Optimal Adjacency Labels for Subgraphs of Cartesian Products

Louis Esperet
Laboratoire G-SCOP, Grenoble, France

Nathaniel Harms
EPFL, Lausanne, Switzerland

Viktor Zamaraev
University of Liverpool, UK

Abstract

For any hereditary graph class $\mathcal{F}$, we construct optimal adjacency labeling schemes for the classes of subgraphs and induced subgraphs of Cartesian products of graphs in $\mathcal{F}$. As a consequence, we show that, if $\mathcal{F}$ admits efficient adjacency labels (or, equivalently, small induced-universal graphs) meeting the information-theoretic minimum, then the classes of subgraphs and induced subgraphs of Cartesian products of graphs in $\mathcal{F}$ do too. Our proof uses ideas from randomized communication complexity and hashing, and improves upon recent results of Chepoi, Labourel, and Ratel [Journal of Graph Theory, 2020].

2012 ACM Subject Classification Mathematics of computing → Graph theory; Mathematics of computing → Combinatorics

Keywords and phrases Adjacency labeling schemes, Cartesian product, Hypercubes

Digital Object Identifier 10.4230/LIPIcs.ICALP.2023.57

Category Track A: Algorithms, Complexity and Games

Funding *Louis Esperet*: Partially supported by the French ANR Projects GATO (ANR-16-CE40-0009-01), GrR (ANR-18-CE40-0032), TWIN-WIDTH (ANR-21-CE48-0014-01) and by LabEx PERSYVAL-lab (ANR-11-LABX-0025).
Nathaniel Harms: This work was partly funded by NSERC, and was done while the author was a student at the University of Waterloo, visiting Laboratoire G-SCOP and the University of Liverpool.

Acknowledgements We are very grateful to Sebastian Wild, who prevented us trying to reinvent perfect hashing.

1 Introduction

In this paper, we present optimal adjacency labeling schemes (equivalently, induced-universal graph constructions) for subgraphs of Cartesian products, which essentially closes a recent line of work studying these objects [1, 2, 3, 4, 8, 10].

Adjacency labeling

A *class* of graphs is a set $\mathcal{F}$ of graphs closed under isomorphism, where the set $\mathcal{F}_n \subseteq \mathcal{F}$ of graphs on n vertices has vertex set $[n]$. It is *hereditary* if it is also closed under taking induced subgraphs, and *monotone* if it is also closed under taking subgraphs. An *adjacency labeling scheme* for a class $\mathcal{F}$ consists of a *decoder* $D : \{0,1\}^* \times \{0,1\}^* \to \{0,1\}$ such that for every $G \in \mathcal{F}$ there exists a *labeling* $\ell : V(G) \to \{0,1\}^*$ satisfying

$$\forall x, y \in V(G) : \qquad D(\ell(x), \ell(y)) = 1 \iff xy \in E(G)\,.$$

50th International Colloquium on Automata, Languages, and Programming (ICALP 2023).
Editors: Kousha Etessami, Uriel Feige, and Gabriele Puppis; Article No. 57; pp. 57:1–57:11
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

The *size* of the adjacency labeling scheme (or *labeling scheme* for short) is the function $n \mapsto \max_{G \in \mathcal{F}_n} \max_{x \in V(G)} |\ell(x)|$, where $|\ell(x)|$ is the number of bits of $\ell(x)$. Labeling schemes have been studied extensively since their introduction by Kannan, Naor, & Rudich [13] and Muller [15]. If $\mathcal{F}$ admits a labeling scheme of size $s(n)$, then a graph $G \in \mathcal{F}_n$ can be recovered from the $n \cdot s(n)$ total bits in the adjacency labels of its vertices, so a labeling scheme is an encoding of the graph, distributed among its vertices. The information-theoretic lower bound on any encoding is $\log |\mathcal{F}_n|$, so the question is, when can the distributed adjacency labeling scheme approach this bound? In other words, which classes of graphs admit labeling schemes of size $O(\frac{1}{n} \log |\mathcal{F}_n|)$? We will say that a graph class has an *efficient* labeling scheme if it either has a labeling scheme of size $O(1)$ (i.e. it satisfies $\log |\mathcal{F}_n| = o(n \log n)$ [16]), or $O(\frac{1}{n} \log |\mathcal{F}_n|)$.

Cartesian products

Write $G \square H$ for the Cartesian product of G and H, write G^d for the d-wise Cartesian product of G, and for any class $\mathcal{F}$ write $\mathcal{F}^\square = \{G_1 \square G_2 \square \cdots \square G_d : d \in \mathbb{N}, G_i \in \mathcal{F}\}$ for the class of Cartesian products of graphs in $\mathcal{F}$. A vertex x of $G_1 \square \cdots \square G_d$ can be written $x = (x_1, \ldots, x_d)$ where $x_i \in V(G_i)$ and two vertices x, y are adjacent if and only if they differ on exactly one coordinate $i \in [d]$, and on this coordinate $x_i y_i \in E(G_i)$. Write $\mathsf{mon}(\mathcal{F}^\square)$ and $\mathsf{her}(\mathcal{F}^\square)$, respectively, for the monotone and hereditary closures of this class, which are the sets of all graphs G that are a subgraph (respectively, induced subgraph) of some $H \in \mathcal{F}^\square$.

We will construct optimal labeling schemes for $\mathsf{mon}(\mathcal{F}^\square)$ and $\mathsf{her}(\mathcal{F}^\square)$ from an optimal labeling scheme for $\mathcal{F}$. Cartesian products appear several times independently in the recent literature on labeling schemes [3, 8, 2] (and later in [10, 1, 4]), and are extremely natural for the problem of adjacency labeling for a few reasons.

First, for example, if $\mathcal{F}$ is the class of complete graphs, a labeling scheme for $\mathsf{her}(\mathcal{F}^\square)$ is equivalent to an encoding $\ell : T \to \{0,1\}^*$ of strings $T \subseteq \Sigma^*$, with Σ being an arbitrarily large finite alphabet, such that a decoder who doesn't know T can decide whether $x, y \in T$ have Hamming distance 1, using only the encodings $\ell(x)$ and $\ell(y)$. Replacing complete graphs with, say, paths, one obtains induced subgraphs of grids in arbitrary dimension. Switching to $\mathsf{mon}(\mathcal{F}^\square)$ allows arbitrary edges of these products to be deleted.

Second, Cartesian product graphs admit, by definition, a natural but inefficient "implicit representation", meaning (informally) that the adjacency between two vertices x and y can be verified by examining their representation (in this case, the tuples $x = (x_1, \ldots, x_d)$ and $y = (y_1, \ldots, y_d)$). Formalizing and quantifying this general notion was the motivation for labeling schemes in [13], who also observed that adjacency labeling schemes are equivalent to *induced-universal graphs* (or simply *universal graphs*). A sequence of graphs $(U_n)_{n \in \mathbb{N}}$ are universal graphs of size $n \mapsto |U_n|$ for a class $\mathcal{F}$ if each n-vertex graph $G \in \mathcal{F}$ is an induced subgraph of U_n. A labeling scheme of size $s(n)$ is equivalent to a universal graph of size $2^{s(n)}$, and Cartesian product graphs admit natural but inefficient universal graphs: if $(U_n)_{n \in \mathbb{N}}$ are universal graphs for $\mathcal{F}$ then for large enough $d = d(n)$, the graphs $(U_n^d)_{n \in \mathbb{N}}$ are universal for $\mathsf{her}(\mathcal{F}^\square)$. In general, this construction has exponential size: the hypercubes K_2^d are themselves universal for $\mathsf{her}(\{K_2\}^\square)$, but a star with $n-1$ leaves cannot be embedded in K_2^d for $d < n-1$, so these universal graphs are of size at least 2^{n-1}. It is not clear *a priori* whether it is possible to use the universal graphs for the base class $\mathcal{F}$ to obtain more efficient universal graphs for $\mathsf{her}(\mathcal{F}^\square)$, and even less clear for $\mathsf{mon}(\mathcal{F}^\square)$, but we will show in this paper how to do so.

Finally, there was the possibility that *subgraphs* of Cartesian products could provide the first explicit counterexample to the Implicit Graph Conjecture (IGC) of [13, 17], which suggested that the condition $\log |\mathcal{F}_n| = O(n \log n)$ was *sufficient* for $\mathcal{F}$ to admit a labeling scheme of size $O(\log n)$; this was refuted by a non-constructive counting argument in a recent breakthrough of Hatami & Hatami [11]. There is a labeling scheme of size $O(\log^2 n)$ for the subgraphs of hypercubes, due to a folklore bound of $\log n$ on the degeneracy of this class (see [5]) and a general $O(k \log n)$ labeling scheme for classes of degeneracy k [13]. Designing an efficient labeling scheme for *induced* subgraphs of hypercubes (rather, the weaker question of proving bounds on $|\mathcal{F}_n|$ for this family) was an open problem of Alecu, Atminas, & Lozin [2], resolved concurrently and independently in [8]; this also gave an example of a class with an efficient labeling scheme but unbounded *functionality*, answering another open question of [2]. Also independently, Chepoi, Labourel, & Ratel [3] studied the structure of general Cartesian products, motivated by the problem of designing labeling schemes for the classes $\mathsf{mon}(\mathcal{F}^\square)$. They give upper bounds (via bounds on the degeneracy) for a number of special cases but do not improve on the $O(\log^2 n)$ bound for hypercubes. The following 3 observations then suggested that subgraphs of Cartesian products could give the first explicit counterexample to the IGC (and this was posed as an open problem in [4]):

1. It is shown in [4] that, while *induced* subgraphs of hypercubes have a constant-size *adjacency sketch* (a probabilistic version of a labeling scheme), the *subgraphs* of hypercubes do not, so, with respect to *randomized* labels, subgraphs are more complex than induced subgraphs.
2. The above result shows that the class of subgraphs of hypercubes is a counterexample to a conjecture of [10]. That conjecture was refuted earlier by a construction of [7] that, with some extension, refuted the IGC itself [11].
3. The previous work considering Cartesian products [3, 8, 10, 2, 1] had not improved on the $O(\log^2 n)$ bound for subgraphs.

Alas, a consequence of our main result is that subgraphs of Cartesian products are *not* counterexamples to the IGC.

Results and techniques

We improve the best-known $O(\log^2 n)$ bound for subgraphs of hypercubes to the optimal $O(\log n)$, and in general show how to construct optimal labels for all subgraphs and induced subgraphs of Cartesian products. Our proof is short, and departs significantly from standard techniques in the field of labeling schemes: we do not rely on any structural results, graph width parameters, or decompositions, and instead use communication complexity (as in [8, 10]), encoding, and hashing arguments, which may be useful for future work on labeling schemes. We prove:

▶ **Theorem 1.** *Let $\mathcal{F}$ be a hereditary class with an adjacency labeling scheme of size $s(n)$. Then:*

1. $\mathsf{her}(\mathcal{F}^\square)$ *has a labeling scheme of size at most $4s(n) + O(\log n)$.*
2. $\mathsf{mon}(\mathcal{F}^\square)$ *has a labeling scheme where each $G \in \mathsf{mon}(\mathcal{F}^\square)$ on n vertices is given labels of size at most $4s(n) + O(k(G) + \log n)$, where $k(G)$ is the degeneracy of G.*

We allow $\mathcal{F}$ to be finite, in which case $s(n) = O(1)$; in particular, setting $\mathcal{F} = \{K_2, K_1\}$, we get the result for hypercubes:

▶ **Corollary 2.** *Let $\mathcal{H}$ be the class of hypercube graphs. Then $\mathsf{mon}(\mathcal{H})$ has a labeling scheme of size $O(\log n)$.*

All of the labeling schemes of Chepoi, Labourel, & Ratel [3] are obtained by bounding $k(G)$ and applying the black-box $O(k(G) \cdot \log n)$ bound of [13]. For example, they get labels of size $O(d \log^2 n)$ when the base class $\mathcal{F}$ has degeneracy d, by showing that $\mathsf{mon}(\mathcal{F}^\square)$ has degeneracy $O(d \log n)$. Our result can be substituted for that black-box, replacing the multiplicative $O(\log n)$ with an *additive* $O(\log n)$, thereby improving all of the results of [3] when combined with their bounds on $k(G)$; for example, achieving $O(d \log n)$ when $\mathcal{F}$ has degeneracy d.

For subgraphs of hypercubes, [3] observed that a bound of $O(\mathsf{vc}(G) \log n)$ follows from the inequality $k(G) \leq \mathsf{vc}(G)$ due to Haussler [12], where $\mathsf{vc}(G)$ is the VC dimension[1], which can be as large as $\log n$ but is often much smaller; they generalize this inequality in various ways to other Cartesian products. Our result supercedes the VC dimension result for hypercubes.

Theorem 1 is optimal up to constant factors (which we have not tried to optimize), and yields the following corollary (see Section 3 for proofs).

▶ **Corollary 3.** *If a hereditary class $\mathcal{F}$ has an efficient labeling scheme, then so do* $\mathsf{her}(\mathcal{F}^{\square})$ *and* $\mathsf{mon}(\mathcal{F}^{\square})$.

One of our main motivations was to find explicit counterexamples to the IGC; a consequence of the above corollary is that, counterexamples to the IGC cannot be obtained by taking the monotone closure of Cartesian products of some hereditary class $\mathcal{F}$, unless $\mathcal{F}$ itself is already a counterexample. This leaves open the problem of finding an explicit counterexample to the IGC, which would require developing the first lower-bound technique for adjacency labeling schemes.

2 Adjacency Labeling Scheme

Notation

For two binary strings x, y, we write $x \oplus y$ for the bitwise XOR. For two graphs G and H, we will write $G \subset H$ if G is a subgraph of H, and $G \subset_I H$ if G is an *induced* subgraph of H. We will write $V(G)$ and $E(G)$ as the vertex and edge set of a graph G, respectively. All graphs in this paper are simple and undirected. A graph G has *degeneracy* k if all subgraphs of G have a vertex of degree at most k.

Strategy

Suppose $G \subset G_1 \square \cdots \square G_d$ is a subgraph of a Cartesian product. Then $V(G) \subseteq V(G_1) \times \cdots \times V(G_d)$. Let $H \subset_I G_1 \square \cdots \square G_d$ be the subgraph induced by $V(G)$, so that $E(G) \subseteq E(H)$. One may think of G as being obtained from the induced subgraph H by deleting some edges. Then two vertices $x, y \in V(G)$ are adjacent if and only if:

1. There exists exactly one coordinate $i \in [d]$ where $x_i \neq y_i$;
2. On this coordinate, $x_i y_i \in E(G_i)$; and,
3. The edge $xy \in E(H)$ has not been deleted in $E(G)$.

We construct the labels for vertices in G in three phases, which check these conditions in sequence.

2.1 Phase 1: Exactly One Difference

We give two proofs for Phase 1. The first is a reduction to the k-Hamming Distance communication protocol. The second proof is direct and self-contained; it is an extension of the proof of the labeling scheme for induced subgraphs of hypercubes, in the unpublished note [9] (adapted from [8, 10]). In both cases the labels are obtained by the probabilistic method, and are efficiently computable by a randomized algorithm.

[1] See [3] for the definition of VC dimension

For any alphabet Σ and any two strings $x, y \in \Sigma^d$ where $d \in \mathbb{N}$, write $\mathsf{dist}(x,y)$ for the Hamming distance between x and y, i.e. $\mathsf{dist}(x,y) = |\{i \in [d] : x_i \neq y_i\}|$.

For the first proof, we require a result in communication complexity (which we translate into our terminology). A version with two-sided error appears in [18], the one-sided error version below is implicit in [10] (and may appear elsewhere in the literature, which we did not find).

▶ **Theorem 4** ([18, 10]). *There exists a constant $c > 0$ satisfying the following. For any $k \in \mathbb{N}$, there exists a function $D : \{0,1\}^* \times \{0,1\}^* \to \{0,1\}$ such that, for any $d \in \mathbb{N}$ and set $S \subseteq \{0,1\}^d$ of size $|S| = n$, there exists a probability distribution L over functions $\ell : S \to \{0,1\}^{ck^2}$, where for all $x, y \in S$,*

1. *If $\mathsf{dist}(x,y) \leq k$ then $\underset{\ell \sim L}{\mathbb{P}}\left[D(\ell(x), \ell(y)) = 1\right] = 1$; and,*
2. *If $\mathsf{dist}(x,y) > k$ then $\underset{\ell \sim L}{\mathbb{P}}\left[D(\ell(x), \ell(y)) = 0\right] \geq 2/3$.*

We transform these randomized labels into deterministic labels using standard arguments:

▶ **Proposition 5.** *There exists a constant $c > 0$ satisfying the following. For any $k \in \mathbb{N}$, there exists a function $D : \{0,1\}^* \times \{0,1\}^* \to \{0,1\}$ such that, for any $d \in \mathbb{N}$ and set $S \subseteq \{0,1\}^d$ of size $|S| = n$, there exists a function $\ell : S \to \{0,1\}^{ck^2 \log n}$ where for all $x, y \in S$, $D(\ell(x), \ell(y)) = 1$ if and only if $\mathsf{dist}(x,y) \leq k$.*

Proof. Let $D' : \{0,1\}^* \times \{0,1\}^* \to \{0,1\}$, $c > 0$, and L be the function, the constant, and the probability distribution given for S by Theorem 4. Let $q = \lceil 2 \log_3 n \rceil$, and let L' be the distribution over functions defined by choosing $\ell_1, \dots, \ell_q \sim L$ independently at random, and setting $\ell(x) = (\ell_1(x), \ell_2(x), \dots, \ell_q(x))$ for each $x \in S$. Define $D : \{0,1\}^* \times \{0,1\}^* \to \{0,1\}$ such that

$$D(\ell(x), \ell(y)) = \bigwedge_{i=1}^{q} D'(\ell_i(x), \ell_i(y)) .$$

Observe that, if $x, y \in S$ have $\mathsf{dist}(x,y) \leq k$ then $\mathbb{P}\left[D(\ell(x), \ell(y)) = 1\right] = 1$ since for each $i \in [q]$ we have $\mathbb{P}\left[D'(\ell_i(x), \ell_i(y)) = 1\right] = 1$. On the other hand, if $x, y \in S$ have $\mathsf{dist}(x,y) > k$, then

$$\mathbb{P}\left[D(\ell(x), \ell(y)) = 1\right] < (1/3)^q \leq 1/n^2 .$$

By the union bound, the probability that there exist $x, y \in S$ such that $D(\ell(x), \ell(y))$ takes the incorrect value is strictly less than 1. Therefore there exists a fixed function $\ell : S \to \{0,1\}^{ck^2 q}$ satisfying the required conditions, where $ck^2 q = Ck^2 \log n$ for an appropriate constant C. ◀

We reduce the problem for alphabets Σ to the 2-Hamming Distance labeling problem above.

▶ **Lemma 6.** *There exists a function $D : \{0,1\}^* \times \{0,1\}^* \to \{0,1\}$ and a constant $c > 0$ such that, for any countable alphabet Σ, any $d \in \mathbb{N}$, and any set $S \subseteq \Sigma^d$ of size $|S| = n$, there exists a function $\ell : S \to \{0,1\}^k$ for $k \leq c \log n$, where $D(\ell(x), \ell(y)) = 1$ if and only if $\mathsf{dist}(x,y) = 1$.*

Proof. Since $\lceil \log n \rceil$ bits can be added to any $\ell(x)$ to ensure that $\ell(x)$ is unique, it suffices to construct functions D, ℓ where $D(\ell(x), \ell(y)) = 1$ if and only if $\mathsf{dist}(x,y) \leq 1$, instead of $\mathsf{dist}(x,y) = 1$ exactly.

Since S has at most n elements, we may assume that Σ has a finite number N of elements, since we may reduce to the set of elements which appear in the strings S. We may then identify Σ with $[N]$ and define an encoding $\mathsf{enc} : [N] \to \{0,1\}^N$ where for any $\sigma \in [N]$, $\mathsf{enc}(\sigma)$ is the string that takes value 1 on coordinate σ, and all other coordinates take value 0.

Abusing notation, for any $x \in \Sigma^d$, we may now define the concatenated encoding $\mathsf{enc}(x) = \mathsf{enc}(x_1) \circ \mathsf{enc}(x_2) \circ \cdots \circ \mathsf{enc}(x_d)$, where $\circ$ denotes concatenation. It is easy to verify that for any $x, y \in \Sigma^d$, $\mathsf{dist}(\mathsf{enc}(x), \mathsf{enc}(y)) = 2 \cdot \mathsf{dist}(x, y)$. We may therefore apply Proposition 5 with $k = 2$ on the set $S' = \{\mathsf{enc}(x) : x \in S\}$ to obtain a function $D : \{0,1\}^* \times \{0,1\}^* \to \{0,1\}$, a constant $C > 0$, and a function $\ell' : S' \to \{0,1\}^{C \log n}$ such that for all $x, y \in S$,

$$D(\ell'(\mathsf{enc}(x)), \ell'(\mathsf{enc}(y))) = 1 \iff \mathsf{dist}(\mathsf{enc}(x), \mathsf{enc}(y)) \leq 2 \iff \mathsf{dist}(x, y) \leq 1 \,.$$

We may then conclude the proof by setting $\ell(x) = \ell'(\mathsf{enc}(x))$ for each $x \in S$. ◀

Below, we give an alternative, direct proof that does not reduce to k-Hamming Distance.

▶ **Proposition 7.** *For any set $S \subseteq \{0,1\}^d$, there exists a random function $\ell : S \to \{0,1\}^4$ such that, for all $x, y \in S$,*

1. *If $\mathsf{dist}(x, y) \leq 1$ then $\underset{\ell}{\mathbb{P}}[\mathsf{dist}(\ell(x), \ell(y)) \leq 1] = 1$, and*
2. *If $\mathsf{dist}(x, y) > 1$ then $\underset{\ell}{\mathbb{P}}[\mathsf{dist}(\ell(x), \ell(y)) \leq 1] \leq 3/4$.*

Proof. Choose a uniformly random map $p : [d] \to [4]$ and partition $[d]$ into four sets $P_j = p^{-1}(j)$. For each $i \in [4]$, define $\ell(x)_i := \bigoplus_{j \in P_i} x_j$.

Let $x, y \in S$ and write $w = \ell(x) \oplus \ell(y)$. Note that $\mathsf{dist}(\ell(x), \ell(y)) = |w|$, which is the number of 1s in w. If $\mathsf{dist}(x, y) = 0$ then $\mathsf{dist}(\ell(x), \ell(y)) = 0 \leq 1$. Now suppose $\mathsf{dist}(x, y) = 1$. For any choice of $p : [d] \to [4]$, one of the sets P_i contains the differing coordinate and will have $w_i = 1$, while the other three sets P_j will have $w_j = 0$, so $\underset{\ell}{\mathbb{P}}[\mathsf{dist}(\ell(x), \ell(y)) \leq 1] = 1$.

Now suppose $\mathsf{dist}(x, y) = t \geq 2$. We will show that $|w| \leq 1$ with probability at most $3/4$. Note that w is obtained by the random process where $\vec{0} = w^{(0)}, w = w^{(t)}$, and $w^{(i)}$ is obtained from $w^{(i-1)}$ by flipping a uniformly random coordinate.

Observe that, for $i \geq 1$, $\mathbb{P}\left[w^{(i)} = \vec{0}\right] \leq 1/4$. This is because $w^{(i)} = \vec{0}$ can occur only if $|w^{(i-1)}| = 1$, so the probability of flipping the 1-valued coordinate is $1/4$. If $|w^{(i-1)}| \geq 1$ then $\mathbb{P}\left[|w^{(i)}| \leq 1 \;\;|\;\; |w^{(i-1)}| \geq 1\right] \leq 1/2$ since either $|w^{(i-1)}| = 1$ and then $|w^{(i)}| = 0 \leq 1$ with probability $1/4$, or $|w^{(i-1)}| \geq 2$ and $|w^{(i)}| = 1$ with probability at most $1/2$. Then, for $t \geq 2$,

$$\begin{aligned}\mathbb{P}\left[|w^{(t)}| \leq 1\right] &= \mathbb{P}\left[w^{(t-1)} = \vec{0}\right] + \mathbb{P}\left[|w^{(t-1)}| \geq 1\right] \cdot \mathbb{P}\left[|w^{(t)}| = 1 \;\;|\;\; |w^{(t-1)}| \geq 1\right] \\ &\leq \frac{1}{4} + \frac{1}{2} = \frac{3}{4}\,.\end{aligned}$$ ◀

▶ **Proposition 8.** *There exists a function $D : \{0,1\}^4 \times \{0,1\}^4 \to \{0,1\}$ such that, for any countable alphabet, Σ, any $d \in \mathbb{N}$, and any $S \subseteq \Sigma^d$ of size $n = |S|$, there exists a random function $\ell : S \to \{0,1\}^4$ such that, for all $x, y \in S$,*

1. *If $\mathsf{dist}(x, y) \leq 1$, then $\underset{\ell}{\mathbb{P}}[D(\ell(x), \ell(y)) = 1] = 1$, and*
2. *If $\mathsf{dist}(x, y) > 1$, then $\underset{\ell}{\mathbb{P}}[D(\ell(x), \ell(y)) = 1] \leq 15/16$.*

Proof. For each $\sigma \in \Sigma$ and $i \in [d]$, generate an independently and uniformly random bit $q_i(\sigma) \sim \{0,1\}$. Then for each $x \in S$ define $p(x) = (q_1(x_1), \ldots, q_d(x_d)) \in \{0,1\}^d$ and $S' = \{p(x) : x \in S\}$, and let ℓ' be the random function $S' \to \{0,1\}^4$ guaranteed to exist by Proposition 7. We define the random function $\ell : S \to \{0,1\}^4$ as $\ell(x) = \ell'(p(x))$. We define $D(\ell(x), \ell(y)) = 1$ if and only if $\mathsf{dist}(\ell'(p(x)), \ell'(p(y))) \leq 1$.

Let $x, y \in S$. If $\mathsf{dist}(x,y) \leq 1$, so there is a unique $i \in [d]$ with $x_i \neq y_i$, then

$$\mathbb{P}\left[\mathsf{dist}(p(x),p(y)) = 1\right] = \mathbb{P}\left[q_i(x_i) \neq q_i(y_i)\right] = \mathbb{P}\left[\mathsf{dist}(p(x),p(y)) = 0\right] = 1/2\,,$$

so $\mathbb{P}\left[\mathsf{dist}(p(x),p(y)) \leq 1\right] = 1$. Then by Proposition 7,

$$\mathbb{P}\left[D(\ell(x),\ell(y)) = 1\right] = \mathbb{P}\left[\mathsf{dist}(\ell'(p(x)),\ell'(p(y))) \leq 1\right] = 1\,.$$

If $\mathsf{dist}(x,y) > 1$ so that there are distinct $i, i' \in [d]$ such that $x_i \neq y_i$ and $x_{i'} \neq y_{i'}$, then

$$\mathbb{P}\left[\mathsf{dist}(p(x),p(y)) \geq 2\right] \geq \mathbb{P}\left[q_i(x_i) \neq q_i(y_i) \wedge q_{i'}(x_{i'}) \neq q_{i'}(y_{i'})\right] = 1/4\,.$$

Then by Proposition 7,

$$\begin{aligned}\mathbb{P}\left[D(\ell(x),\ell(y)) = 1\right] &= \mathbb{P}\left[\mathsf{dist}(\ell'(p(x)),\ell'(p(y))) \leq 1\right] \\ &= \mathbb{P}\left[\mathsf{dist}(p(x),p(y)) \leq 1 \vee \mathsf{dist}(\ell'(p(x)),\ell'(p(y))) \leq 1\right] \\ &\leq 3/4 + (1-3/4)(3/4) = 15/16\,.\end{aligned}$$ ◀

The alternative proof of Lemma 6 now concludes by using Proposition 8 with a nearly identical derandomization argument as in Proposition 5.

2.2 Phase 2: Induced Subgraphs

After the first phase, we are guaranteed that there is a unique coordinate $i \in [d]$ where $x_i \neq y_i$. In the second phase we wish to determine whether $x_i y_i \in E(G_i)$. It is convenient to have labeling schemes for the factors $G_1, \dots, G_d$ where we can XOR the labels together while retaining the ability to compute adjacency. Define an *XOR-labeling scheme* the same as an adjacency labeling scheme, with the restriction that for each $s \in \mathbb{N}$ there is some function $g_s : \{0,1\}^s \to \{0,1\}$ such that on any two labels $\ell(x), \ell(y)$ of size s, the decoder outputs $D(\ell(x),\ell(y)) = g_s(\ell(x) \oplus \ell(y))$. Any labeling scheme can be transformed into an XOR-labeling scheme with at most a constant-factor loss:

▶ **Lemma 9.** *Let $\mathcal{F}$ be any class of graphs with an adjacency labeling scheme of size $s(n)$. Then $\mathcal{F}$ admits an XOR-labeling scheme of size at most $4s(n)$.*

Proof. Let $D : \{0,1\}^* \times \{0,1\}^* \to \{0,1\}$ be the decoder of the adjacency labeling scheme for $\mathcal{F}$, fix any $n \in \mathbb{N}$, and write $s = s(n)$. Note that D must be symmetric, so $D(a,b) = D(b,a)$ for any $a, b \in \{0,1\}^s$. Let $\phi : \{0,1\}^s \to \{0,1\}^{4s}$ be uniformly randomly chosen, so that for every $z \in \{0,1\}^s$, $\phi(z) \sim \{0,1\}^{4s}$ is a uniform and independently random variable. For any two distinct pairs $\{z_1,z_2\}, \{z_1',z_2'\} \in \binom{\{0,1\}^s}{2}$ where $z_1 \neq z_2$, $z_1' \neq z_2'$, and $\{z_1,z_2\} \neq \{z_1',z_2'\}$, the probability that $\phi(z_1) \oplus \phi(z_2) = \phi(z_1') \oplus \phi(z_2')$ is at most 2^{-4s}, since at least one of the variables $\phi(z_1), \phi(z_2), \phi(z_1'), \phi(z_2')$ is independent of the other ones. Therefore, by the union bound,

$$\mathbb{P}\left[\exists \{z_1,z_2\}, \{z_1',z_2'\} : \phi(z_1) \oplus \phi(z_2) = \phi(z_1') \oplus \phi(z_2')\right] \leq \binom{2^s}{2}^2 2^{-4s} \leq \frac{1}{4}\,.$$

Then there is $\phi : \{0,1\}^s \to \{0,1\}^{4s}$ such that each distinct pair $\{z_1,z_2\} \in \binom{\{0,1\}^s}{2}$ is assigned has a distinct unique value $\phi(z_1) \oplus \phi(z_2)$. So the function $\Phi(\{z_1,z_2\}) := \phi(z_1) \oplus \phi(z_2)$ is a one-to-one map $\binom{\{0,1\}^s}{2} \to \{0,1\}^{4s}$. Then for any graph $G \in \mathcal{F}$ on n vertices, with labeling $\ell : V(G) \to \{0,1\}^s$, we may assign the new label $\phi(\ell(x))$ to each vertex x. On labels $\phi(\ell(x)), \phi(\ell(y)) \in \{0,1\}^s$, the decoder for the XOR-labeling scheme simply computes $\{\ell(x),\ell(y)\} = \Phi^{-1}(\phi(\ell(x)) \oplus \phi(\ell(y)))$ and outputs $D(\ell(x),\ell(y))$, where we are using the fact that $D(\ell(x),\ell(y)) = D(\ell(y),\ell(x))$, so that the ordering of the pair $\{\ell(x),\ell(y)\}$ does not matter. ◀

We can now prove the first part of Theorem 1.

▶ **Lemma 10.** *Let $\mathcal{F}$ be a hereditary class of graphs that admits an adjacency labeling scheme of size $s(n)$. Then $\mathsf{her}(\mathcal{F}^\square)$ admits an adjacency labeling scheme of size $4s(n) + O(\log n)$.*

Proof. By Lemma 9, there is an XOR-labeling scheme for $\mathcal{F}$ with labels of size $4s(n)$. Let $D : \{0,1\}^* \times \{0,1\}^* \to \{0,1\}$ be the decoder for this scheme, with $D(a, b) = g(a \oplus b)$ for some function g. Design the labels for $\mathsf{her}(\mathcal{F}^\square)$ as follows. Consider a graph $G \in \mathsf{her}(\mathcal{F}^\square)$, so that $G \subset_I G_1 \square G_2 \square \cdots \square G_d$ for some $d \in \mathbb{N}$ and $G_i \in \mathcal{F}$ for each $i \in [d]$. Since $\mathcal{F}$ is hereditary, we may assume that each G_i has at most n vertices; otherwise we could simply replace it with the subgraph of G_i induced by the vertices $\{x_i : x \in V(G)\}$. For each $x = (x_1, \dots, x_d) \in V(G)$, construct the label as follows:

1. Treating the vertices in each G_i as characters of the alphabet $[n]$, use $O(\log n)$ bits to assign the label given to $x = (x_1, \dots, x_d) \in [n]^d$ by Lemma 6.
2. Using $4s(n)$ bits, append the vector $\bigoplus_{i \in [d]} \ell_i(x_i)$, where $\ell_i(x_i)$ is the label of $x_i \in V(G_i)$ in graph G_i, according to the XOR-labeling scheme for $\mathcal{F}$.

The decoder operates as follows. Given the labels for $x, y \in V(G)$:

1. If x and y differ on exactly one coordinate, as determined by the first part of the label, continue to the next step. Otherwise output "not adjacent".
2. Now guaranteed that there is a unique $i \in [d]$ such that $x_i \neq y_i$, output "adjacent" if and only if the following is 1:

$$\begin{aligned} D\left(\bigoplus_{j \in [d]} \ell_j(x_j), \bigoplus_{j \in [d]} \ell_j(y_j)\right) &= g\left(\bigoplus_{j \in [d]} \ell_j(x_j) \oplus \bigoplus \ell_j(y_j)\right) \\ &= g\left(\ell_i(x_i) \oplus \ell_i(y_i) \oplus \bigoplus_{j \neq i} \ell_j(x_j) \oplus \ell_j(y_j)\right) \\ &= g(\ell_i(x_i) \oplus \ell_i(y_i)), \end{aligned}$$

   where the final equality holds because $x_j = y_j$ for all $j \neq i$, so $\ell_j(x_j) = \ell_j(y_j)$. Then the output value is 1 if and only $x_i y_i$ is an edge of G_i; equivalently, xy is an edge of G.

This concludes the proof. ◀

The XOR-labeling trick can also be used to simplify the proof of [10] for adjacency *sketches* of Cartesian products. That proof is similar to the one above, except it uses a two-level hashing scheme and some other tricks to avoid destroying the labels of x_i and y_i with the XOR (with sufficiently large probability of success). This two-level hashing approach does not succeed in our current setting, and we avoid it with XOR-labeling.

2.3 Phase 3: Subgraphs

Finally, we must check whether the edge $xy \in E(H)$ in the *induced* subgraph $H \subset_I G_1 \square \cdots \square G_d$ has been deleted in $E(G)$. There is a minimal and perfect tool for this task:

▶ **Theorem 11** (Minimal Perfect Hashing). *For every $m, k \in \mathbb{N}$, there is a family $\mathcal{P}_{m,k}$ of hash functions $[m] \to [k]$ such that, for any $S \subseteq [m]$ of size k, there exists $h \in \mathcal{P}_{m,k}$ where the image of S under h is $[k]$ and for every distinct $i, j \in S$ we have $h(i) \neq h(j)$. The function h can be stored in $k \ln e + \log\log m + o(k + \log\log m)$ bits of space and it can be computed by a randomized algorithm in expected time $O(k + \log\log m)$.*

Minimal perfect hashing has been well-studied. A proof of the space bound appears in [14] and significant effort has been applied to improving the construction and evaluation time. We take the above statement from [6]. We now conclude the proof of Theorem 1 by applying the next lemma to the class $\mathcal{G} = \mathsf{her}(\mathcal{F}^{\square})$, using the labeling scheme for $\mathsf{her}(\mathcal{F}^{\square})$ obtained in Lemma 10 (note that $\mathsf{mon}(\mathsf{her}(\mathcal{F}^{\square})) = \mathsf{mon}(\mathcal{F}^{\square})$).

▶ **Lemma 12.** *Let $\mathcal{G}$ be any graph class which admits an adjacency labeling scheme of size $s(n)$. Then $\mathsf{mon}(\mathcal{G})$ admits an adjacency labeling scheme where each $G \in \mathsf{mon}(\mathcal{G})$ on n vertices has labels of size $s(n) + O(k(G) + \log n)$, where $k(G)$ is the degeneracy of G.*

Proof. Let $G \in \mathsf{mon}(\mathcal{G})$ have n vertices, so that it is a subgraph of $H \in \mathcal{G}$ on n vertices. The labeling scheme is as follows.

1. Fix a total order $\prec$ on $V(H)$ such that each vertex x has at most $k = k(G)$ neighbors y in H such that $x \prec y$; this exists by definition. We will identify each vertex x with its position in the order.
2. For each vertex x, assign the label as follows:
   a. Use $s(n)$ bits for the adjacency label of x in H.
   b. Use $\log n$ bits to indicate x (the position in the order).
   c. Let $N^+(x)$ be the set of neighbors $x \prec y$. Construct a perfect hash function $h_x : N^+(x) \to [k]$ and store it, using $O(k + \log\log n)$ bits.
   d. Use k bits to write the function $\mathsf{edge}_x : [k] \to \{0, 1\}$ which takes value 1 on $i \in [k]$ if and only if xy is an edge of G, where y is the unique vertex in $N^+(x)$ satisfying $h_x(y) = i$.

Given the labels for x and y, the decoder performs the following:

1. If xy are not adjacent in H, output "not adjacent".
2. Otherwise xy are adjacent. If $x \prec y$, we are guaranteed that y is in the domain of h_x, so output "adjacent" if and only if $\mathsf{edge}_x(h_x(y)) = 1$. If $y \prec x$, output "adjacent" if and only if $\mathsf{edge}_y(h_y(x)) = 1$.

This concludes the proof. ◀

3 Optimality

We now prove the optimality of our labeling schemes, and Corollary 3. We require:

▶ **Proposition 13.** *For any hereditary class $\mathcal{F}$, let $k(n)$ be the maximum degeneracy of an n-vertex graph $G \in \mathsf{her}(\mathcal{F}^{\square})$. Then $\mathsf{her}(\mathcal{F}^{\square})$ contains a graph H on n vertices with at least $n \cdot k(n)/4$ edges, so $\mathsf{mon}(\mathcal{F}^{\square})$ contains all $2^{n \cdot k(n)/4}$ spanning subgraphs of H.*

Proof. Since G has degeneracy $k = k(n)$, it contains an induced subgraph $G' \subset_I G$ with minimum degree k and $n_1 \leq n$ vertices. If $n_1 \geq n/2$ then G itself has at least $kn_1/2 \geq kn/4$ edges, and we are done. Now assume $n_1 < n/2$. Since $G \in \mathsf{her}(\mathcal{F}^{\square})$, $G \subset_I H_1 \square \cdots \square H_t$ for some $t \in \mathbb{N}$ and $H_i \in \mathcal{F}$. So for any $d \in \mathbb{N}$, the graph $(G')^d \subset_I (H_1 \square \cdots \square H_t)^d$ belongs to $\mathsf{her}(\mathcal{F}^{\square})$. Consider the graph $H \subset_I (G')^d$ defined as follows. Choose any $w \in V(G')$, and for each $i \in [d]$ let

$$V_i = \{(v_1, v_2, \ldots, v_d) : v_i \in V(G') \text{ and } \forall j \neq i, v_j = w\},$$

and let H be the graph induced by vertices $V_1 \cup \cdots \cup V_d$. Then H has dn_1 vertices, each of degree at least k, since each $v \in V_i$ is adjacent to k other vertices in V_i. Set $d = \lceil n/n_1 \rceil$, so that H has at least n vertices, and let $m = dn_1 - n$, which satisfies $m < n_1$. Remove any m vertices of V_1. The remaining graph H' has n vertices, and at least $(d-1)n_1 \geq n - n_1 > n/2$ vertices of degree k. Then H' has at least $kn/4$ edges. ◀

The next proposition shows that Theorem 1 is optimal up to constant factors. It is straightforward to check that this proposition implies Corollary 3.

▶ **Proposition 14.** *Let $\mathcal{F}$ be a hereditary class whose optimal adjacency labeling scheme has size $s(n)$ and which contains a graph with at least one edge. Then any adjacency labeling scheme for $\mathsf{her}(\mathcal{F}^\square)$ has size at least $\Omega(s(n) + \log n)$, and any adjacency labeling scheme for $\mathsf{mon}(\mathcal{F}^\square)$ has size at least $\Omega(s(n) + k(n) + \log n)$, where $k(n)$ is the maximum degeneracy of any n-vertex graph in $\mathsf{mon}(\mathcal{F}^\square)$.*

Proof. Since $\mathcal{F} \subseteq \mathsf{her}(\mathcal{F}^\square)$ and $\mathcal{F} \subseteq \mathsf{mon}(\mathcal{F}^\square)$, we have a lower bound of $s(n)$ for the labeling schemes for both of these classes. Since $\mathcal{F}$ contains a graph G with at least one edge, the Cartesian products contain the class of hypercubes: $\mathsf{her}(\{K_2\}^\square) \subseteq \mathsf{her}(\mathcal{F}^\square) \subseteq \mathsf{mon}(\mathcal{F}^\square)$. A labeling scheme for $\mathsf{her}(\{K_2\}^\square)$ must have size $\Omega(\log n)$ (which can be seen since each vertex of K_2^d has a unique neighborhood and thus requires a unique label). This establishes the lower bound for $\mathsf{her}(\mathcal{F}^\square)$, since the labels must have size $\max\{s(n), \Omega(\log n)\} = \Omega(s(n) + \log n)$. Finally, by Proposition 13, the number of n-vertex graphs in $\mathsf{mon}(\mathcal{F}^\square)$ is at least $2^{\Omega(nk(n))}$, so there is a lower bound on the label size of $\Omega(k(n))$, which implies a lower bound of $\max\{s(n), \Omega(\log n), \Omega(k(n))\} = \Omega(s(n) + k(n) + \log n)$ for $\mathsf{mon}(\mathcal{F}^\square)$. ◀

References

1 Bogdan Alecu, Vladimir E. Alekseev, Aistis Atminas, Vadim Lozin, and Viktor Zamaraev. Graph parameters, implicit representations and factorial properties. In submission, 2022.

2 Bogdan Alecu, Aistis Atminas, and Vadim Lozin. Graph functionality. *Journal of Combinatorial Theory, Series B*, 147:139–158, 2021.

3 Victor Chepoi, Arnaud Labourel, and Sébastien Ratel. On density of subgraphs of Cartesian products. *Journal of Graph Theory*, 93(1):64–87, 2020.

4 Louis Esperet, Nathaniel Harms, and Andrey Kupavskii. Sketching distances in monotone graph classes. In *International Conference on Randomization and Computation (RANDOM 2022)*, 2022.

5 Ron L Graham. On primitive graphs and optimal vertex assignments. *Annals of the New York academy of sciences*, 175(1):170–186, 1970.

6 Torben Hagerup and Torsten Tholey. Efficient minimal perfect hashing in nearly minimal space. In *Annual Symposium on Theoretical Aspects of Computer Science (STACS 2001)*, pages 317–326. Springer, 2001.

7 Lianna Hambardzumyan, Hamed Hatami, and Pooya Hatami. A counter-example to the probabilistic universal graph conjecture via randomized communication complexity. *Discrete Applied Math.*, 2022.

8 Nathaniel Harms. Universal communication, universal graphs, and graph labeling. In *11th Innovations in Theoretical Computer Science Conference (ITCS 2020)*, volume 151, page 33. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2020.

9 Nathaniel Harms. Adjacency labeling and sketching for induced subgraphs of the hypercube. `https://cs.uwaterloo.ca/~nharms/downloads/hypercube_sketch.pdf`, 2022. URL: `https://cs.uwaterloo.ca/~nharms/downloads/hypercube_sketch.pdf`.

10 Nathaniel Harms, Sebastian Wild, and Viktor Zamaraev. Randomized communication and implicit graph representations. In *54th Annual Symposium on Theory of Computing (STOC 2022)*, 2022.

11 Hamed Hatami and Pooya Hatami. The implicit graph conjecture is false. In *63rd IEEE Symposium on Foundations of Computer Science (FOCS 2022)*, 2022.

12 David Haussler. Sphere packing numbers for subsets of the Boolean n-cube with bounded Vapnik-Chervonenkis dimension. *Journal of Combinatorial Theory, Series A*, 69(2):217–232, 1995.

13 Sampath Kannan, Moni Naor, and Steven Rudich. Implicit representation of graphs. *SIAM Journal on Discrete Mathematics*, 5(4):596–603, 1992.

14 Kurt Mehlhorn. *Data Structures and Algorithms 1 Sorting and Searching*. Monographs in Theoretical Computer Science. An EATCS Series, 1. Springer Berlin Heidelberg, Berlin, Heidelberg, 1st ed. 1984. edition, 1984.

15 John H Muller. Local structure in graph classes, 1989.

16 Edward R Scheinerman. Local representations using very short labels. *Discrete mathematics*, 203(1-3):287–290, 1999.

17 Jeremy P Spinrad. *Efficient graph representations*. American Mathematical Society, 2003.

18 Andrew Chi-Chih Yao. On the power of quantum fingerprinting. In *Proceedings of the Thirty-Fifth Annual ACM Symposium on Theory of Computing (STOC)*, pages 77–81, 2003.