithms, then we should not expect a reduction from A to B to exist. Such a reduction would allow the nondeterministic speedups for problem B to carry over to A. To formalize this connection, Carmosino et al. introduce NSETH: the hypothesis that SETH holds against co-nondeterministic algorithms. NSETH is plausible because it is not clear how a powerful prover could convince a sub- 2^n -time verifier that a given CNF formula is *not* satisfiable. Moreover, it is known that refuting NSETH requires new techniques since it implies new circuit lower bounds. Then, Carmosino et al. exhibited nondeterministic (and conondeterministic) speedups for problems such as 3-SUM and Max-Flow (using LP duality), showing that a reduction from SAT to these problems would refute NSETH.

Our final result builds on Theorem 1.2 to design a quasi-linear time⁶ nondeterministic algorithm for constructing a cut-equivalent tree. This algorithm can perform nondeterministic choices and in the end, outputs either a correct cut-equivalent tree or "don't know" (i.e., aborts), however we are guaranteed that for every input graph there is a at least one sequence of nondeterministic choices leads to a correct output. Our result could have applications in computation-delegation settings and may be of interest in other contexts. In particular, since our nondeterministic witness can be constructed deterministically efficiently, namely, in polynomial but super-linear time, it provides a potentially interesting *certifying algorithm* [MMNS11, ABMR11] (see [Kün18] for a recent paper with a further discussion of the connections to fine-grained complexity). Our final non-reducibility result is as follows.

THEOREM 1.3. If for some $\varepsilon > 0$ there is a deterministic fine-grained reduction proving an $\Omega(m^{1+\varepsilon})$ lower bound under SETH for constructing a cut-equivalent tree of an undirected unit edge-capacitated graph on m edges, then NSETH is false.

Our result (and this framework for non-reducibility) does not address the possibility of proving a SETH based lower bound with a randomized fine-grained reduction. This is because NSETH does not remain plausible when faced against randomization (see [CGI+16, Wil16]). That said, we are not aware of any examples where this barrier has been successfully bypassed with randomization.

Roadmap. Our main algorithm is described in the Section 2. The nondeterministic algorithm and non-reducibility result are presented in Section 3. We then present our lower bounds in Section 4. The last section discusses open questions.

2 Algorithm for a Cut-Equivalent Tree

The basic strategy in our algorithm for unit edge capacities is to handle separately nodes whose connectivity (to other nodes) is high from those whose connectivity is low. The motivation comes from the simple observation that the degree of a node is an upper bound on the maximum flow from this node to any other node in the graph. Specifically, our algorithm has two stages. The first stage uses one method (of partial trees [HKP07, BHKP07]), to compute the parts of the tree that correspond to small connectivities, and the second stage uses another method (the classical Gomory-Hu algorithm [GH61]) to complete it to a cut-equivalent tree (see Figure 1). Let us briefly review these two methods.

The Gomory-Hu algorithm. This algorithm constructs a cut-equivalent tree \mathcal{T} in iterations. Initially, \mathcal{T} is a single node associated with V (the node set of G), and the execution maintains the invariant that \mathcal{T} is a tree; each tree node i is a *super-node*, which means that it is associated with a subset $V_i \subseteq V$; and these super-nodes form a partition $V = V_1 \sqcup \cdots \sqcup V_l$. At each iteration, the algorithm picks arbitrarily two graph nodes s, t that lie in the same tree super-node i,

⁶We say that a time bound T(n) is quasi-linear if it is bounded by $O(n \log^c n)$ for some positive constant c > 0.



Figure 1: An illustration of the construction of \mathcal{T} . Left: \mathcal{T} right before the partition of the super-node V_i . Middle: after the partitioning of V_i Right: \mathcal{T} as it unfolds after the Gomory-Hu algorithm finishes.

i.e., $s, t \in V_i$. The algorithm then constructs from G an auxiliary graph G' by merging nodes that lie in the same connected component of $\mathcal{T} \setminus \{i\}$ and invokes a Max-Flow algorithm to compute in this G' a minimum st-cut, denoted C'. (For example, if the current tree is a path on super-nodes $1, \ldots, l$, then G' is obtained from G by merging $V_1 \cup \cdots \cup V_{i-1}$ into one node and $V_{i+1} \cup \cdots \cup V_l$ into another node.) The submodularity of cuts ensures that this cut is also a minimum st-cut in the original graph G, and it clearly induces a partition $V_i = S \sqcup T$ with $s \in S$ and $t \in T$. The algorithm then modifies \mathcal{T} by splitting super-node *i* into two super-nodes, one associated with S and one with T, that are connected by an edge whose weight is the value of the cut C', and further connecting each neighbor of i in \mathcal{T} to either Sor T (viewed as super-nodes), depending on its side in the minimum st-cut C' (more precisely, neighbor j is connected to the side containing V_i).

The algorithm performs these iterations until all super-nodes are singletons, and then \mathcal{T} is a weighted tree with effectively the same node set as G. It can be shown [GH61] that for every $s, t \in V$, the minimum st-cut in \mathcal{T} , viewed as a bipartition of V, is also a minimum st-cut in G, and of the same cut value. We stress that this property holds regardless of the choice made at each step of two nodes $s \neq t \in V_i$.

Partial Tree. A k-partial tree, formally defined below, can also be thought of as the result of contracting all edges of weight greater than k in a cutequivalent tree of G. Such a tree can obviously be constructed using the Gomory-Hu algorithm, but as stated below (in Lemma 2.2), faster algorithms were designed in [HKP07, BHKP07], see also [Pan16, Theorem 3]. We show below (in Lemma 2.3) that such a tree can be obtained also by a truncated execution of the Gomory-Hu algorithm, and finally we use this simple but crucial fact to prove our main theorem.

DEFINITION 2.1. (k-PARTIAL TREE [HKP07]) A k-partial tree of a graph G = (V, E) is a tree on $l \leq |V|$ super-nodes constituting a partition $V = V_1 \sqcup \cdots \sqcup V_l$, with the following property: For every two nodes $s, t \in V$ whose minimum-cut value in G is at most k, let S,T be the super-nodes for which $s \in S$ and $t \in T$, then the minimum ST-cut in the tree defines a bipartition of V which is a minimum st-cut in G and has the same value.

LEMMA 2.2. ([BHKP07]) There is an algorithm that given an undirected graph with n nodes and m edges with unit edge capacities and an integer $k \in [n]$, constructs a k-partial tree in time $\tilde{O}(mk)$.

LEMMA 2.3. Given a k-partial tree T_{low} of a graph G = (V, E), there is a truncated execution of the Gomory-Hu algorithm that produces T_{low} (i.e., its auxiliary tree \mathcal{T} becomes T_{low}).

Proof. Consider an execution of the Gomory-Hu algorithm with the following choices. At each iteration, pick any two nodes $s, t \in V$ that lie in the same super-node i of the current tree \mathcal{T} (hence they are feasible choice in a Gomory-Hu execution) but furthermore lie in different super-nodes of T_{low} , as long as such s, t exist. Then split super-node i of \mathcal{T} using the minimum st-cut induced by T_{low} (rather than an arbitrary minimum st-cut). As this cut corresponds to an edge in T_{low} , it cannot split any super-node of T_{low} , which implies, by an inductive argument, that the super-nodes of T_{low} are subsets of

ble choice for a Gomory-Hu execution. Notice also that a pair s, t as required above can be chosen as long as \mathcal{T} is not equal to T_{low} , hence the Gomory-Hu execution continues until \mathcal{T} becomes exactly T_{low} . We are now ready to prove our main theorem. *Proof.* [Proof of Theorem 1.2] Let G = (V, E) be an

input undirected graph with unit edge capacities, and denote by V_{low} all the nodes in G whose degrees are at most the chosen parameter $d \in [n]$, and by $V_{high} =$ $V \setminus V_{low}$ the nodes whose degrees are greater than d.

the super-nodes of \mathcal{T} , and thus our chosen cut is a feasi-

П

First use Lemma 2.2 to construct a *d*-partial tree T_{low} , and treat it as the auxiliary tree computed by a truncated execution of the Gomory-Hu algorithm. Then continue a Gomory-Hu execution (using this tree) to complete the construction of a cut-equivalent tree. Note that every node in V_{low} is in a singleton super-node of T_{low} , since its minimum cut value to any other node is at most d; thus a super-node V_i in T_{low} has more than one node if and only if it contains only nodes in V_{high} . Moreover, by the properties of T_{low} , two nodes have minimum-cut value greater than d if and only if they are in the same super-node V_i . Since by Lemma 2.3 there exists a truncated Gomory-Hu execution that produces T_{low} , a Gomory-Hu execution starting with T_{low} as the auxiliary tree will result in a cut-equivalent tree and the correctness follows. The running time bound follows as the first step of constructing T_{low} takes O(md) time, and the second step of the Gomory-Hu execution takes $|V_{high}|$ invocations of Max-Flow, that is running time $\sum_{i=1}^{m/d} T(m, n, F_i)$. Since every invocation of maximum st-flow with value F_i in our algorithm determines a unique edge with capacity F_i in the final cut-equivalent tree, and the sum of the capacities over all the edges of the cut-equivalent tree is at most 2m(see Claim 3.9) it holds for the invocations of Max-Flow that $\sum_{i=1}^{m/d} T(m, n, F_i) \leq 2m$. Thus, the proof of Theorem 1.2 is concluded.

We use the $T(m, n, F) = O(m^{3/4}n^{1/4}F^{1/2})$ time algorithm by [ST18] to optimize our running time. By the concavity of $F^{1/2}$, the maximum of $\sum_{i=1}^{m/d} T(m, n, F_i)$ is attained when all $F_i = d$. By setting $d = \sqrt{mn^{1/6}}$ we get $\sum_{i=1}^{\sqrt{m}/n^{1/6}} m^{3/4} n^{1/4} m^{1/4} n^{1/12} = \sum_{i=1}^{\sqrt{m}/n^{1/6}} m n^{1/3} =$ $m^{3/2}n^{1/6}$, which is faster than the known $\tilde{O}(mn)$ -time algorithm of [BHKP07] whenever $m \in [n, n^{5/3}]$

Finally, relying on a hypothetical $m^{1+o(1)}$ -time algorithm for Max-Flow, we could set $d = \sqrt{m}$ to get a total running time of $m^{1+o(1)} \cdot m/\sqrt{m} + \tilde{O}(m \cdot \sqrt{m}) \leq$ $m^{3/2+o(1)}$, as claimed immediately after Theorem 1.2.

Quasi-Linear Nondeterministic Algorithm 3 for Cut-Equivalent Tree

As no conditional lower bounds are known for the problem of constructing a cut-equivalent tree, one potentially promising approach is to design a reduction from SAT to prove that running time $n^{1+\delta-o(1)}$, for a fixed $\delta > 0$, is not possible assuming SETH. However, in this section we show that the existence of such a reduction (at least in the case of unit edge-capacities) would refute NSETH. This proves our Theorem 1.3.

Our main technical result in this section (Theorem 3.2) is a fast *nondeterministic* algorithm for constructing a cut-equivalent tree (the meaning of this notion will be formalized shortly). We then reach the conclusion about NSETH by following an argument first made in $[CGI^+16]$, however we have to rewrite their argument (rather than use their definitions and results directly), in order to adapt it from decision problems or functions (where each input has exactly one output) to total functions, since every graph has at least one cut-equivalent tree (see Section 3.2).

Generally speaking, a search problem P is a binary relation, and we say that S is a solution to instance x iff $(x, S) \in P$. Let $SOL(x) = \{S : (x, S) \in P\}$ denote the set of solutions for instance x. We say that P is a *total* function⁷ if every instance x has at least one solution, i.e., $SOL(x) \neq \emptyset$. Let \perp be the "don't know" symbol and assume that $\perp \notin \text{SOL}(x)$ for all x. For example, in our problem of constructing a cut-equivalent tree, x is a graph and SOL(x) is the set of all cut-equivalent trees for x.

DEFINITION 3.1. (NONDETERMINISTIC COMPLEXITY OF A TOTAL FUNCTION)

We say that a total function P has nondeterministic time complexity T(n) if there is a deterministic Turing Machine M such that for every instance x of P with size |x| = n:

- a. $\forall g, \mathsf{DTIME}(M(x,g)) \leq T(n), \text{ i.e., the time com-}$ plexity of M is bounded by T(n);
- b. $\exists q, M(x, q) \in SOL(x)$, i.e., at least one guess leads M to output a solution;
- c. $\forall g, M(x,g) \in \{\bot\} \cup \text{SOL}(X), i.e., every guess leads$ M to output either a solution or "don't know".

We can now state the main technical result of this section. We prove it in Section 3.1, and then use it in Section 3.2 to prove Theorem 1.3.

⁷We use this name for consistency with previous literature, although it is really a relation rather than a function.

THEOREM 3.2. The nondeterministic complexity of constructing a cut-equivalent tree for an input graph with unit edge-capacities is $\tilde{O}(m)$, where m is the number edges in the graph.

This algorithm employs the Gomory-Hu algorithm in a very specific manner, where the vertices chosen at each iteration are "centroids" (see below). The same choice was previously used by Anari and Vazirani [AV18] in the context of parallel algorithms (for planar edge-capacitated graphs), to achieve a logarithmic recursion depth, which is key for parallel time. However, since our goal is different (we want near-linear total time) we have to worry about additional issues, besides the depth of the recursion. Many auxiliary graphs must be handled throughout the execution of the algorithm, and for each one we need to verify multiple minimum cuts. This is done by guessing cuts and flows, and the main challenge is to argue that the total size of all these objects (the auxiliary graphs, and the cuts and flows within them) is only O(m). Towards overcoming this challenge, we show a basic structural result about cutequivalent trees (see Claim 3.9 below) which may have other applications. Prior to our work, it seemed unlikely that the Gomory-Hu approach could come close to near-linear time, even if Max-Flow could be computed in linear time, since a $\mathsf{Max}\text{-}\mathsf{Flow}$ computation is executed many times in many auxiliary graphs. However, our analysis shows that the total size of all these auxiliary graphs can be near-linear (if the right vertices are chosen at each iteration), giving hope that this approach may still achieve the desired upper bound.

3.1 The Nondeterministic Algorithm We now prove Theorem 3.2. Let G = (V, E) be the input graph, and let n = |V| and m = |E|.

Overview. At a high level, the nondeterministic algorithm first guesses nondeterministically a cutequivalent tree \mathcal{T}^* , and then verifies it by a (nondeterministic) process that resembles an execution of the Gomory-Hu algorithm that produces \mathcal{T}^* . Similarly to the actual Gomory-Hu algorithm, our verification process is iterative and maintains a tree \mathcal{T} of super-nodes, which means, as described in Section 2, that every tree node *i* is associated with $V_i \subseteq V$, and these super-nodes form a partition $V = V_1 \sqcup \cdots \sqcup V_l$. This tree \mathcal{T} is initialized to have a single super-node corresponding to V and then modified at each iteration, hence we shall call it the *intermediate tree*. If all guesses work well, then eventually every super-node is a singleton and the tree \mathcal{T} corresponds to \mathcal{T}^* . Otherwise (some step in the verification fails), the algorithm outputs \perp .

In a true Gomory-Hu execution, every iteration partitions some super-node into exactly two super-nodes connected by an edge (say $V_i = S \sqcup T$). In contrast, every iteration of our verification process partitions some super-node into multiple super-nodes that form a star topology, whose center is a singleton (say $V_i =$ $\{w\} \sqcup V_{i,1} \sqcup \cdots \sqcup V_{i,d}$, where super-node $\{w\}$ has edges to all super-nodes $V_{i,1}, \ldots, V_{i,d}$). We call this an *expansion* step (see Figure 2), and the node in the center of the star (i.e., w) the *expanded node*. These expansion steps will be determined from the guess \mathcal{T}^* . For example, in the extreme case that \mathcal{T}^* itself is a star, our verification process will take only one expansion step instead of |V| - 1 Gomory-Hu steps.

To prove that our algorithm is correct, we will show that every expansion step corresponds to a valid sequence of steps in the Gomory-Hu algorithm. As the latter relies on minimum-cut computations in some auxiliary graph G', also our verification will need minimumcut computations, which can be easily performed in nondeterministic linear time. However, this will not achieve overall running time O(m), because in some scenarios (e.g., in the above example where \mathcal{T}^* is a star), most of the |V| - 1 minimum-cut computations are performed on an auxiliary graph G' of size that is comparable to G, i.e., $\Omega(m)$. We overcome this obstacle using two ideas. First, we compute simultaneously all the minimum-cuts of the same expansion step in nondeterministic time that is *linear* in the size of G'. Second, we design a specific sequence of expansion steps such that the total size of all auxiliary graphs G' is O(m).

Detailed Algorithm. The algorithm first guesses nondeterministically an edge-capacitated tree \mathcal{T}^* , and then verifies, as explained below, that it is a cutequivalent tree. Here, verification means that upon the failure of any step, e.g., verifying some equality (say between the cut and flow values), the algorithm terminates with output \perp . (By the same reasoning, we may assume that all guesses are proper, e.g., a guessed tree is indeed a tree). The verification process starts by picking a sequence of nodes c_0, c_1, c_2, \ldots using the guess \mathcal{T}^* , as follows. Recall that a *centroid* of a tree is a node whose removal disconnects the tree into connected components (subtrees), each containing at most half the nodes in the tree. It is well-known that in every tree, a centroid exists and can be found in linear time. In a recursive centroid decomposition of a tree, one finds a centroid of the given tree, removes it and then repeats the process recursively in every connected component, until all remaining components are singletons (have size one). Our verification process computes this decomposition for the guess \mathcal{T}^* , which takes time $O(n \log n)$. For each recursion depth $i \geq 0$ (where clearly $i \leq \log n$), denote the set of centroids computed at depth i by $D_i \subset V$. For example, D_0



Figure 2: An illustration of the verification of a guessed tree \mathcal{T}^* . Left: the intermediate tree \mathcal{T} right before an expansion step of the node c_j in the super-node $V(\mathcal{T}^*_{c_j})$. Middle: after the expansion step (of c_j , in the dashed circle) where $U_1, ..., U_4$ are c_j 's neighbors in $\mathcal{T}^{(j+1)}$ such that $\bigcup_{i=1}^4 U_i \cup \{c_j\} = V(\mathcal{T}^*_{c_j})$. Right: the guessed cut-equivalent tree \mathcal{T}^* .

contains exactly one centroid, of the entire \mathcal{T}^* . Now let c_0, c_1, c_2, \ldots be the centroids in this decomposition in order of increasing depth, i.e., starting with the one centroid $c_0 \in D_0$, followed by the centroids from D_1 (ordered arbitrarily), and so forth. Let $\mathcal{T}^*_{c_j}$ be the subtree of \mathcal{T}^* in which the centroid c_j was computed; for example $\mathcal{T}^*_{c_0} = \mathcal{T}^*$.

OBSERVATION 3.3. For every two centroids from the same depth, namely, $c_j \neq c_{j'} \in D_i$, the corresponding subtrees $\mathcal{T}_{c_j}^*$ and $\mathcal{T}_{c_{j'}}^*$ are node disjoint.

The verification process now initializes a tree \mathcal{T} , called the intermediate tree, to consist of a single super-node associated with V, and then performs on it expansion steps for nodes c_0, c_1, c_2, \ldots (in this order) as explained below.

We now explain how to perform an expansion step for node c_j . Recall that c_j is a centroid of the subtree $\mathcal{T}_{c_i}^*$, therefore it defines a partition $V(\mathcal{T}_{c_i}^*) = \{c_j\} \sqcup U_1 \sqcup$ $\cdots \sqcup U_d$, where U_1, \ldots, U_d are the connected components after removing c_j . Notice that $d = \deg_{\mathcal{T}_{c_i}^*}(c_j) \leq$ $\deg_{\mathcal{T}^*}(c_j)$, and that each $U_k, k \in [d]$, contains exactly one node $u_k \in U_k$ that is a neighbor of c_j in $\mathcal{T}_{c_j}^*$. The expansion step replaces the super-node $V(\mathcal{T}_{c_j}^*)$ in \mathcal{T} with d + 1 super-nodes $\{c_j\}, U_1, \ldots, U_d$. (We slightly abuse notation and use a subset of nodes like $V(\mathcal{T}_{c_i}^*)$ also to refer to the super-node in \mathcal{T} associated with this subset.) These d + 1 new super-nodes are connected by a star topology, where the singleton $\{c_j\}$ at the center and each newly-added edge $(\{c_i\}, U_k)$ is set to the same capacity as the edge (c_i, u_k) in the guess \mathcal{T}^* . In addition, every edge that was incident to super-node $V(\mathcal{T}_{c_j}^*)$, say $(V(\mathcal{T}_{c_j}^*), W)$, is modified to an edge (U, W), where U is one of the new super-nodes $\{c_j\}, U_1, \ldots, U_d$, chosen according to the edge in \mathcal{T}^* that was used to set a capacity for $(V(\mathcal{T}_{c_j}^*), W)$. (We will explain how the algorithm verifies the correctness of these edge weights shortly.)

It is easy to verify that the modifications to \mathcal{T} (due to expansion steps) maintain the following property: Every super-node U in \mathcal{T} induces a subtree of \mathcal{T}^* , i.e., the induced subgraph $\mathcal{T}^*[U]$ is connected. Moreover, eventually every super-node will be a singleton, and the intermediate tree will exactly match the guess \mathcal{T}^* . When we need disambiguation, we may use $\mathcal{T}^{(j)}$ to denote the tree's state before the expansion step for c_j . For example, $\mathcal{T}^{(0)}$ is the initial tree with a single super-node V.

Informally, the verification algorithm still has to check that the capacities of the newly-added tree edges correctly represent minimum-cut values. To this end, the algorithm now constructs an auxiliary graph G'_j just as in the Gomory-Hu algorithm (see Section 2). Specifically, G'_j is constructed by taking G, and then for each connected component of $\mathcal{T}^{(j)} \setminus \{V(\mathcal{T}^*_{c_j})\}$ (i.e., after removing super-node $V(\mathcal{T}^*_{c_j})$ from $\mathcal{T}^{(j)}$), merging the nodes in (all the super-nodes in) this component into a single node. Our analysis shows (in Claim 3.6) that for all $s, t \in V(\mathcal{T}^*_{c_j})$, every minimum st-cut in the auxiliary graph G'_j is also a minimum st-cut in G. In addition, all the auxiliary graphs of a single depth q can be constructed in quasi-linear time (Lemma 3.10).

Observe that each neighbor u_k of c_j in $\mathcal{T}_{c_j}^*$ defines a (c_j, u_k) -cut in the auxiliary graph G'_j , given by the two

connected components of $\mathcal{T}^* \setminus \{(c_j, u_k)\}$. The algorithm evaluates for each u_k the capacity of this cut in G'_j , and verifies that it is equal to the capacity of the newlyadded edge $(\{c_j\}, U_k)$ (set to be the same as of edge (c_j, u_k) in \mathcal{T}^*). In fact, all these cuts evaluations are performed not sequentially but rather simultaneously for all $k \in [d]$, as follows. The key observation is that if we denote each aforementioned (c_j, u_k) -cut by $(V(G'_j) \setminus C'_k, C'_k)$, where $u_k \in C'_k$, then $\{c_j\}, C'_1, \ldots, C'_d$ are disjoint subsets of $V(G'_j)$. One can clearly evaluate the capacity of all these d cuts in a single pass over the edges of G'_j , and since each edge contributes to at most two cuts (by the disjointness), this entire pass takes only linear time $O(|E(G'_j)|)$.

Next, to verify that each (c_i, u_k) -cut exhibited above, namely, each $(V(G'_j) \setminus C'_k, C'_k)$, is actually a minimum (c_j, u_k) -cut in G'_j , the algorithm finds a flow whose value is equal to the cut capacity. In order to perform this task simultaneously for all $k \in [d]$, our verification algorithm employs a known result about disjoint trees, as a witness for maximum-flow values in a graph with unit edge-capacities (strictly speaking, this witness provides lower bounds on maximum-flow values). In the following theorem, a *directed tree rooted* at r is a directed graph arising from an undirected tree rall of whose edges are then directed away from r. This is equivalent to an arborescence (having exactly one path from r to every node other than r), however we will not require that it spans all the graph nodes. In the following, Max-Flow_G(s, t) is the maximum st-flow value in a graph G.

LEMMA 3.4. Given an undirected multigraph $H = (V_H, E_H)$, a root node $r \in V_H$, and a function $\lambda : V_H \rightarrow [|E_H|]$, it is possible to nondeterministically verify in time $\tilde{O}(|E_H|)$ that

(3.1)
$$\forall v \in V_H \setminus \{r\}, \quad Max-Flow_H(r,v) \ge \lambda(v).$$

Here, nondeterministic verification means that if (3.1) holds then there exists a guess that leads to output "yes"; and if (3.1) does not hold then every guess leads to output "no".

Proof. We use the following theorem known from [BFJ95, Theorem 2.7], in its variation from [CH03] as the Tree Packing Theorem.

THEOREM 3.5. Let H_e be an Eulerian directed graph, and r_e be a node in H_e . Then there exist $\max_{v \neq r_e} \{ Max-Flow_{H_e}(r_e, v) \}$ edge-disjoint directed trees rooted at r_e , such that each node $v \in H_e$ appears in exactly Max-Flow_{H_e}(r_e, v) trees.

Given the undirected multigraph H, first subdivide each edge into two edges with a new node in between them, then orient each edge in both directions⁸ to obtain an Eulerian directed graph H_e . Observe that the minimum-cut values between pairs of original nodes in H_e are the same as in H. Now find all maximum-flow lower-bound values from r in H_e by guessing $|V_H|$ edgedisjoint trees and then counting occurrences of each node in those trees. By Theorem 3.5, these counts correspond to maximum-flow lower-bound values from r. And so if the guessed trees support the values given by λ , then answer "yes", and otherwise answer "no". Note that the conversion to directed Eulerian graph multiplied the amount of edges by 2, and so the running time is still near linear. \Box

The verification algorithm then applies Lemma 3.4 to G'_j with c_j as the root, and verifies in time $\tilde{O}(|E(G'_j)|)$ that the maximum-flow from c_j to each u_k is at least the capacity of the (c_j, u_k) -cut exhibited above (in turn verified to be equal to the capacity of edge (c_j, u_k) in \mathcal{T}^*).

Correctness. We begin by claiming that if the guessed tree \mathcal{T}^* is a correct cut-equivalent tree of G, then our algorithm outputs \mathcal{T}^* ; we discuss the complement case afterwards. Since \mathcal{T}^* is a cut-equivalent tree, every verification step of an expansion will not fail and so the algorithm will not terminate and output \mathcal{T}^* at the end, as required.

Next, we show that if \mathcal{T}^* is not a cut equivalent tree, then our algorithm will not succeed. This is proved mainly by the claim below, which is proved in the full version, that an intermediate tree attained by expansion steps can be attained also by a sequence of Gomory-Hu steps.

CLAIM 3.6. If there is a sequence of Gomory-Hu steps simulating expansions attaining $\mathcal{T}^{(j)}$, and another expansion step is being done to attain $\mathcal{T}^{(j+1)}$, then there is a sequence of Gomory-Hu steps simulating this last step too.

Now, assume for the contrary that \mathcal{T}^* is not a cutequivalent tree of G and our algorithm still produces it. As a consequence of Claim 3.6, there is a sequence of Gomory-Hu steps attaining \mathcal{T}^* , contradicting the proof of correctness of the Gomory-Hu algorithm (which cannot produce \mathcal{T}^*). Thus, it is impossible that our algorithm finishes and produces \mathcal{T}^* , and so in one of the minimum-cut verifications after an expansion step, the cut witness inspired from \mathcal{T}^* would not be correct, or there would not be a set of directed trees to testify that the corresponding cuts are minimal. This completes the proof of correctness.

⁸The subdivision and orientation are used to transform the undirected multigraph to a directed graph.

Running Time. Observe that the running time of a single expansion step, i.e., verifying its corresponding minimum cuts by evaluating cuts and flows, is quasilinear in the size of the auxiliary graph. Thus, we only have to show that the total size of all the auxiliary graphs (over all the expansions) is quasi-linear in m. We prove in Lemma 3.7 below an O(m) bound for a single depth q, and since the depth of the decomposition is $O(\log n)$, we immediately conclude in Corollary 3.8 that the total size of all auxiliary graphs over all depths is $\tilde{O}(m)$.

LEMMA 3.7. Let $D_q = \{c_{j_1}, \ldots, c_{j_2}\}$ contain the centroids at depth q. Then the total size of $G'_{j_1}, \ldots, G'_{j_2}$ is at most O(m).

COROLLARY 3.8. The total size of all auxiliary graphs (over all depths) is $\tilde{O}(m)$.

Proof. [Proof of Lemma 3.7] Let us count for each edge $uv \in E(G)$ in how many auxiliary graphs of depth q it appears. This quantity turns out to be at most $2 + (\operatorname{dist}_{\mathcal{T}}(u, v) - 1)$, where $\operatorname{dist}_{\mathcal{T}}(u, v)$ is the hopdistance, i.e., the minimum number of edges (ignoring weights or capacities) in a path between u and v in the tree \mathcal{T} . The summand 2 comes from edges uv such that either u or v belong to $V(\mathcal{T}_{c_i})$ for some auxiliary graph G'_i . Clearly, every such edge is in at most two auxiliary graphs at depth q, because there is at most one index $j' \in D_q$ where $u \in V(\mathcal{T}_{c_{i'}})$ and at most one index $j'' \in D_q$ where $v \in V(\mathcal{T}_{c_{j''}})$. The summand $\operatorname{dist}_{\mathcal{T}}(u, v) - 1$ bounds the other appearances of edge uv, i.e., when neither u nor v belongs to some $V(\mathcal{T}_{c_i})$, and is stated in the claim below, whose proof is in the full version. While our graph has unit capacities, the claim holds for general capacities.

CLAIM 3.9. For every cut-equivalent tree \mathcal{T} of a graph G with edge capacities $c_G: E \to \mathbb{R}_+$,

$$\sum_{uv \in E(G)} c_G(u, v) \cdot \operatorname{dist}_{\mathcal{T}}(u, v) \le 2 \sum_{uv \in E(G)} c_G(u, v).$$

To complete the proof of Lemma 3.7, recall that by Observation 3.3 the super-nodes $V(\mathcal{T}_{c_{j_1}}), \ldots, V(\mathcal{T}_{c_{j_2}})$ of the same depth q are pairwise disjoint. Thus, an edge uv appears in at most $dist_{\mathcal{T}^*}(u, v) - 1$ auxiliary graphs of depth q, which totals to O(m) for all the edges in this depth according to the unit edge-capacity special case of the above Claim 3.9. This concludes Lemma 3.7.

Next, we bound the time it takes to construct all the auxiliary graphs. The proof appears in the full version.

LEMMA 3.10. The total time it takes to construct the auxiliary graphs for all the expansions in the centroid decomposition is $\tilde{O}(m)$.

3.2 Reduction from a Decision Problem to a Total Function Let us start with the formal statement of NSETH.

Hypothesis 3.11. (Nondeterministic Strong Exponential-Time Hypothesis (NSETH))

For every $\varepsilon > 0$ there exists $k = k(\varepsilon)$ such that k-TAUT (the language of all k-DNF formulas that are tautologies) is not in NTIME $(2^{n(1-\varepsilon)})$.

Note that deciding if a k-DNF formula is a tautology is equivalent to deciding if a k-CNF formula is satisfiable, thus the above hypothesis could be stated also using k-CNF appropriately. Next, we define (deterministic) fine-grained reductions from a decision problem to a total function. Note that these are Turing reductions.

DEFINITION 3.12. (FINE-GRAINED REDUCTION FROM A DECISION PROBLEM TO A TOTAL FUNCTION)

Let L be a language and P be a total function, and let $T_L(\cdot)$ and $T_P(\cdot)$ be time bounds. We say that (L, T_L) admits a fine-grained reduction to (P, T_P) if for all $\varepsilon > 0$ there is a $\gamma > 0$ and a deterministic Turing machine M^P (with an access to an oracle that generates a solution to every instance of P) such that:

- a. M^P decides L correctly on all inputs when given a correct oracle for P.
- b. Let $\tilde{Q}(M^P, x)$ denote the set of oracle queries made by M^P on input x of length n. Then the query lengths obey the bound

$$\forall x, \qquad \mathsf{DTIME}(M^P, |x|) + \sum_{q \in \tilde{Q}(M, x)} (T_P(|q|))^{1-\varepsilon}$$
$$\leq (T_L(n))^{1-\gamma}.$$

We are now ready to prove the non-reducibility result under NSETH for total functions with small nondeterministic complexity. The proof appears in the full version, and its arguments are similar to those of Carmosino et al. $[CGI^+16]$.

THEOREM 3.13. Suppose P is a total function with nondeterministic time complexity T(m). If for some $\delta > 0$ there is a deterministic fine-grained reduction from k-SAT with time-bound 2^n to P with time bound $T(m)^{1+\delta}$, i.e., from (k-SAT, $2^n)$ to $(P, T(m)^{1+\delta})$, then NSETH is false. Since the construction of a cut-equivalent tree is a total function, and by theorem 1.2 its nondeterministic complexity is $\tilde{O}(m)$, applying Theorem 3.13 implies that any deterministic reduction from SETH to the construction of a cut-equivalent tree that implies a lower bound of $\Omega(m^{1+\delta})$, for some $\delta > 0$, would refute NSETH, concluding Theorem 1.3.

4 Conditional Lower Bound for All-Pairs Max-Flow

In this section we prove a conditional lower bound for All-Pairs Max-Flow in undirected graphs with node capacities. Our construction is inspired by the one in [KT18], which was designed for directed graphs with edge capacities, but it adopts it using our new trick described in the introduction. In fact, readers familiar with the reduction in [KT18] may notice that we had to tweak it a little, making it simpler in certain ways but more complicated in others. This was necessary in order to apply our new trick successfully to it.

The starting point for our reduction is the 3OV problem.

DEFINITION 4.1. (3OV) Given three sets $U_1, U_2, U_3 \subseteq \{0, 1\}^d$ containing n binary vectors each, over dimension d, decide if there is a triple (α, β, γ) of vectors in $U_1 \times U_2 \times U_3$, whose dot product is 0. That is, a triple for which for all $i \in [d]$ at least one of $\alpha[i], \beta[i], \gamma[i]$ is equal to 0.

An adaptation of the reduction by Williams [Wil05] shows that 3OV cannot be solved in $O(n^{3-\varepsilon})$ time for any $\varepsilon > 0$ and $d = \omega(\log n)$, unless SETH is false (see [ABW15]). For us, it suffices to assume the milder conjecture that 3OV cannot be solved in $O(n^{3-\varepsilon})$ time when $d = n^{\delta}$, for all $\varepsilon, \delta > 0$. Refuting this conjecture has important implications beyond refuting SETH [GIKW17, ABDN18], e.g. it refutes the Weighted Clique Conjecture.

The high level structure of the reduction is the following: create three layers of nodes that correspond to the three sets of vectors, with additional two layers in between them that correspond to the coordinates. These additional layers help keep the number of edges small by avoiding direct edges between pairs of vectors. Among other things, we utilize the trick described in the introduction and set the capacity of the nodes in the leftmost and rightmost sides to be 1, while making the other capacities much larger. This way a flow would not gain too much from crisscrossing through these nodes. Formally, we prove the following.

LEMMA 4.2. 3OV over vector sets of size n and dimension d can be reduced to All-Pairs Max-Flow in undirected graphs with $\Theta(n \cdot d)$ nodes, $\Theta(n \cdot d)$ edges, and node capacities in $[2n^2d]$.

Proof. Given a 3OV instance F we construct a graph G with maximum flow size between some pair (among a certain set of pairs) bounded by a certain amount if and only if F is a yes instance. For simplicity, we first provide a construction that has some of the edges directed (only where we will specifically mention that), and then we show how to avoid these directions. In addition, some of the edges will be capacitated as well, however the amount of such edges is small enough so that subdividing them with appropriate capacitated nodes will work too without a significant change to the size of the constructed graph.

An Intermediate Construction with Few Directed Edges. To simplify the exposition, we start with a construction of a graph G' in which most of the edges are undirected, but some are still directed (see figure 3).

Our final graph G will be very similar to G'. It will have the same nodes and edges except that all edges will be undirected and the capacities on the nodes will be a little different.

We construct the graph G' on N nodes $V_1 \cup V_2 \cup V_3 \cup$ $A \cup B \cup \{v_B\}$. The layer V_1 contains a node α of capacity 1 for every vector $\alpha \in U_1$. V_2 contains d+1 nodes for every vector $\beta \in U_2$, d nodes denoted by β_i for every $i \in [d]$ and their capacity is 1, plus a node denoted by β' of capacity d-1. V_3 contains a node γ of capacity 1 for every vector γ in U_3 . The intermediate layer A contains 2d nodes: two nodes C_i^0 and C_i^1 of capacity n for every coordinate $i \in [d]$. The other intermediate layer B contains a node C_i of capacity *n* for every coordinate $i \in [d]$. Finally, the auxiliary node v_B has capacity n(d-1). With a slight abuse of notation, we will use the following symbols in the following ways: α will be either a node in V_1 or a vector in U_1 ; β will be a vector in U_2 ; γ will be either a node in V_3 or a vector in U_3 ; and C_i will be either a node in B or a coordinate in [d]. The usage will be clear from context.

The edges of the network will be defined as follows. First, we describe the edges that depend on the given 30V instance.

- For every α and $i \in [d]$, we add a *directed* edge from α to C_i^0 if $\alpha[i] = 0$, and a directed edge from α to C_i^1 if $\alpha[i] = 1$.
- For every β , we add an (undirected) edge from β_i to C_i if $\beta[i] = 1$.
- For every γ and $i \in [d]$, we add an (undirected) edge from C_i to γ if $\gamma[i] = 1$.



Figure 3: An illustration of part of the reduction. Here, U_1 , U_2 , and U_3 have two vectors each; α and $\tilde{\alpha}$ in U_1 , β and $\tilde{\beta}$ in U_2 , γ and $\tilde{\gamma}$ in U_3 . Bolder nodes correspond to nodes of higher capacity, and dashed edges are conditional on the input instance. For simplicity, we omit the edges not relevant to α and $\tilde{\gamma}$, and also the edges from nodes in $\{C_i^0\}_{i \in [3]}$ to nodes in $\{\beta', \tilde{\beta}'\}$. In this illustration, $\alpha = 110$, $\beta = 101$, $\tilde{\beta} = 001$, and $\tilde{\gamma} = 101$. Note that the triple α , $\tilde{\beta}$, and $\tilde{\gamma}$ has an inner product 0, and indeed the maximum flow from α to $\tilde{\gamma}$ is $2 \cdot 3 - 1 = 5$.

Moreover, there will be some (undirected) edges that are independent of the vectors. For every β , we have an edge of capacity 1 from C_i^0 to β' , and an edge of capacity 1 from C_i^1 to β_i . Also, for every β , we have an edge from β_i to β' , and an edge from β' to v_B . Finally, for every γ , we have an edge from v_B to $\gamma \in V_3$. (Unless specified otherwise, these edges have no capacity constraints.)

The graph built has $N = n+2d+n \cdot d+n+1+d+n = \Theta(nd)$ nodes, at most O(nd) edges, all of its capacities are in [N], and its construction time is O(Nd).

The following two claims whose proofs appear in the full version, prove the correctness of this intermediate reduction.

CLAIM 4.3. If every triple of vectors in (U_1, U_2, U_3) have inner product at least 1, then for all pairs $\alpha \in V_1, \gamma \in V_3$ the maximum-flow in G' is at least $n \cdot d$.

CLAIM 4.4. If there is a triple of vectors $(\alpha_{\Phi}, \beta_{\Phi}, \gamma_{\Phi}) \in (U_1, U_2, U_3)$ whose inner product is 0, then the maximum-flow in G' from $\alpha_{\Phi} \in V_1$ to $\gamma_{\Phi} \in V_3$ is at most nd - 1.

The Final Construction. The main issue with avoiding the directions on the edges between nodes in

 V_1 and A, is that additional α 's might participate in the flow as well, potentially allowing one additional unit of flow to pass through. As described in the introduction, the solution is to multiply the capacities of all nodes that are not in $V_1 \cup V_3$ by 2n. This is how we get our final graph G from G'. In the following we show how this modification concludes the proof of Lemma 4.2.

CLAIM 4.5. If every triple of vectors in (U_1, U_2, U_3) has inner product at least 1, then for all pairs $\alpha \in V_1, \gamma \in V_3$ the maximum-flow in G is at least $2n^2d$.

Proof. Since the flow that was defined in Claim 4.3 does not touch nodes in $V_1 \cup V_3$, considering the same flow in G but multiplied by 2n, we get a new flow that is of size $nd \cdot (2n)$, concluding the proof. \Box

CLAIM 4.6. If there is a triple of vectors $(\alpha_{\Phi}, \beta_{\Phi}, \gamma_{\Phi}) \in (U_1, U_2, U_3)$ whose inner product is 0, then the maximum-flow in G from $\alpha_{\Phi} \in V_1$ to $\gamma_{\Phi} \in V_3$ is at most $2n^2d - 1$.

Proof. Let f be the maximum flow from α_{Φ} to γ_{Φ} in G. The paths in f can be divided into two kinds: paths that pass through nodes in $(V_1 \cup V_3) \setminus \{\alpha_{\Phi}, \gamma_{\Phi}\}$, and

paths that do not. The total contribution of paths of the first kind can be upper bounded by the size of $(V_1 \cup V_3) \setminus \{\alpha_{\Phi}, \gamma_{\Phi}\}$, which is 2n - 2, since the capacity of all nodes in this set is 1. On the other hand, paths from the second kind must obey the directions of the directed edges in G' and can therefore be used in G', except that in G their multiplicity (the amount of flow we push through them) can be larger by a factor of 2n. Therefore, we can upper bound the total contribution of paths of the second kind by 2n times the maximum flow in G', which is (nd-1)(2n). Thus, the overall flow is at most $(nd-1)(2n)+2n-2=2n^2d-2$, which proves Claim 4.6. \Box

Since we showed a gap of at least one unit of flow between the yes and the no instances, the proof of Lemma 4.2 is concluded.

5 Open Problems

Many gaps and open questions around the complexity of maximum flow remain after this work. We highlight a few for which our intuitions may have changed following our discoveries.

- Can we break the O(mn) barrier also when the graphs have arbitrary (polynomial) capacities? Our result gives hope that this may be possible.
- Can we reduce the directed case to the undirected, node-capacitated case? Because of our lower bound, it is likely that both of these cases will end up having the same time complexity, and so such a reduction may be possible.
- Can we generalize the nondeterministic algorithm to arbitrary edge-capacities? Notice that one obstacle for achieving that goal is finding lower bounds witness for flows from a certain source to other nodes.
- Can we prove any conditional lower bound for All-Pairs Max-Flow in undirected graphs with edge capacities? This is obviously the most important and intriguing open question in this context. Our new deterministic and nondeterministic upper bounds make this task more challenging than previously thought.

6 Acknowledgements

We would like to thank Arturs Backurs for asking about the nondeterministic complexity of the problems, Marvin Kunnemann for pointing out the connection to certifying algorithms, and Richard Peng for helpful comments on the different known upper bounds for Max-Flow.

References

- [ABDN18] Amir Abboud, Karl Bringmann, Holger Dell, and Jesper Nederlof. More consequences of falsifying SETH and the orthogonal vectors conjecture. In Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, pages 253–266, 2018.
- [ABMR11] Eyad Alkassar, Sascha Böhme, Kurt Mehlhorn, and Christine Rizkallah. Verification of certifying computations. In *International Conference on Computer Aided Verification*, pages 67–82. Springer, 2011.
- [ABW15] Amir Abboud, Arturs Backurs, and Virginia Vassilevska Williams. Tight hardness results for LCS and other sequence similarity measures. In *IEEE 56th Annual Symposium on Foundations of Computer Science*, FOCS 2015, pages 59–78, 2015.
- [ACZ98] Srinivasa Rao Arikati, Shiva Chaudhuri, and Christos D. Zaroliagis. All-pairs min-cut in sparse networks. J. Algorithms, 29(1):82–110, 1998.
- [AGI⁺19] Amir Abboud, Loukas Georgiadis, Giuseppe F. Italiano, Robert Krauthgamer, Nikos Parotsidis, Ohad Trabelsi, Przemysław Uznanski, and Daniel Wolleb-Graf. Faster Algorithms for All-Pairs Bounded Min-Cuts. In 46th International Colloquium on Automata, Languages, and Programming (ICALP 2019), volume 132, pages 7:1–7:15, 2019.
- [AV18] Nima Anari and Vijay V. Vazirani. Planar graph perfect matching is in NC. In 59th IEEE Annual Symposium on Foundations of Computer Science, FOCS '18, pages 650–661. IEEE Computer Society, 2018.
- [AVY15] Amir Abboud, Virginia Vassilevska Williams, and Huacheng Yu. Matching triangles and basing hardness on an extremely popular conjecture. In Proc. of 47th STOC, pages 41–50, 2015.
- [BENW16] Glencora Borradaile, David Eppstein, Amir Nayyeri, and Christian Wulff-Nilsen. All-pairs minimum cuts in near-linear time for surface-embedded graphs. In 32nd International Symposium on Computational Geometry, SoCG 2016, pages 22:1–22:16, 2016.
- [BFJ95] Jorgen Bang-Jensen, Andras Frank, and Bill Jackson. Preserving and increasing local edge-connectivity in mixed graphs. SIAM J. Discret. Math., 8(2):155– 178, 1995.
- [BHKP07] Anand Bhalgat, Ramesh Hariharan, Telikepalli Kavitha, and Debmalya Panigrahi. An O(mn) Gomory-Hu tree construction algorithm for unweighted graphs. In 39th Annual ACM Symposium on Theory of Computing, STOC'07, pages 605–614. ACM, 2007.
- [CGI⁺16] Marco L. Carmosino, Jiawei Gao, Russell Impagliazzo, Ivan Mihajlin, Ramamohan Paturi, and Stefan Schneider. Nondeterministic extensions of the strong exponential time hypothesis and consequences for non-reducibility. In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science*, ITCS '16, pages 261–270. ACM, 2016.
- [CH03] Richard Cole and Ramesh Hariharan. A fast algorithm for computing steiner edge connectivity. In *Pro*-

- [Chi60] Robert T. Chien. Synthesis of a communication net. IBM Journal of Research and Development, 4(3):311– 320, 1960.
- [CLL13] Ho Yee Cheung, Lap Chi Lau, and Kai Man Leung. Graph connectivities, network coding, and expander graphs. SIAM Journal on Computing, 42(3):733–751, 2013.
- [Edm70] Jack Edmonds. Submodular functions, matroids, and certain polyhedra. *Combinatorial structures and their applications*, pages 69–87, 1970.
- [FF56] Lester R. Ford and Delbert R. Fulkerson. Maximal flow through a network. *Canadian journal of Mathematics*, 8(3):399–404, 1956.
- [Gab95] Harold N. Gabow. A matroid approach to finding edge connectivity and packing arborescences. J. Comput. Syst. Sci., 50(2):259–273, 1995.
- [GGI⁺17] Loukas Georgiadis, Daniel Graf, Giuseppe F. Italiano, Nikos Parotsidis, and Przemysław Uznanski. All-Pairs 2-Reachability in $O(n^{\omega} \log n)$ Time. In 44th International Colloquium on Automata, Languages, and Programming (ICALP 2017), volume 80 of Leibniz International Proceedings in Informatics (LIPIcs), pages 74:1–74:14. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2017.
- [GH61] Ralph E. Gomory and Te C. Hu. Multi-terminal network flows. Journal of the Society for Industrial and Applied Mathematics, 9:551–570, 1961.
- [GH86] Frieda Granot and Refael Hassin. Multi-terminal maximum flows in node-capacitated networks. *Discrete Applied Mathematics*, 13(2-3):157–163, 1986.
- [GIKW17] Jiawei Gao, Russell Impagliazzo, Antonina Kolokolova, and R. Ryan Williams. Completeness for first-order properties on sparse structures with algorithmic applications. In 28th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, pages 2162– 2181, 2017.
- [GT01] Andrew V. Goldberg and Kostas Tsioutsiouliklis. Cut tree algorithms: an experimental study. *Journal of Algorithms*, 38(1):51–83, 2001.
- [Gus90] Dan Gusfield. Very simple methods for all pairs network flow analysis. SIAM Journal on Computing, 19(1):143–155, 1990.
- [HKP07] Ramesh Hariharan, Telikepalli Kavitha, and Debmalya Panigrahi. Efficient algorithms for computing all low s - t edge connectivities and related problems. In Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms, pages 127–136. SIAM, 2007.
- [HL07] Refael Hassin and Asaf Levin. Flow trees for vertex-capacitated networks. Discrete Appl. Math., 155(4):572–578, 2007.
- [IP01] Russell Impagliazzo and Ramamohan Paturi. On the complexity of k-SAT. Journal of Computer and System Sciences, 62(2):367–375, March 2001.
- [Jel63] Frederick Jelinek. On the maximum number of

different entries in the terminal capacity matrix of oriented communication nets. *IEEE Transactions on Circuit Theory*, 10(2):307–308, 1963.

- [KT18] Robert Krauthgamer and Ohad Trabelsi. Conditional lower bounds for all-pairs max-flow. ACM Trans. Algorithms, 14(4):42:1–42:15, 2018.
- [Kün18] Marvin Künnemann. On nondeterministic derandomization of Freivalds' algorithm: Consequences, avenues and algorithmic progress. In 26th Annual European Symposium on Algorithms, ESA 2018, August 20-22, 2018, Helsinki, Finland, pages 56:1–56:16, 2018.
- [LNSW12] Jakub Lacki, Yahav Nussbaum, Piotr Sankowski, and Christian Wulff-Nilsen. Single source - all sinks Max Flows in planar digraphs. In Proc. of the 53rd FOCS, pages 599–608, 2012.
- [LS14] Yin Tat Lee and Aaron Sidford. Path finding methods for linear programming: Solving linear programs in $\tilde{O}(\sqrt{\text{rank}})$ iterations and faster algorithms for maximum flow. In *Proceedings of the 2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, FOCS '14, pages 424–433. IEEE Computer Society, 2014.
- [Mad16] Aleksander Madry. Computing maximum flow with augmenting electrical flows. In IEEE 57th Annual Symposium on Foundations of Computer Science, FOCS 2016, pages 593–602, 2016.
- [May60] Wataru Mayeda. Terminal and branch capacity matrices of a communication net. *IRE Transactions* on Circuit Theory, 7(3):261–269, 1960.
- [May62] Wataru Mayeda. On oriented communication nets. IRE Transactions on Circuit Theory, 9(3):261–267, 1962.
- [MMNS11] Ross M. McConnell, Kurt Mehlhorn, Stefan Näher, and Pascal Schweitzer. Certifying algorithms. *Computer Science Review*, 5(2):119–161, 2011.
- [Pan16] Debmalya Panigrahi. Gomory-Hu trees. In Ming-Yang Kao, editor, *Encyclopedia of Algorithms*, pages 858–861. Springer New York, 2016.
- [ST18] Aaron Sidford and Kevin Tian. Coordinate methods for accelerating ℓ_{∞} regression and faster approximate maximum flow. In *FOCS '18*, pages 922–933. IEEE Computer Society, 2018.
- [Wil05] R. Ryan Williams. A new algorithm for optimal 2-constraint satisfaction and its implications. *Theor. Comput. Sci.*, 348(2-3):357–365, 2005.
- [Wil16] R. Ryan Williams. Strong ETH breaks with Merlin and Arthur: Short non-interactive proofs of batch evaluation. In 31st Conference on Computational Complexity, CCC 2016, pages 2:1–2:17, 2016.
- [WL93] Zhenyu Wu and Richard Leahy. An optimal graph theoretic approach to data clustering: Theory and its application to image segmentation. *IEEE trans*actions on pattern analysis and machine intelligence, 15(11):1101–1113, 1993.