

Conditional lower bounds for sparse parameterized 2-CSP: A streamlined proof

Karthik C. S.* Dániel Marx† Marcin Pilipczuk‡ Uéverton Souza§

Abstract

Assuming the Exponential Time Hypothesis (ETH), a result of Marx (ToC'10) implies that there is no $f(k) \cdot n^{o(k/\log k)}$ time algorithm that can solve 2-CSPs with k constraints (over a domain of arbitrary large size n) for any computable function f . This lower bound is widely used to show that certain parameterized problems cannot be solved in time $f(k) \cdot n^{o(k/\log k)}$ time (assuming the ETH). The purpose of this note is to give a streamlined proof of this result.

1 Introduction

The goal of this note is to discuss a simple proof for a widely used conditional lower bound on the time needed to solve Constraint Satisfaction Problems (CSP). Our focus is on 2-CSPs, that is, the special case where each constraint involves two variables. Formally, an instance Γ of 2-CSP consists of a (constraint) graph H , an alphabet set Σ , and, for each edge $\{u, v\} \in E(H)$, a constraint $C_{uv} \subseteq \Sigma \times \Sigma$. An *assignment* for Γ is a function $\sigma : V(H) \rightarrow \Sigma$. An edge $\{u, v\} \in E(H)$ is said to be *satisfied* by an assignment σ if and only if $(\sigma(u), \sigma(v)) \in C_{uv}$. The goal of the 2-CSP is to determine if there is an assignment σ that satisfies all the edges (i.e., constraints). Such an assignment is referred to as a satisfying assignment. A 2-CSP admitting a satisfying assignment is said to be satisfiable. Some authors use the name *binary CSP* for 2-CSP, emphasizing that each constraint is defined by a binary relation on two variables. The problem can be equivalently stated as PARTITIONED SUBGRAPH ISOMORPHISM, but here we prefer to present the results using the language of 2-CSP.

Let us consider two simple examples how a 2-CSP can model other algorithmic problems. If $\Sigma = [3]$ and every constraint C_{uv} is the inequality relation (i.e., $([3] \times [3]) \setminus \{(1, 1), (2, 2), (3, 3)\}$), then the problem is the same as 3-COLORING: the instance is satisfiable if and only if the constraint graph has a proper 3-coloring of the vertices. The k -CLIQUE problem on a graph G can be reduced to 2-CSP in the following way. Let the alphabet set be $\Sigma = V(G)$ and let the constraint graph be a clique on k variables v_1, \dots, v_k . For every $1 \leq i < j \leq k$, the constraint $C_{v_i v_j}$ is the relation $\{(x, y) \in V(G) \times V(G) \mid x \text{ and } y \text{ are adjacent in } G\}$. Then it is easy to see that the 2-CSP instance is satisfiable if and only if G has a clique of size k . Note that in the second example, the alphabet size can be arbitrarily large. The lower bounds we consider are mostly relevant for such instances.

Given a 2-CSP instance with k variables, an exhaustive search algorithm can solve it by considering each of the $|\Sigma|^k$ possible assignments. In a sense, the reduction from k -CLIQUE shows that this algorithm is essentially optimal. Chen et al. [CHKX06a, CHKX06b] showed that there is no $f(k) \cdot n^{o(k)}$ time algorithm for k -CLIQUE for any computable function f , unless the Exponential-Time Hypothesis (ETH) fails (see Section 2 for the formal definition of ETH). The reduction given above translates this result to a lower bound for 2-CSP in a transparent way.

THEOREM 1.1. (CHEN ET AL. [CHKX06A, CHKX06B]) *If there is an $f(k) \cdot |\Sigma|^{o(k)}$ time algorithm for 2-CSP, where Σ is the alphabet set, k is the number of variables, and f is a computable function, then the ETH fails.*

Theorem 1.1 is essentially tight (up to a constant factor in the exponent of $|\Sigma|$), but let us observe that the reduction from k -CLIQUE creates dense 2-CSP instances having k variables and $\binom{k}{2}$ constraints. The lower bound does not remain valid if we consider sparse instances. Statements about sparse instances can be formulated by

*Rutgers University, USA, karthik.cs@rutgers.edu

†CISPA, Germany, marx@cispa.de

‡Institute of Informatics, University of Warsaw, Poland, m.pilipczuk@mimuw.edu.pl

§Universidade Federal Fluminense (UFF), Brazil, ueverton@ic.uff.br

expressing the running time as a function of the number m of constraints. As the reduction above creates an instance with $m = \Omega(k^2)$ constraints, the exponent cannot be better than $O(\sqrt{m})$.

COROLLARY 1.1. *If there is an $f(m) \cdot |\Sigma|^{o(\sqrt{m})}$ time algorithm for 2-CSP, where Σ is the alphabet set, m is the number of constraints, and f is a computable function, then the ETH fails.*

A different way to state the results about sparse instances is to consider only instances where the constraint graph has bounded degree. There is a very simple transformation that turns a 2-CSP instance with k variables and m constraints into a 2-CSP instance with $2m$ variables and a 3-regular constraint graph: if a variable appears in c constraints, then let us replace it with c copies, connect these copies with a cycle of equality constraints, and let us introduce the original constraints in such a way that each copy of a variable appears in exactly one such constraint. A further simple transformation can be used to make the constraint graph bipartite, at the cost of increasing the number of variables by a constant factor. Corollary 1.1 and these transformations give the following slightly stronger statement.

COROLLARY 1.2. *If there is an $f(k) \cdot |\Sigma|^{o(\sqrt{k})}$ time algorithm for 2-CSP on bipartite 3-regular constraint graphs, where Σ is the alphabet set, k is the number of constraints, and f is a computable function, then the ETH fails.*

Marx [Mar10] improved the lower bound to a result that is tight up to a logarithmic factor in the exponent.

THEOREM 1.2. (MARX [MAR10]) *If there is an $f(k) \cdot |\Sigma|^{o(k/\log k)}$ time algorithm for 2-CSP on bipartite 3-regular constraint graphs, where Σ is the alphabet set, k is the number of constraints, and f is a computable function, then the ETH fails.*

Theorem 1.2 was obtained as a special case of a result in a much more general setting. Informally, Marx [Mar10] considered the 2-CSP problem restricted to an arbitrary class \mathcal{H} of constraint graphs and showed that, under any such restriction, an algorithm with running time $f(k) \cdot |\Sigma|^{o(t/\log t)}$ would violate the ETH, where k is the number of variables, t is the treewidth¹ of the constraint graph, and f is an arbitrary computable function. Specializing this result to a class of 3-regular bipartite graphs where treewidth is linear in the number of vertices (explicit constructions of bounded-degree expanders can be used to construct such classes) yields Theorem 1.2.

Applications. An important aspect of Theorem 1.1 is that it can be used to obtain lower bounds for other parameterized problems. W[1]-hardness proofs are typically done by parameterized reductions from k -CLIQUE. It is easy to observe that if a parameterized reduction increases the parameter only at most by a constant multiplicative factor, then this implies a lower bound similar to Theorem 1.1 for the target problem. In many cases, the reduction from k -CLIQUE constructs an instance consisting of k vertex-selection gadgets such that each gadget increases the parameter by a constant, and thus the parameter of the target instance is indeed linear in k . However, many of the more involved reductions use edge selection gadgets (see e.g., [FHRV09]). As the k -clique has $\Theta(k^2)$ edges, this means that the reduction increases the parameter to $\Theta(k^2)$ and we can conclude only the weaker bound stating that there is no $f(k) \cdot n^{o(\sqrt{k})}$ -time algorithm for the target problem (unless the ETH fails). If we want to obtain stronger bounds on the exponent, then we have to avoid the quadratic blow-up of the parameter and do the reduction from a different problem. One possibility is to reduce from 2-CSP on a 3-regular constraint graph and use the lower bound of Theorem 1.2. In such a reduction, we would need k vertex-selection gadgets and $3k/2$ edge-selection gadgets. If each gadget increases the target parameter only by a constant, then the target parameter is again $O(k)$. Therefore, such a reduction and Theorem 1.2 allows us to conclude that there is no $f(k) \cdot n^{o(k/\log k)}$ -time algorithm for the target problem. The use of Theorem 1.2 has become a standard technique when proving almost-tight lower bounds ruling out $f(k) \cdot n^{o(k/\log k)}$ time algorithms [MP15, JKMS13, CDM17, JLR⁺17, CX15, BM16, GHNS13, PW18, BS17, BKMN16, EL18, LRSZ20, EKPS18, BGL17, CdVMdM21, CFHM20, CFM21, APSZ21, CDH21, DEG22, KSS22, EGN⁺23, AKMR19, GHNS13, FPRS23, CCG⁺13, BIJK19, BCMP20, NS22]. It seems that some kind of edge representation is required in many W[1]-hardness proofs, and for such problems basing the reduction on Theorem 1.1 would be able to rule out only algorithms with running time $f(k) \cdot n^{o(\sqrt{k})}$. For these problems, Theorem 1.2 is the only known way of obtaining a lower bound (almost) matching the $n^{O(k)}$ upper bounds.

¹Informally, treewidth is a measure of the proximity of the graph to a tree, and more formally is the size of the largest vertex set in a tree decomposition of the graph.

Our contribution. Our main contribution is to provide a streamlined proof of Theorem 1.2 and state some stronger formulations. As mentioned above, Theorem 1.2 was obtained by considering the special case of expander graphs in the more general result of [Mar10]. It turns out that the proof can be significantly shortened in this case: a multi-step combinatorial proof can be replaced by a known result about routing in expanders with congestion $O(\log n)$ [LR99].

Instead of saying that the exponent cannot be $o(k/\log k)$, we state the result in a stronger form by stating a lower bound for every fixed k . A similar formulation was stated by Cohen-Addad et al. [CdVMdM21]. It seems that such a statement is particularly useful to obtain in a clear way lower bounds that involve multiple parameters (where the little- o notation is unclear).

THEOREM 1.3. *There exists $\alpha > 0$ and $k_0 \in \mathbb{N}$ such that the following holds. If there exists some fixed integer $k \geq k_0$ and an algorithm \mathcal{A}_k with the following guarantees, then ETH is false.*

Input: \mathcal{A}_k takes as input a 2-CSP instance $\Gamma(H, \Sigma, \{C_{uv}\}_{\{u,v\} \in E(H)})$ where H is a 3-regular simple bipartite graph on k vertices.

Output: \mathcal{A}_k outputs 1 if and only if there is a satisfying assignment to Γ .

Run time: \mathcal{A}_k runs in time $O(|\Sigma|^{\alpha k / \log k})$.

Clearly, Theorem 1.3 implies Theorem 1.2: for a fixed sufficiently large value of k , the hypothetical algorithm of Theorem 1.2 would satisfy the guarantees in Theorem 4.2.

Jaffke et al. [JLM⁺23] proved a version of Theorem 1.2 where there is an upper bound on the alphabet size by a function of k ; this result was used to establish a lower bound for the problem GLOBAL LABEL MINCUT. Intuitively, if we set the alphabet size to $|\Sigma| = 2^k$, then we expect that $|\Sigma|^{O(k)} = 2^{O(k^2)}$ is the best possible running time. However, the formulations of Theorems 1.2 and 1.3 cannot be used to establish such results, as the number of possible instances with this restriction is bounded for each value of k . The following theorem proves bounds of this form; for example, it shows that, assuming ETH, there is no $2^{o(k^2/\log k)}$ time algorithm even with the restriction $|\Sigma| \leq 2^k$.

THEOREM 1.4. *Let $f : \mathbb{N} \rightarrow \mathbb{N}$. We say f is good if it is non-decreasing, unbounded, and for all $n \in \mathbb{N}$, $f(n)$ can be computed in $\text{poly}(f(n))$ time. If there is a function f that is good and an algorithm that given a 2-CSP on k variables and alphabet set Σ where $|\Sigma| < f(k)$ that solves it in time at most $f(k-1)^{o(k/\log k)}$, then ETH is false.*

Original Proof in [Mar10]. At a high level, the proof of Marx [Mar10] can be described in the following way. It is known that a $2^{o(n)}$ algorithm for 3-coloring an n -vertex 4-regular graph would violate the ETH. Our goal is to use the hypothetical $f(k) \cdot n^{o(t/\log t)}$ algorithm for some graph class \mathcal{H} to efficiently solve the 3-coloring problem (given the reduction earlier, it will be convenient to consider 3-coloring as a 2-CSP on $|\Sigma| = 3$). Let us choose an $H \in \mathcal{H}$ with sufficiently large treewidth. The core of the result is a combinatorial embedding result showing that a graph with $O(n)$ edges can be efficiently embedded in graphs with large treewidth.

The notion of embedding we need is the following: it is a mapping f from the vertices of G to connected vertex sets of H such that if two vertices x and y of G are adjacent, then $f(x)$ and $f(y)$ intersect or have an edge between them (such an embedding is referred to as *connected embedding* in this paper). The depth of the embedding is the maximum number of times each vertex of H is used as an image of some vertex of G . The main combinatorial result of [Mar10] shows that a bounded-degree graph G with n vertices has an embedding into H of depth $O(n \log t / t)$, where t is the treewidth of H (and this embedding can be found efficiently). The embedding is constructed via the following steps: (1) large treewidth implies the existence of a large set without balanced separators², (2) a set without balanced separators implies the existence of a large uniform concurrent flow, (3) the paths appearing in this large concurrent flow can be used to embed a blowup of the line graph of a complete graph, (4) the bounded-degree graph G can be embedded into this blowup of the line graph of a clique by a simple routing scheme.

Given such an embedding, we can reduce a 2-CSP on $|\Sigma| = 3$ and having constraint graph G to a 2-CSP on constraint graph H , but having larger alphabet Σ' . If the embedding maps vertex v of G to a subset of

²Given a nonempty set W of vertices, we say that a set S of vertices is a balanced separator (with respect to W) if $|W \cap C| \leq |W|/2$ for every connected component C of $G \setminus S$.

vertices $f(v)$ of H , then every variable in $f(v)$ “simulates” variable v in the constructed instance. If the depth of the embedding is at most d , then each variable of the constructed instance needs to simulate at most d original variables. Therefore, if the alphabet size of the new instance is $|\Sigma'| = |\Sigma|^d = 3^d$, then each variable of H can simulate the required number of original variables. As $f(v)$ is connected in H , we can introduce constraints that ensure that every variable in $f(v)$ expresses the same value for v . Furthermore, when u and v are adjacent in G , then $f(u)$ and $f(v)$ intersect or adjacent in H , so we can introduce a constraint to ensure that the constraint C_{uv} on u and v are respected. Thus the created instance is equivalent to the original one. By carefully following how the instance parameters change during the reduction, we can argue that the hypothetical algorithm on H would solve the original 3-coloring instance in $2^{o(n)}$ time.

One technical difficulty in the proof is choosing $H \in \mathcal{H}$. To obtain $2^{o(n)}$ running time for 3-COLORING, the choice of H should depend on n : for larger n , we want to find an H with larger treewidth in order to have a larger advantage compared to brute force. Thus one has to take into account the time required to enumerate graphs from \mathcal{H} . On the other hand, graph H cannot be very large, otherwise the factor $f(k)$ in the running time could be too large. Thus the correct choice of H creates an additional layer of difficulties in the proof.

Our Proof. Our main observation is that if the goal is to obtain Theorem 1.2, then it can be reached in a much simpler way than going through a general bound for treewidth and restricting it to the special case of expanders. In particular, the required embedding result is much simpler if the target graph is an expander. The key step is provided by Leighton and Rao [LR99, Theorem 22], showing the following routing result: given disjoint pairs of vertices $(x_1, y_1), \dots, (x_t, y_t)$ in a k -vertex expander graph, we can find in polynomial time t paths P_1, \dots, P_t such that P_i has endpoints (x_i, y_i) and every edge is used by $O(\log k)$ of the paths.

Given an n -vertex bounded-degree graph G and a k -vertex bounded-degree expander H , we can use this routing result to find an embedding of G in H having depth $O(n \log k/k)$. First, let f be an arbitrary mapping from the vertices of G to vertices of H such that at most n/k vertices are mapped to each vertex of H . This mapping can be used to define a “demand multigraph” D : it has the same vertex set as H , and for every edge uv of G , we introduce a new edge $f(u)f(v)$ into D . As G has bounded degree and f maps at most n/k vertices to each vertex of H , it follows that D has maximum degree $O(n/k)$ and hence its edge set can be partitioned into $O(n/k)$ matchings. For each such matching, we can use the routing result of Leighton and Rao [LR99, Theorem 22] to find paths connecting the endpoints. This means that for every edge uv of G , we find a path in H connecting $f(u)$ and $f(v)$ in such a way that every edge of H is used by $O(n \log k/k)$ of the paths. As H has bounded degree, it also follows that every vertex is used by $O(n \log k/k)$ paths. Now we can define the embedding by letting $h(u)$ be the union of the paths corresponding to the edges of G incident to u . It is clear that every vertex of H is used by $O(n \log k/k)$ such sets and if uv is an edge of G then $h(u)$ and $h(v)$ intersect.

Another technical difference compared to the original proof of Marx [Mar10] is that the original proof started with the assumption that there is no $2^{o(n)}$ time algorithm for n -variable 3SAT. This is a somewhat weaker assumption than the ETH, which states that there is no $O(2^{\epsilon n})$ algorithm for some $\epsilon > 0$. Using the ETH as the starting assumption and stating the result for a fixed k as in Theorem 1.3 allows a proof with significantly fewer technicalities.

Obstacles to obtain tight result It remains an important open question whether the $\log k$ factor in Theorem 1.2 can be reduced or even eliminated. One obvious approach for improvement would be to improve the embedding result and show that an n -vertex bounded-degree graph G has an embedding with depth $o(n \log k/k)$ into some k -vertex expander. Note that the proof of our embedding result starts with an arbitrary balanced mapping of vertices of G to vertices of H , so there seems to be a lot of room for improvement for optimizing this scheme by a more careful grouping of vertices. However, by now, it is known that the $\log k$ factor in the embedding result cannot be improved in general.

THEOREM 1.5. (ALON AND MARX [AM11]) *Let H be a 3-regular graph with k vertices. Then, for all even $n > n_0(k)$, there exists a 3-regular graph G on n vertices so that any embedding into H is of depth at least $\Omega(\frac{n \log k}{k})$.*

Thus any improvement of Theorem 1.2 should involve significantly different techniques, going beyond the simple notion of embedding we have here.

1.1 Organization of Paper In Section 2 we introduce the problems and hypotheses of interest to this paper. In Section 3 we prove the main technical tool (a graph embedding theorem) needed to prove Theorems 1.3 and 1.4.

Finally, in Section 4, we prove Theorems 1.3 and 1.4.

2 Preliminaries

In this section, we state a few definitions of relevance to the rest of the paper.

3-SAT. In the 3-SAT problem, we are given a CNF formula φ over n variables x_1, \dots, x_n , such that each clause contains at most 3 literals. Our goal is to decide if there exist an assignment to x_1, \dots, x_n which satisfies φ .

HYPOTHESIS 2.1. (EXPONENTIAL TIME HYPOTHESIS (ETH) [IP01, IPZ01, Tov84]) *There exists $\epsilon > 0$ such that no algorithm can solve 3-SAT on n variables in time $O(2^{\epsilon n})$. Moreover, this holds even when restricted to formulae in which each variable appears in at most three clauses.*

Note that the original version of the hypothesis from [IP01] does not enforce the requirement that each variable appears in at most three clauses. To arrive at the above formulation, we first apply the Sparsification Lemma of [IPZ01], which implies that we can assume without loss of generality that the number of clauses m is $O(n)$. We then apply Tovey's reduction [Tov84] which produces a 3-CNF instance with at most $3m + n = O(n)$ variables and every variable occurs in at most three clauses. This means that the bounded occurrence restriction is also without loss of generality.

2-CSP. An instance Γ of 2-CSP consists of a (constraint) graph H , an alphabet set Σ , and, for each edge $\{u, v\} \in E(H)$, a constraint $C_{uv} \subseteq \Sigma \times \Sigma$. An *assignment* for Γ is a function $\sigma : V(H) \rightarrow \Sigma$. An edge $\{u, v\} \in E(H)$ is said to be *satisfied* by an assignment σ if and only if $(\sigma(u), \sigma(v)) \in C_{uv}$. The goal of the 2-CSP is to determine if there is an assignment σ that satisfies all the edges (i.e., constraints). Such an assignment is referred to as a satisfying assignment. A 2-CSP admitting a satisfying assignment is said to be satisfiable.

Cheeger Constant. Let G be a graph. For any subset $S \subseteq V(G)$, the edge boundary or cut of S , denoted $\delta(S)$, is the set of edges which have one endpoint in S and the other endpoint not in S , i.e.,

$$\forall S \subseteq V(G), \delta(S) := \{\{u, v\} \in E(G) \mid u \in S, v \notin S\}.$$

The Cheeger constant of G , denoted $\alpha(G)$, is defined by:

$$\alpha(G) := \min_{\substack{S \subseteq V(G) \\ |S| \leq |V(G)|/2}} \left\{ \frac{|\delta(S)|}{|S|} \right\}.$$

Expanders. We will now discuss about the construction of a bipartite 3-regular expander H on k vertices, keeping in mind that we will later provide \hat{G} and H as input to \mathcal{A}' . Our starting is the following result of Alon [Alo21].

THEOREM 2.1. (ALON [ALO21]) *There is an algorithm \mathcal{A} such that the following holds for every degree $d \geq 3$, every $\epsilon > 0$ and all sufficiently large $n \geq n_0(d, \epsilon)$, where $n \cdot d$ is even. \mathcal{A} on input n, d , and ϵ (n is provided in unary to \mathcal{A}) outputs in polynomial time a d -regular graph G on n vertices with at most one self-loop on each vertex³ such that the the absolute value of every non-trivial eigenvalue of the adjacency matrix of G is at most $2 \cdot \sqrt{d-1} + \epsilon$.*

Alon communicated to us [Alo23] that the above result can be used to prove the existence of 3-regular balanced bipartite expanders of every order (greater than 4).

THEOREM 2.2. (ALON [ALO23]) *There is a universal constant $\alpha_0 > 0$ and an algorithm \mathcal{A} such that the following holds. \mathcal{A} on input an even integer $n \geq 6$ (n is provided in unary to \mathcal{A}) outputs in polynomial time a 3-regular balanced bipartite graph G on n vertices such that $\alpha(G) \geq \alpha_0$.*

Proof. Let \mathcal{A}' be the algorithm given by Theorem 2.1 and let $n_0 \in \mathbb{N}$ be the integer guaranteed in that theorem statement for $d := 3$ and $\epsilon := 2.85 - 2\sqrt{2}$. The α_0 in this theorem statement is then set to $\min(1/n_0, 0.015)$.

³In [Alo21], the author uses the convention that a self-loop adds one to the degree of the vertex.

If the input n to \mathcal{A} is less than $\max(2n_0, 12)$ then simply output the following graph⁴ on n vertices: $\forall i \in [n/2]$, and $\forall t \in \{-1, 0, 1\}$, we have $(i, ((i+t) \bmod (n/2)) + \frac{n}{2} + 1) \in E(G_n)$. Note that G_n is connected, and thus for all non-empty $S \subseteq V(G_n)$ such that $|S| \leq n/2$, we have that there is an edge in $\delta(S)$. Thus, we have:

$$\alpha(G_n) \geq \left(\min_{\substack{S \subseteq V(G_n) \\ |S| \leq n/2}} \frac{1}{|S|} \right) \geq \frac{2}{n} \geq \frac{1}{n_0}.$$

If the input n to \mathcal{A} is at least $\max(2n_0, 12)$ and n is divisible by 4, then let G_n^0 be the output of \mathcal{A}' on input $n/2, d = 3, \varepsilon := 2.85 - 2\sqrt{2}$ (we can invoke \mathcal{A}' on input $n/2$ as $n/2$ is even). Let G_n be the double cover of G_n^0 , i.e., if A is the adjacency matrix of G_n^0 then the adjacency matrix of G_n is $\begin{bmatrix} \mathbf{0} & A \\ A & \mathbf{0} \end{bmatrix}$. Note that G_n^0 is a 3-regular graph on $n/2$ vertices and G_n is a 3-regular simple⁵ balanced bipartite graph on n vertices. Moreover, if $-2.85 \leq \lambda_{\frac{n}{2}}^0 \leq \lambda_{\frac{n}{2}-1}^0 \leq \dots \leq \lambda_2^0 \leq 2.85 < \lambda_1^0 = 3$ were the eigenvalues of G_n^0 then the set of eigenvalues of G_n is $\{(-1)^a \cdot \lambda_j^0 \mid j \in [n/2], a \in \{0, 1\}\}$. Thus the second highest eigenvalue of G_n is at most 2.85. From Cheeger inequality [Che70, Alo86, AM85], we have:

$$\alpha(G_n) \geq (3 - 2.85)/2 = 0.075.$$

If the input n to \mathcal{A} is at least $\max(2n_0, 12)$ and n is not divisible by 4 (i.e., $n \not\equiv 2 \pmod{4}$), then let G_{n+2} be the output of \mathcal{A} on input $n + 2$. Note that G_{n+2} is a 3-regular balanced bipartite graph on $n + 2$ vertices with $\alpha(G_{n+2}) \geq 0.075$. Let $e := \{u, v\}$ be an arbitrary edge in G_{n+2} . Let v, v_1 , and v_2 (resp. u, u_1 , and u_2) be the neighbors of u (resp. v). We remove u and v (and the edges incident to these two vertices), and insert an edge between u_1 and v_1 , and between u_2 and v_2 , to obtain G_n , a 3-regular balanced bipartite graph on n vertices.

Let $S \subseteq V(G_n)$ such that $|S| \leq |V(G_n)|/2$ and $\alpha(G_n) = \frac{|\delta(S)|}{|S|}$. Let $T := S \cap \{u_1, u_2, v_1, v_2\}$. Suppose $S = T$, then we note that $\delta(T) > 0$ in G_n , as otherwise, we have that the set of vertices $\{u, v, u_1, v_1, u_2, v_2\}$ is disconnected from the rest of the graph in G_{n+2} , and this implies $\alpha(G_{n+2}) = 0$ (as $n \geq 12$). But if $\delta(T) > 0$ then we have:

$$\alpha(G_n) = \frac{|\delta(S)|}{|S|} = \frac{|\delta(T)|}{|T|} \geq \frac{1}{4}.$$

On the other hand, suppose $S \neq T$, then let $\tilde{S} := S \setminus T$. First note that $\delta(S)$ in G_n is not empty as then we have that either $\delta(S)$ or $\delta(S \cup \{u, v\})$ is empty in G_{n+2} . Let $F_{\tilde{S}} \subseteq E(G_n)$ (resp. $F_T \subseteq E(G_n)$) be the set of edges with one endpoint in \tilde{S} (resp. T) and another endpoint in $V(G_n) \setminus S$. Since there are no edges between the set \tilde{S} and $\{u, v\}$ in G_{n+2} , we have that $F_{\tilde{S}}$ is also the set of edges with one endpoint in \tilde{S} and another endpoint in $V(G_{n+2}) \setminus S$. On the other hand, let $\tilde{F}_T \subseteq E(G_{n+2})$ be the set of edges with one endpoint in T and another endpoint in $V(G_{n+2}) \setminus S$, then we know that $0 \leq |\tilde{F}_T| - |F_T| \leq 4$. Also since $\delta(S)$ in G_n is not empty, we have that $|F_{\tilde{S}}| + |F_T| \geq 1$. Thus, we have:

$$|\tilde{F}_T| + |F_{\tilde{S}}| \leq 4 + |F_T| + |F_{\tilde{S}}| \leq 4 \cdot (|F_T| + |F_{\tilde{S}}|) + |F_T| + |F_{\tilde{S}}| = 5(|F_{\tilde{S}}| + |F_T|).$$

This implies:

$$\alpha(G_n) = \frac{|\delta_n(S)|}{|S|} \geq \frac{|F_{\tilde{S}}| + |F_T|}{|S|} \geq \frac{|F_{\tilde{S}}| + |\tilde{F}_T|}{5|S|} \geq \frac{\alpha(G_{n+2})}{5} \geq 0.015.$$

Finally, we note that \mathcal{A} runs in polynomial time because \mathcal{A}' runs in polynomial time. □

3 Proof of Main Embedding Theorem

In this section, we prove a graph embedding theorem which will be used to prove Theorems 1.3 and 1.4 in the next section. First we introduce a notion of graph embedding below.

⁴Any arbitrary connected balanced bipartite graph G_n on n vertices would suffice for us.

⁵Recall that in a simple graph there are no self-loops and no multiple edges between a pair of vertices.

Two connected subgraphs H_1 and H_2 of a graph H *touch* if they share a vertex or if there is an edge of H with one endpoint in $V(H_1)$ and another endpoint in $V(H_2)$. A *connected embedding* of a graph G in a graph H is a function $\Psi : V(G) \rightarrow \mathfrak{N}(V(H))$ that maps every $v \in V(G)$ to a nonempty connected subgraph in H such that for every edge $uv \in E(G)$, the subgraphs $\Psi(u)$ and $\Psi(v)$ touch. The *depth* of an embedding Ψ , denoted $\Delta(\Psi)$, is:

$$\Delta(\Psi) := \max_{x \in V(H)} |\{v \in V(G) \mid x \in \Psi(v)\}|,$$

that is, the maximum number of subgraphs $\Psi(v)$ that meet in a single vertex of H . In literature, such embedding also appear in the context of *congested minor models* where depth is called *congestion* or *ply*.

The main embedding result of this section is the following.

THEOREM 3.1. *There exists a constant $Z \geq 1$ and an algorithm \mathcal{A} that takes as input a graph G and an integer $k \geq 6$, and outputs a bipartite 3-regular graph H with no isolated vertices and a connected embedding $\Psi : V(G) \rightarrow \mathfrak{N}(V(H))$ such that the following holds.*

Size: $|V(H)| \leq k$.

Depth Guarantee: $\Delta(\Psi) \leq Z \cdot \left(1 + \frac{|V(G)| + |E(G)|}{k}\right) \cdot \log k$.

Runtime: \mathcal{A} runs in time $(|V(G)| + |E(G)|)^{O(1)}$.

Before, we proceed with the proof of the above theorem, we first state a key result of Leighton and Rao [LR99] on expander routing that will be used in our proof.

THEOREM 3.2. (LEIGHTON AND RAO [LR99]) *There exists an algorithm \mathcal{A} that takes as input a k -node bounded degree graph G , and a k -node bounded degree graph H with Cheeger constant $\alpha := \alpha(H)$, and a bijection $h : V(G) \rightarrow V(H)$, and outputs in deterministic polynomial time, a path P_{uv} for each edge uv of G such that the following holds.*

1. Path P_{uv} is a path between $h(u)$ and $h(v)$ in H ,
2. Each edge of H is in at most $O(\alpha^{-1} \log k)$ of the paths,
3. Path P_{uv} has length at most $O(\alpha^{-1} \log k)$.

Proof. [Proof of Theorem 3.1] Let \mathcal{A}' (resp. \mathcal{A}'') be the algorithm given by Theorem 3.2 (resp. Theorem 2.2). Given as input a d -regular graph G and an integer $k \geq 6$, the algorithm \mathcal{A} does the following.

\mathcal{A} first provides k as input to \mathcal{A}'' . Let H be the balanced bipartite 3-regular expander H on k vertices with Cheeger constant bounded away from α_0 (some positive universal constant) that is given as output.

\mathcal{A} then fixes an arbitrary function $g : V(G) \rightarrow [k]$ such that for all $j \in [k]$, we have $|\{v \in V(G) \mid g(v) = j\}| \leq \lceil |V(G)|/k \rceil$, i.e., g is balanced.

Consider the following multigraph (*demand graph*) D on k vertices, where for every $\{u, v\} \in E(G)$, if $g(u) \neq g(v)$, then there is a unique edge in $E(D)$ between $g(u)$ and $g(v)$. In D , there are at most $\frac{|V(G)|d}{k}$ edges incident on any vertex. Thus from [dW72] $E(D)$ admits a decomposition into matchings, where the size of each matching is at least $\frac{k}{2} - 1$. For each such matching, say $\{(x_1, y_1), \dots, (x_\ell, y_\ell)\}$, where $\ell \geq \frac{k}{2} - 1$, we construct the degree 1 graph \widehat{G} given by the union of the ℓ edges of the matching.

\mathcal{A}' on input \widehat{G} and H (and some canonical bijection between $V(\widehat{G})$ and $V(H)$) outputs ℓ paths in H whose congestion is at most $O(\log k)$ (since the Cheeger constant of H is bounded below by a positive constant). We do the above for each matching (a different \widehat{G} for each matching) and using the same graph H as input to \mathcal{A}' to obtain a collection of paths. More importantly, there is a 1-to-1 correspondence between the edges of G and this collection of paths.

The final connected embedding Ψ is given as follows: for each $v \in V(G)$, we have d many edges incident on it and each edge corresponds to a path of length at most $O(\log k)$. Then $\Psi(v)$ is the union of the nodes in all the d paths corresponding to the d edges incident of v .

The runtime claim is straightforward. To see the claim on the depth of the embedding, notice that the congestion on each node of H in each iteration of the calling of \mathcal{A}' is at most $O(\log k)$. We make $O(|V(G)|d/k)$ many calls and thus the total congestion is $O(\frac{|V(G)|d \log k}{k})$.

Finally, if $|V_1| \neq |V_2|$, the only difference, would be in the selection of g_1 and g_2 , where instead of the range being $[k/2]$ for both, it would instead need to be $k|V_1|/|V|$ and $k|V_2|/|V|$ respectively. \square

4 ETH Lower Bound for 2-CSP

In this section we first prove a bridging theorem to connect the embedding result from previous section with 2-CSPs, and then in the subsequent subsection prove Theorems 1.3 and 1.4.

4.1 Embedding 2-CSP to smaller constraint graph The below theorem is essentially proved in [Mar10] but we provide a proof here for the sake of completeness.

THEOREM 4.1. *There is an algorithm \mathcal{A} that takes as input a 2-CSP instance $\Gamma(G(V, E), \Sigma, \{C_{uv}\}_{\{u,v\} \in E})$, a graph $H(X, F)$ without isolated vertices, and a connected embedding $\Psi : V \rightarrow \blacktriangleleft(X)$ of depth d , and outputs a 2-CSP instance $\Phi(H(X, F), \Sigma^d, \{\tilde{C}_{xy}\}_{\{x,y\} \in F})$ such that the following holds.*

Reduction Guarantee: Γ is satisfiable if and only if Φ is satisfiable.

Runtime: \mathcal{A} runs in time $O(|E| \cdot |F| \cdot |\Sigma|^{2d+2})$.

Proof. Given a 2-CSP instance $\Gamma(G(V, E), \Sigma, \{C_{uv}\}_{\{u,v\} \in E})$, a graph $H(X, F)$, and a connected embedding $\Psi : V \rightarrow \blacktriangleleft(X)$ of depth d , the algorithm \mathcal{A} outputs the following 2-CSP instance $\Phi(H(X, F), \Sigma^d, \{\tilde{C}_{xy}\}_{\{x,y\} \in F})$, where we only need to specify $\{\tilde{C}_{xy}\}_{\{x,y\} \in F}$.

For every $x \in X$ define $V_x \subseteq V$ as follows: for all $v \in V$ we have $v \in V_x$ if and only if $x \in \Psi(v)$. We further define E_x to be all the edges in E whose both end points are in V_x . Finally, define $d_x := |V_x|$ and $\pi_x : V_x \rightarrow [d_x]$ to be some canonical bijection. Note that since Ψ has depth d we know that $d_x \leq d$ for all $x \in X$.

Fix some $\{x, y\} \in F$. Let $V_{xy} \subseteq V$ and $E_{xy} \subseteq E$ be defined as follows.

$$V_{xy} := V_x \cap V_y, \quad E_{xy} := \{\{u, v\} \in E \mid u \in V_x, v \in V_y\}.$$

Let $(a, b) \in \Sigma^d \times \Sigma^d$. We include $(a, b) \in \tilde{C}_{xy}$ if and only if both the following hold:

Vertex Consistency: For every $v \in V_{xy}$ we have $a_{\pi_x(v)} = b_{\pi_y(v)}$.

Edge Consistency: For every $\{u, v\} \in E_{xy}$ (resp. $\{u, u'\} \in E_x$ or $\{v, v'\} \in E_y$) we have $(a_{\pi_x(u)}, b_{\pi_y(v)}) \in C_{uv}$ (resp. $(a_{\pi_x(u)}, a_{\pi_x(u')}) \in C_{uu'}$ or $(b_{\pi_y(v)}, b_{\pi_y(v')}) \in C_{vv'}$).

Completeness Analysis: Suppose $\sigma : V \rightarrow \Sigma$ is a satisfying assignment to Γ then consider the assignment $\tilde{\sigma} : X \rightarrow \Sigma^d$ to Φ constructed from σ as follows. Let $\alpha \in \Sigma$ be an arbitrary element. For all $x \in X$ we have:

$$\tilde{\sigma}(x)_i := \begin{cases} \sigma(\pi_x^{-1}(i)) & \text{if } i \leq d_x \\ \alpha & \text{if } i > d_x \end{cases}.$$

We claim that $\tilde{\sigma}$ is a satisfying assignment to Φ . To see this, consider an arbitrary $\{x, y\} \in F$. We verify that $(\tilde{\sigma}(x), \tilde{\sigma}(y)) \in \tilde{C}_{xy}$ by checking the conditions of vertex consistency and edge consistency are met.

For every $v \in V_{xy}$ we have by definition that $\tilde{\sigma}(x)_{\pi_x(v)} = \sigma(\pi_x^{-1}(\pi_x(v))) = \sigma(v) = \sigma(\pi_y^{-1}(\pi_y(v))) = \tilde{\sigma}(y)_{\pi_y(v)}$. This verifies vertex consistency. Next we verify edge consistency.

For every $\{u, v\} \in E_{xy}$ (resp. $\{u, u'\} \in E_x$ or $\{v, v'\} \in E_y$) we have $(\tilde{\sigma}(x)_{\pi_x(u)}, \tilde{\sigma}(y)_{\pi_y(v)}) = (\sigma(u), \sigma(v))$ (resp. $(\tilde{\sigma}(x)_{\pi_x(u)}, \tilde{\sigma}(x)_{\pi_x(u')}) = (\sigma(u), \sigma(u'))$ or $(\tilde{\sigma}(y)_{\pi_y(v)}, \tilde{\sigma}(y)_{\pi_y(v')}) = (\sigma(v), \sigma(v'))$) and this is in C_{uv} (resp. $C_{uu'}$ or $C_{vv'}$) because σ is a satisfying assignment to Γ .

Soundness Analysis: Suppose $\tilde{\sigma} : X \rightarrow \Sigma^d$ is a satisfying assignment to Φ then consider the assignment $\sigma : V \rightarrow \Sigma$ to Γ constructed from $\tilde{\sigma}$ as follows.

Fix any function $\tau : V \rightarrow X$ that satisfies the following: for all $v \in V$, $\tau(v) \in \Psi(v)$. Informally, τ is picking a representative vertex in $\Psi(v)$ for each $v \in V$. For all $v \in V$ we have:

$$\sigma(v) := \tilde{\sigma}(\tau(v))_{\pi_{\tau(v)}(v)}.$$

We first observe that the value $\sigma(v)$ is independent of the choice of $\tau(v)$, that is,

$$(4.1) \quad \forall v \in V, \forall x \in \Psi(v), \quad \tilde{\sigma}(x)_{\pi_x(v)} = \sigma(v).$$

To see (4.1), fix $v \in V$ and $x \in \Psi(v)$. Since $\Psi(v)$ is a connected subgraph of H , we have that there is a path from $\tau(v)$ to x contained in $\Psi(v)$, and let $\tau(v) =: x_1 - x_2 - \dots - x_r := x$ be one such path where for all $i \in [r]$ we have $x_i \in \Psi(v)$. Since $\tilde{\sigma}$ is a satisfying assignment, for every $i \in [r-1]$, by looking at the vertex consistency condition, we have $\tilde{\sigma}(x_i)_{\pi_{x_i}(v)} = \tilde{\sigma}(x_{i+1})_{\pi_{x_{i+1}}(v)}$. Thus, we have $\sigma(v) = \tilde{\sigma}(\tau(v))_{\pi_{\tau(v)}(v)} = \tilde{\sigma}(x_1)_{\pi_{x_1}(v)} = \tilde{\sigma}(x_r)_{\pi_{x_r}(v)} = \tilde{\sigma}(x)_{\pi_x(v)}$. This proves (4.1).

We claim that σ is a satisfying assignment to Γ . To see this, consider an arbitrary $\{u, v\} \in E$. We verify below that $(\sigma(u), \sigma(v)) \in C_{uv}$.

From the definition of Ψ , we have that either (i) there exists $x \in X$ such that $x \in \Psi(u) \cap \Psi(v)$, or, (ii) there exists $\{x, y\} \in F$ such that $x \in \Psi(u)$ and $y \in \Psi(v)$.

Case (i): By (4.1), we have $\sigma(u) = \tilde{\sigma}(x)_{\pi_x(u)}$ and $\sigma(v) = \tilde{\sigma}(x)_{\pi_x(v)}$. Since $\{u, v\} \in E_x$ and $H(X, F)$ contains at least one edge incident with x , by the edge consistency condition on any such edge we have $(\sigma(u) = \tilde{\sigma}(x)_{\pi_x(u)}, \sigma(v) = \tilde{\sigma}(x)_{\pi_x(v)}) \in C_{uv}$, as desired.

Case (ii): By (4.1), we have $\sigma(u) = \tilde{\sigma}(x)_{\pi_x(u)}$ and $\sigma(v) = \tilde{\sigma}(y)_{\pi_y(v)}$. Since $\{u, v\} \in E_{xy}$ and noting the edge consistency condition we have $(\sigma(u) = \tilde{\sigma}(x)_{\pi_x(u)}, \sigma(v) = \tilde{\sigma}(y)_{\pi_y(v)}) \in C_{uv}$, as desired.

Runtime Analysis: Finally, for a fixing of $\{x, y\}$ in F and a fixing of $(a, b) \in \Sigma^d \times \Sigma^d$, the time needed to check if $(a, b) \in \tilde{C}_{xy}$ is at most $O(d^2 \cdot |\Sigma|^2)$. Thus, the running time to produce Φ given Ψ and Γ is $O(|E| \cdot |\Sigma|^2 + |F| \cdot |\Sigma|^{2d} \cdot d^2 \cdot |\Sigma|^2) = O(|E| \cdot |F| \cdot |\Sigma|^{2d+2})$. \square

An immediate corollary that will be useful to us is the following.

COROLLARY 4.1. (PUTTING TOGETHER THEOREM 3.1 AND THEOREM 4.1) *There exists a constant $Z > 1$ and an algorithm \mathcal{A} that takes as input a 2-CSP instance $\Gamma(G(V, E), \Sigma, \{C_{uv}\}_{\{u,v\} \in E})$ and an integer $k \geq 6$, and outputs a 2-CSP instance $\Phi(H(X, F), \tilde{\Sigma}, \{\tilde{C}_{xy}\}_{\{x,y\} \in F})$ such that the following holds.*

Size: H is a 3-regular bipartite graph with $|X| \leq k$. We have $|\tilde{\Sigma}| = |\Sigma|^{Z \cdot (1+k^{-1}(|V|+|E|) \cdot \log k)}$.

Reduction Guarantee: Γ is satisfiable if and only if Φ is satisfiable.

Runtime: There is a polynomial $p : \mathbb{N} \rightarrow \mathbb{N}$ such that \mathcal{A} runs in time bounded above by $p\left(|E| \cdot k \cdot |\Sigma|^{\left(\frac{(|V|+|E|) \cdot \log k}{k} + 1\right)}\right)$.

Proof. Given as input a 2-CSP instance $\Gamma(G(V, E), \Sigma, \{C_{uv}\}_{\{u,v\} \in E})$ and an integer k the algorithm \mathcal{A} of the corollary statement does the following. Let \mathcal{A}_0 and \mathcal{A}_1 be the algorithms guaranteed through Theorem 3.1 and Theorem 4.1 respectively.

\mathcal{A} first runs \mathcal{A}_0 on input $(G(V, E), k)$ and obtains as output a bipartite 3-regular graph $H(X, F)$ with no isolated vertices and a connected embedding $\Psi : V \rightarrow \mathfrak{N}(X)$ such that $|X| \leq k$ and $\Delta(\Psi) \leq Z \cdot (1 + k^{-1}(|V| + |E|)) \cdot \log k$.

\mathcal{A} then runs \mathcal{A}_1 on input $(\Gamma(G(V, E), \Sigma, \{C_{uv}\}_{\{u,v\} \in E}), H(X, F), \Psi : V \rightarrow \mathfrak{N}(X))$, to obtain a 2-CSP instance $\Phi(H(X, F), \tilde{\Sigma}, \{\tilde{C}_{xy}\}_{\{x,y\} \in F})$ where $|\tilde{\Sigma}|$ is some integer bounded above by $|\Sigma|^{Z \cdot (1+k^{-1}(|V|+|E|) \cdot \log k)}$.

Finally, \mathcal{A} outputs Φ . The reduction guarantee of the corollary statement follows from the reduction guarantee of Theorem 4.1.

The runtime of the algorithm is upper bounded by the sum of the runtimes of \mathcal{A}_0 and \mathcal{A}_1 . The output of \mathcal{A}_0 is produced in time $(|V| + |E|)^{O(1)}$. The output of \mathcal{A}_1 is produced in time $O(|E| \cdot |F| \cdot |\Sigma|^{2 \cdot (1+Z \cdot (1+k^{-1}(|V|+|E|) \cdot \log k))})$. The claim on the runtime in the corollary statement follows by noting that $|F| \leq 1.5 \cdot |X| \leq 1.5 \cdot k$. \square

4.2 ETH Lower Bounds Our starting point for the lower bounds in this subsection is the following result on 3-coloring:

THEOREM 4.2. (ETH LOWER BOUND FOR 3-COLORING [CFG⁺16]) *Assuming ETH, there is some $\delta > 0$ such that no algorithm can decide all 2-CSP instances $\Gamma(G(V, E), \Sigma, \{C_{uv}\}_{\{u,v\} \in E})$, where $|\Sigma| = 3$ and G is 4-regular, while running in $2^{\delta|V|}$ time.*

We remark that in [CFG⁺16] the lower bound is proven for 2-CSP instances with max-degree 4, but this can be easily extended to 4-regular graphs by adding dummy constraints⁶ between pairs of vertices which are both of degree less than 4.

THEOREM 4.3. *There exists $\alpha > 0$ and $k_0 \in \mathbb{N}$ such that the following holds. If there exists some fixed integer $k \geq k_0$ and an algorithm \mathcal{A}_k with the following guarantees, then ETH is false.*

Input: \mathcal{A}_k takes as input a 2-CSP instance $\Gamma(H, \Sigma, \{C_{uv}\}_{\{u,v\} \in E(H)})$ where H is a 3-regular simple bipartite graph on k vertices.

Output: \mathcal{A}_k outputs 1 if and only if there is a satisfying assignment to Γ .

Run time: \mathcal{A}_k runs in time $O(|\Sigma|^{\alpha k / \log k})$.

Proof. Suppose there is some $k \geq k_0$ and an algorithm \mathcal{A}_k as suggested in this theorem statement for $\alpha := \frac{\delta}{4Z}$ and $k_0 := Z^2/\delta^2$, where δ is the constant referred to in Theorem 4.2 and Z is the constant referred to in Corollary 4.1. We will then design an algorithm $\tilde{\mathcal{A}}$ that runs in $2^{\delta n/2}$ time and can decide 2-CSPs on n variables/vertices even when the alphabet set is of size 3 and the constraint graph is 4-regular. This would imply that ETH is false (from Theorem 4.2).

Given $\tilde{\mathcal{A}}$ as input a 2-CSP instance $\Gamma(G(V, E), \Sigma, \{C_{uv}\}_{\{u,v\} \in E})$ where $|\Sigma| = 3$ and 4-regular graph G , $\tilde{\mathcal{A}}$ first runs the algorithm in Corollary 4.1 (Γ, k) as input and obtains as output a 2-CSP instance $\Phi(H(X, F), \tilde{\Sigma}, \{\tilde{C}_{xy}\}_{\{x,y\} \in F})$ with the guarantee that $|X| \leq k$ and $|\tilde{\Sigma}| \leq 3^{Z \cdot (1 + \frac{|V| \log k}{k})}$, for some universal constant $Z \geq 1$ (also note that the size of Φ is at most $O(k \cdot |\tilde{\Sigma}|) = O\left(k \cdot 3^{Z \cdot (1 + \frac{|V| \log k}{k})}\right) = O\left(k \cdot 3^{Z \cdot (1 + \frac{|V| \log k_0}{k_0})}\right) < 2^{\delta|V|}$).

Then $\tilde{\mathcal{A}}$ feeds Φ as input to \mathcal{A} and obtains as output $\lambda \in \{0, 1\}$. $\tilde{\mathcal{A}}$ outputs λ . The reduction guarantee of the theorem statement follows from the reduction guarantee of Corollary 4.1.

The runtime of the algorithm is upper bounded by the sum of the runtimes of \mathcal{A}_k and the algorithm in Corollary 4.1. The output of the algorithm in Corollary 4.1 is produced in time $p(|V| \cdot k \cdot 3)$ for some polynomial p . The output of \mathcal{A} is produced in time $|\tilde{\Sigma}|^{\alpha \cdot k / \log k} \leq 3^{Z \cdot (1 + \frac{|V| \log k}{k}) \cdot \alpha \cdot \frac{k}{\log k}} = 3^{\alpha \cdot Z \cdot \frac{k}{\log k} + \alpha \cdot Z \cdot |V|} \leq 3^{2\alpha \cdot Z \cdot |V|} \leq 2^{4\alpha \cdot Z \cdot |V|} = 2^{\delta \cdot |V|}$. \square

For most applications, the following corollary is sufficient.

COROLLARY 4.2. *Assuming ETH, there is no $f(k) \cdot |\Sigma|^{o(k/\log k)}$ time algorithm that can decide all 2-CSP instances on 3-regular bipartite constraint graphs, where Σ is the alphabet set, k is the number of constraints, and f is an arbitrary function.*

We now turn to prove Theorem 1.4 .

THEOREM 4.4. *Let $f : \mathbb{N} \rightarrow \mathbb{N}$. We say f is good if it is non-decreasing, unbounded, and for all $n \in \mathbb{N}$, $f(n)$ can be computed in $\text{poly}(f(n))$ time. If there is a function f that is good and an algorithm that given a 2-CSP on k variables and alphabet set Σ where $|\Sigma| < f(k)$ that solves it in time at most $f(k-1)^{o(k/\log k)}$, then ETH is false.*

⁶A dummy constraint on two variables accepts every possible assignment to the two variables.

Proof. Suppose there is a function f that is good and an algorithm \mathcal{A}_f as suggested in this theorem statement for $\alpha := \frac{\delta}{4Z}$ where δ is the constant referred to in Theorem 4.2 and Z is the constant referred to in Corollary 4.1. We will then design an algorithm $\tilde{\mathcal{A}}$ that runs in $2^{\delta n/2}$ time and can decide 2-CSPs on n variables/vertices even when the alphabet set is of size 3 and the constraint graph is 4-regular. This would imply that ETH is false (from Theorem 4.2).

Given $\tilde{\mathcal{A}}$ as input a 2-CSP instance $\Gamma(G(V, E), \Sigma, \{C_{uv}\}_{\{u,v\} \in E})$ where $|\Sigma| = 3$ and 4-regular graph G , $\tilde{\mathcal{A}}$ first computes an appropriate integer k which it will later feed as input to the algorithm in Corollary 4.1. Let k_0 be the smallest integer such that $f(k_0) \geq 2$. Let k be the largest integer such that $f(k-1) < 3^{Z \cdot (1 + \frac{|V| \log k}{k})}$. Such an integer k can be computed in time $|V|^{O(1)}$ in the following way: We evaluate $f(1), f(2), \dots$, and stop when either $f(k) \geq 3^{Z \cdot (1 + \frac{|V| \log k}{k})}$ or if computing $f(k)$ does not terminate after $|V|^2 + k_0$ steps (in which case we can conclude that $f(k) > 3^{Z \cdot (1 + \frac{|V| \log k}{k})}$).

Now $\tilde{\mathcal{A}}$ first runs the algorithm in Corollary 4.1 (Γ, k) as input and obtains as output a 2-CSP instance $\Phi(H(X, F), \tilde{\Sigma}, \{\tilde{C}_{xy}\}_{\{x,y\} \in F})$ with the guarantee that $|X| \leq k$, $H(X, F)$ is 3-regular bipartite, and $|\tilde{\Sigma}| \leq 3^{Z \cdot (1 + \frac{|V| \log k}{k})} \leq f(k)$, for some universal constant $Z \geq 1$ (also note that the size of Φ is at most $O(k \cdot |\tilde{\Sigma}|) = O(k \cdot 3^{Z \cdot (1 + \frac{|V| \log k}{k})}) < 2^{\delta |V|}$).

Then $\tilde{\mathcal{A}}$ feeds Φ as input to \mathcal{A} and obtains as output $\lambda \in \{0, 1\}$. $\tilde{\mathcal{A}}$ outputs λ . The reduction guarantee of the theorem statement follows from the reduction guarantee of Corollary 4.1.

The runtime of the algorithm is upper bounded by the sum of the runtimes of \mathcal{A}_f and the algorithm in Corollary 4.1. The output of the algorithm in Corollary 4.1 is produced in time $p(|V| \cdot k \cdot 3)$ for some polynomial p . The output of \mathcal{A} is produced in time $|\tilde{\Sigma}|^{\alpha \cdot k / \log k} \leq 3^{Z \cdot (1 + \frac{|V| \log k}{k}) \cdot \alpha \cdot \frac{k}{\log k}} = 3^{\alpha \cdot Z \cdot \frac{k}{\log k} + \alpha \cdot Z \cdot |V|} \leq 3^{2\alpha \cdot Z \cdot |V|} \leq 2^{4\alpha \cdot Z \cdot |V|} = 2^{\delta \cdot |V|}$. \square

The parameter setting of Theorem 1.4 may seem peculiar at first glance, but it has applications in lower bounds for classic (not parameterized) problems under ETH. For an illustration, consider a recent work on GLOBAL LABEL MINCUT [JLM⁺23]. The actual details of the considered problem are not important here; it suffices to say that the main technical result of [JLM⁺23] is a parameterized reduction from 2-CSP to one parameterization of GLOBAL LABEL MINCUT that, given a 2-CSP instance $\Gamma(G(V, E), \Sigma, \{C_{uv}\}_{\{u,v\} \in E})$ with $a := |V| + |E|$ produces an equivalent instance of GLOBAL LABEL MINCUT of size $2^{O(a \log a)} \cdot |\Sigma|$ with parameter value at most a . This reduction, combined with Theorem 1.4, refutes (assuming ETH) an existence of an algorithm for GLOBAL LABEL MINCUT with running time $n^{o(\log n / (\log \log n)^2)}$, where n is the size of the instance, providing a nearly-tight lower bound to an existing quasipolynomial-time algorithm [GKP17].

Indeed, consider a good function $f(n) = n^n$. Let $\Gamma(G(V, E), \Sigma, \{C_{uv}\}_{\{u,v\} \in E})$ be a 2-CSP instance on k variables, $G(V, E)$ 3-regular, and alphabet set Σ where $|\Sigma| < f(k)$. The aforementioned reduction, applied to this 2-CSP instance, produces a GLOBAL LABEL MINCUT instance of size bounded by

$$2^{O(k \log k)} \cdot k^k = 2^{O(k \log k)}.$$

Now, a hypothetical algorithm solving this instance in time $n^{o(\log n / (\log \log n)^2)}$ solves the original 2-CSP instance in time $2^{o(k^2)}$, which is asymptotically smaller than $f(k-1)^{\alpha k / \log k}$ for a constant $\alpha > 0$.

Acknowledgement Karthik C. S. is supported by the National Science Foundation under Grant CCF-2313372 and by the Simons Foundation, Grant Number 825876, Awardee Thu D. Nguyen. Dániel Marx is supported by the European Research Council (ERC) consolidator grant No. 725978 SYSTEMATICGRAPH. During this research Marcin Pilipczuk was part of BARC, supported by the VILLUM Foundation grant 16582. We are also grateful to the Dagstuhl Seminar 23291 for a special collaboration opportunity.

References

- [AKMR19] Saeed Akhoondian Amiri, Stephan Kreutzer, Dániel Marx, and Roman Rabinovich. Routing with congestion in acyclic digraphs. *Inf. Process. Lett.*, 151, 2019.
- [Al086] Noga Alon. Eigenvalues and expanders. *Combinatorica*, 6(2):83–96, 1986.

- [Alo21] Noga Alon. Explicit expanders of every degree and size. *Combinatorica*, pages 1–17, 2021.
- [Alo23] Noga Alon. personal communication, 2023.
- [AM85] Noga Alon and Vitali D Milman. λ_1 , isoperimetric inequalities for graphs, and superconcentrators. *Journal of Combinatorial Theory, Series B*, 38(1):73–88, 1985.
- [AM11] Noga Alon and Dániel Marx. Sparse balanced partitions and the complexity of subgraph problems. *SIAM J. Discret. Math.*, 25(2):631–644, 2011.
- [APSZ21] Akanksha Agrawal, Fahad Panolan, Saket Saurabh, and Meirav Zehavi. Simultaneous feedback edge set: A parameterized perspective. *Algorithmica*, 83(2):753–774, 2021.
- [BCMP20] Édouard Bonnet, Sergio Cabello, Bojan Mohar, and Hebert Pérez-Rosés. The inverse voronoi problem in graphs I: hardness. *Algorithmica*, 82(10):3018–3040, 2020.
- [BGL17] Édouard Bonnet, Panos Giannopoulos, and Michael Lampis. On the parameterized complexity of red-blue points separation. In Daniel Lokshantov and Naomi Nishimura, editors, *12th International Symposium on Parameterized and Exact Computation, IPEC 2017, September 6-8, 2017, Vienna, Austria*, volume 89 of *LIPICs*, pages 8:1–8:13. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2017.
- [BIJK19] Édouard Bonnet, Yoichi Iwata, Bart M. P. Jansen, and Lukasz Kowalik. Fine-grained complexity of k-opt in bounded-degree graphs for solving TSP. In Michael A. Bender, Ola Svensson, and Grzegorz Herman, editors, *27th Annual European Symposium on Algorithms, ESA 2019, September 9-11, 2019, Munich/Garching, Germany*, volume 144 of *LIPICs*, pages 23:1–23:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.
- [BKMN16] Karl Bringmann, László Kozma, Shay Moran, and N. S. Narayanaswamy. Hitting set for hypergraphs of low vc-dimension. In Piotr Sankowski and Christos D. Zaroliagis, editors, *24th Annual European Symposium on Algorithms, ESA 2016, August 22-24, 2016, Aarhus, Denmark*, volume 57 of *LIPICs*, pages 23:1–23:18. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2016.
- [BM16] Édouard Bonnet and Tillmann Miltzow. Parameterized hardness of art gallery problems. In Piotr Sankowski and Christos D. Zaroliagis, editors, *24th Annual European Symposium on Algorithms, ESA 2016, August 22-24, 2016, Aarhus, Denmark*, volume 57 of *LIPICs*, pages 19:1–19:17. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2016.
- [BS17] Édouard Bonnet and Florian Sikora. The graph motif problem parameterized by the structure of the input graph. *Discrete Applied Mathematics*, 231:78–94, 2017.
- [CCG⁺13] Jason Crampton, Robert Crowston, Gregory Z. Gutin, Mark Jones, and M. S. Ramanujan. Fixed-parameter tractability of workflow satisfiability in the presence of seniority constraints. In Michael R. Fellows, Xuehou Tan, and Binhai Zhu, editors, *Frontiers in Algorithmics and Algorithmic Aspects in Information and Management, Third Joint International Conference, FAW-AAIM 2013, Dalian, China, June 26-28, 2013. Proceedings*, volume 7924 of *Lecture Notes in Computer Science*, pages 198–209. Springer, 2013.
- [CDH21] Radu Curticapean, Holger Dell, and Thore Husfeldt. Modular counting of subgraphs: Matchings, matching-splittable graphs, and paths. In Petra Mutzel, Rasmus Pagh, and Grzegorz Herman, editors, *29th Annual European Symposium on Algorithms, ESA 2021, September 6-8, 2021, Lisbon, Portugal (Virtual Conference)*, volume 204 of *LIPICs*, pages 34:1–34:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.
- [CDM17] Radu Curticapean, Holger Dell, and Dániel Marx. Homomorphisms are a good basis for counting small subgraphs. In Hamed Hatami, Pierre McKenzie, and Valerie King, editors, *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 210–223. ACM, 2017.
- [CdVMdM21] Vincent Cohen-Addad, Éric Colin de Verdière, Dániel Marx, and Arnaud de Mesmay. Almost tight lower bounds for hard cutting problems in embedded graphs. *J. ACM*, 68(4):30:1–30:26, 2021.
- [CFG⁺16] Marek Cygan, Fedor V Fomin, Alexander Golovnev, Alexander S Kulikov, Ivan Mihajlin, Jakub Pachocki, and Arkadiusz Socała. Tight bounds for graph homomorphism and subgraph isomorphism. In *Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms*, pages 1643–1649. SIAM, 2016.
- [CFHM20] Rajesh Hemant Chitnis, Andreas Emil Feldmann, Mohammad Taghi Hajiaghayi, and Dániel Marx. Tight bounds for planar strongly connected steiner subgraph with fixed number of terminals (and extensions). *SIAM J. Comput.*, 49(2):318–364, 2020.
- [CFM21] Rajesh Chitnis, Andreas Emil Feldmann, and Pasin Manurangsi. Parameterized approximation algorithms for bidirected steiner network problems. *ACM Trans. Algorithms*, 17(2):12:1–12:68, 2021.
- [Che70] Jeff Cheeger. A lower bound for the smallest eigenvalue of the laplacian, problems in analysis (papers dedicated to salomon bochner, 1969), 1970.
- [CHKX06a] Jianer Chen, Xiuzhen Huang, Iyad A. Kanj, and Ge Xia. On the computational hardness based on linear fpt-reductions. *J. Comb. Optim.*, 11(2):231–247, 2006.
- [CHKX06b] Jianer Chen, Xiuzhen Huang, Iyad A. Kanj, and Ge Xia. Strong computational lower bounds via parameterized complexity. *J. Comput. Syst. Sci.*, 72(8):1346–1367, 2006.
- [CX15] Radu Curticapean and Mingji Xia. Parameterizing the permanent: Genus, apices, minors, evaluation mod 2k. In Venkatesan Guruswami, editor, *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015*,

- Berkeley, CA, USA, 17-20 October, 2015, pages 994–1009. IEEE Computer Society, 2015.
- [DEG22] Argyrios Deligkas, Eduard Eiben, and Tiger-Lily Goldsmith. Parameterized complexity of hotelling-downs with party nominees. In Luc De Raedt, editor, *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022, Vienna, Austria, 23-29 July 2022*, pages 244–250. ijcai.org, 2022.
- [dW72] Dominique de Werra. Decomposition of bipartite multigraphs into matchings. *Zeitschrift für Operations Research*, 16:85–90, 1972.
- [EGN⁺23] Eduard Eiben, Gregory Z. Gutin, Philip R. Neary, Clément Rambaud, Magnus Wahlström, and Anders Yeo. Preference swaps for the stable matching problem. *Theor. Comput. Sci.*, 940(Part):222–230, 2023.
- [EKPS18] Eduard Eiben, Dusan Knop, Fahad Panolan, and Ondrej Suchý. Complexity of the steiner network problem with respect to the number of terminals. *CoRR*, abs/1802.08189, 2018.
- [EL18] David Eppstein and Daniel Lokshtanov. The parameterized complexity of finding point sets with hereditary properties. *CoRR*, abs/1808.02162, 2018.
- [FHRV09] Michael R. Fellows, Danny Hermelin, Frances A. Rosamond, and Stéphane Vialette. On the parameterized complexity of multiple-interval graph problems. *Theor. Comput. Sci.*, 410(1):53–61, 2009.
- [FPRS23] Fedor V. Fomin, Fahad Panolan, M. S. Ramanujan, and Saket Saurabh. On the optimality of pseudo-polynomial algorithms for integer programming. *Math. Program.*, 198(1):561–593, 2023.
- [GHNS13] Jiong Guo, Sepp Hartung, Rolf Niedermeier, and Ondrej Suchý. The parameterized complexity of local search for tsp, more refined. *Algorithmica*, 67(1):89–110, 2013.
- [GKP17] Mohsen Ghaffari, David R. Karger, and Debmalya Panigrahi. Random contractions and sampling for hypergraph and hedge connectivity. In Philip N. Klein, editor, *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16-19*, pages 1101–1114. SIAM, 2017.
- [IP01] Russell Impagliazzo and Ramamohan Paturi. On the complexity of k-sat. *J. Comput. Syst. Sci.*, 62(2):367–375, 2001.
- [IPZ01] Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. Which problems have strongly exponential complexity? *J. Comput. Syst. Sci.*, 63(4):512–530, 2001.
- [JKMS13] Klaus Jansen, Stefan Kratsch, Dániel Marx, and Ildikó Schlotter. Bin packing with fixed number of bins revisited. *J. Comput. Syst. Sci.*, 79(1):39–49, 2013.
- [JLM⁺23] Lars Jaffke, Paloma T. Lima, Tomáš Masarík, Marcin Pilipczuk, and Uéverton S. Souza. A tight quasi-polynomial bound for global label min-cut. In Nikhil Bansal and Viswanath Nagarajan, editors, *Proceedings of the 2023 ACM-SIAM Symposium on Discrete Algorithms, SODA 2023, Florence, Italy, January 22-25, 2023*, pages 290–303. SIAM, 2023.
- [JLR⁺17] Mark Jones, Daniel Lokshtanov, M. S. Ramanujan, Saket Saurabh, and Ondrej Suchý. Parameterized complexity of directed steiner tree on sparse graphs. *SIAM J. Discrete Math.*, 31(2):1294–1327, 2017.
- [KSS22] Dusan Knop, Simon Schierreich, and Ondrej Suchý. Balancing the spread of two opinions in sparse social networks (student abstract). In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022, Thirty-Fourth Conference on Innovative Applications of Artificial Intelligence, IAAI 2022, The Twelfth Symposium on Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event, February 22 - March 1, 2022*, pages 12987–12988. AAAI Press, 2022.
- [LR99] Frank Thomson Leighton and Satish Rao. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *J. ACM*, 46(6):787–832, 1999.
- [LRSZ20] Daniel Lokshtanov, M. S. Ramanujan, Saket Saurabh, and Meirav Zehavi. Parameterized complexity and approximability of directed odd cycle transversal. In Shuchi Chawla, editor, *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5-8, 2020*, pages 2181–2200. SIAM, 2020.
- [Mar10] Dániel Marx. Can you beat treewidth? *Theory of Computing*, 6(1):85–112, 2010.
- [MP15] Dániel Marx and Marcin Pilipczuk. Optimal parameterized algorithms for planar facility location problems using voronoi diagrams. In Nikhil Bansal and Irene Finocchi, editors, *Algorithms - ESA 2015 - 23rd Annual European Symposium, Patras, Greece, September 14-16, 2015, Proceedings*, volume 9294 of *Lecture Notes in Computer Science*, pages 865–877. Springer, 2015.
- [NS22] Jesper Nederlof and Céline M. F. Swennenhuis. On the fine-grained parameterized complexity of partial scheduling to minimize the makespan. *Algorithmica*, 84(8):2309–2334, 2022.
- [PW18] Marcin Pilipczuk and Magnus Wahlström. Directed multicut is w[1]-hard, even for four terminal pairs. *TOCT*, 10(3):13:1–13:18, 2018.
- [Tov84] Craig A. Tovey. A simplified np-complete satisfiability problem. *Discret. Appl. Math.*, 8(1):85–89, 1984.