

UC Irvine

UC Irvine Previously Published Works

Title

The Fibonacci dimension of a graph

Permalink

<https://escholarship.org/uc/item/3sm3t41d>

Journal

Electronic Journal of Combinatorics, 18(1 #P55)

ISSN

1077-8926

Authors

Cabello, Sergio
Eppstein, David
Klavzar, Sandi

Publication Date

2011-03-11

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

The Fibonacci dimension of a graph

Sergio Cabello* David Eppstein† Sandi Klavžar‡

Submitted: Mar 13, 2009; Accepted: Feb 28, 2011; Published: Mar 11, 2011

Mathematics Subject Classification: 05C12, 05C75, 05C85

Abstract

The Fibonacci dimension $\text{fdim}(G)$ of a graph G is introduced as the smallest integer f such that G admits an isometric embedding into Γ_f , the f -dimensional Fibonacci cube. We give bounds on the Fibonacci dimension of a graph in terms of the isometric and lattice dimension, provide a combinatorial characterization of the Fibonacci dimension using properties of an associated graph, and establish the Fibonacci dimension for certain families of graphs. From the algorithmic point of view, we prove that it is NP-complete to decide whether $\text{fdim}(G)$ equals the isometric dimension of G , and show that no algorithm to approximate $\text{fdim}(G)$ has approximation ratio below $741/740$, unless $P=NP$. We also give a $(3/2)$ -approximation algorithm for $\text{fdim}(G)$ in the general case and a $(1+\varepsilon)$ -approximation algorithm for simplex graphs.

1 Introduction

Hypercubes play a prominent role in metric graph theory as well as in several other areas such as parallel computing and coding theory. One of their central features is the ability to compute distances very efficiently because the distance between two vertices is simply the number of coordinates in which they differ; the same ability to compute distances may be transferred to any isometric subgraph of a hypercube. In this way partial cubes appear, a class of graphs intensively studied so far; see the books [12, 19, 32], the recent papers [3, 25, 43, 44], the recent (semi-)survey [42], and references therein. In particular we point out a recent fast recognition algorithm [17] and improvements in classification of cubic partial cubes [16, 37].

*Faculty of Mathematics and Physics, University of Ljubljana, Jadranska 19, 1000 Ljubljana, Slovenia; Institute of Mathematics, Physics and Mechanics, Jadranska 19, 1000 Ljubljana, Slovenia. E-mail: sergio.cabello@fmf.uni-lj.si.

†Computer Science Department, University of California, Irvine, CA 92697-3425, USA. Email: eppstein@uci.edu.

‡Faculty of Mathematics and Physics, University of Ljubljana, Jadranska 19, 1000 Ljubljana, Slovenia; Faculty of Natural Sciences and Mathematics, University of Maribor, Koroška 160, 2000 Maribor, Slovenia; Institute of Mathematics, Physics and Mechanics, Jadranska 19, 1000 Ljubljana, Slovenia. E-mail: sandi.klavzar@fmf.uni-lj.si.

The isometric dimension of a graph G is the smallest (and at the same time the largest) integer d such that G isometrically and irredundantly embeds into the d -dimensional cube. Clearly, the isometric dimension of G is finite if and only if G is a partial cube. This graph dimension is well-understood; for instance, it is equal to the number of steps in Chepoi's expansion procedure [9] and to the number of Θ -equivalence classes [13, 46] of a given graph. Two related graph dimensions need to be mentioned here since they are both defined on the basis of isometric embeddability into graph products. The lattice dimension of a graph is the smallest d such that the graph embeds isometrically into \mathbb{Z}^d (a Cartesian product of paths). Graphs with finite lattice dimension are precisely partial cubes, and the lattice dimension of any partial cube can be determined in polynomial time [15]. Another dimension is the strong isometric dimension—the smallest integer d such that a graph isometrically embeds into the strong product of d paths [20, 21]. In this case every graph has finite dimension, but this universality has a price: it is very difficult to compute the strong isometric dimension.

Fibonacci cubes are a subclass of the partial cubes that were first introduced by Hsu et al. in 1993 [29, 30], although closely related structures had been studied previously [4, 22, 28]. Several papers have investigated the structural properties of this class of graphs [11, 34, 38, 41]. In [8] it was shown that Fibonacci cubes are Θ -graceful while in [45] an efficient recognition algorithm is presented. The original motivation for introducing Fibonacci cubes was as an interconnection network for parallel computers; in that application, it is of interest to study the embeddability of other networks within Fibonacci cubes [10, 24].

In this paper we study this embedding question from the isometric point of view. We introduce the Fibonacci dimension of a graph as the smallest integer f such that the graph admits an isometric embedding into the f -dimensional Fibonacci cube; as we show, a graph G can be embedded in this way if and only if G is a partial cube. In the next section we give definitions, notions, and preliminary results needed in this paper. In Section 3 we give a combinatorial characterization of the Fibonacci dimension using properties of an associated graph. We provide upper and lower bounds showing that the Fibonacci dimension is always within a factor of two of the isometric dimension. We also provide tighter upper bounds based on a combination of the isometric and lattice dimensions, and we discuss the Fibonacci dimension of some particular classes of graphs. In Section 4 we show that computing the Fibonacci dimension is an NP-complete problem, provide inapproximability results, and give approximation algorithms.

2 Preliminaries

We will use the notation $[n] = \{1, \dots, n\}$. For any string u we will use $u^{(i)}$ to denote its i th coordinate. Unless otherwise specified, the distance in this paper is the usual shortest-path distance for unweighted graphs. A graph G is an *isometric subgraph* of another graph H if there is a way of placing the vertices of G in one-to-one correspondence with a subset of vertices of H , such that the distance in G equals the distance between corresponding vertices in H .

The vertex set of the d -cube Q_d consists of all d -tuples $u = u^{(1)}u^{(2)} \dots u^{(d)}$ with




Figure 1: The Fibonacci cube Γ_{10} .

$u^{(i)} \in \{0, 1\}$. Two vertices are adjacent if the corresponding tuples differ in precisely one position. Q_d is also called a *hypercube of dimension d* . Isometric subgraphs of hypercubes are *partial cubes*.

A *Fibonacci string* of length d is a binary string $u^{(1)}u^{(2)} \dots u^{(d)}$ with $u^{(i)} \cdot u^{(i+1)} = 0$ for $i \in [d - 1]$. In other words, a Fibonacci string is a binary string with no two consecutive ones. The set of Fibonacci strings of length d can be decomposed into two subsets: the subset of strings in which a starting 0 is followed by a Fibonacci string of length $d - 1$, and the subset of strings in which a starting 10 is followed by a Fibonacci string of length $d - 2$. For this reason the number of distinct Fibonacci strings of length d satisfies the Fibonacci recurrence and equals a Fibonacci number. The *Fibonacci cube* Γ_d , $d \geq 1$, is the subgraph of Q_d induced by the Fibonacci strings of length d . The Fibonacci cube may alternatively be defined as the graph of the distributive lattice of order-ideals of a fence poset [4, 22, 28] or as the simplex graph of the complement graph of a path graph. Since graphs of distributive lattices and simplex graphs are both instances of median graphs [2, 7], we have:

Theorem 2.1 ([34]) *Fibonacci cubes are median graphs. In particular, Fibonacci cubes are partial cubes and Γ_d isometrically embeds into Q_d .*

We will use the lattice \mathbb{Z}^d equipped with the L_1 -distance. Therefore, the distance between any two elements $(x_1, \dots, x_d), (y_1, \dots, y_d) \in \mathbb{Z}^d$ is given by $\sum_i |x_i - y_i|$. It will be convenient to visualize \mathbb{Z}^d as an infinite graph whose vertex set are elements of \mathbb{Z}^d and where two vertices are adjacent when they are at distance one; with this visualization, L_1 -distance coincides with the shortest path distance in the graph.

Let G be a connected graph. The *isometric dimension*, $\text{idim}(G)$, is the smallest integer k such that G admits an isometric embedding into Q_k . If there is no such k we set




Figure 2: Left: an isometric embedding of Γ_3 into Q_3 , with the complementary semicubes $W_{(3,1)}$ and $W_{(3,0)}$ shown as the shaded regions of the drawing. Right: the semicube graph of the embedding, consisting of a three-vertex path and three isolated vertices.

$\text{idim}(G) = \infty$. By definition, $\text{idim}(G) < \infty$ if and only if G is a partial cube. The *lattice dimension*, $\text{ldim}(G)$, is the smallest integer ℓ such that G admits an isometric embedding into \mathbb{Z}^ℓ . We similarly define the *Fibonacci dimension*, $\text{fdim}(G)$, as the smallest integer f such that G admits an isometric embedding into Γ_f , and set $\text{fdim}(G) = \infty$ if there is no such f .

Let $\beta : V(G) \rightarrow V(Q_k)$ be an isometric embedding. We will denote the i th coordinate of β with $\beta^{(i)}$. The embedding β is called *irredundant* if $\beta^{(i)}(V(G)) = \{0, 1\}$ for each $i \in [k]$. If an embedding is not irredundant, we may find an embedding onto a lower-dimensional hypercube by omitting the redundant coordinates. An isometric embedding $\beta : G \rightarrow Q_k$ is irredundant if and only if $k = \text{idim}(G)$ [46].

Let G be a partial cube with $\text{idim}(G) = k$ and assume that we are given an isometric embedding β of G into Q_k . Each pair $(i, \chi) \in [k] \times \{0, 1\}$ defines the *semicube* $W_{(i,\chi)} = \{u \in V(G) \mid \beta^{(i)}(u) = \chi\}$. For any $i \in [k]$, we refer to $W_{(i,0)}, W_{(i,1)}$ as a *complementary pair of semicubes*. This definition and notation seems to depend on the embedding β . However, any irredundant isometric embedding β' describes the same family of semicubes and pairs of complementary semicubes, possibly indexed in a different way.

For a partial cube G and a complementary pair of semicubes $W_{(i,0)}, W_{(i,1)}$, the set of edges with one endvertex in $W_{(i,0)}$ and the other in $W_{(i,1)}$ constitute a Θ -class of G . The Θ -classes of G form a partition of $E(G)$.

To determine the lattice dimension of a graph G , Eppstein [15] introduced the semicube graph $\text{Sc}(G)$ of a partial cube G as the graph with all the semicubes as nodes, semicubes $W_{(i,\chi)}$ and $W_{(i',\chi')}$ being adjacent if $W_{(i,\chi)} \cup W_{(i',\chi')} = V(G)$ and $W_{(i,\chi)} \cap W_{(i',\chi')} \neq \emptyset$. One can then show that the lattice dimension of G is equal to $\text{idim}(G) - |M|$, where M is a maximum matching of $\text{Sc}(G)$. See also [35] for further work on semicube graphs.

For any graph G , its *simplex graph* $\kappa(G)$ is defined as follows. There is a vertex u_K in $\kappa(G)$ for each clique K of G ; here we regard \emptyset , each vertex, and each edge of G as a clique. There is an edge between vertices u_K and $u_{K'}$ of $\kappa(G)$ whenever the cliques K and K' of G differ by exactly one vertex. In particular, there is an edge between u_\emptyset and u_a for each $a \in V(G)$, and there is an edge between u_a and u_{ab} for each edge $ab \in E(G)$. We will also




Figure 3: A graph G (left) with its corresponding simplex graph $\kappa(G)$ (center) and 2-simplex graph $\kappa_2(G)$.

use the *2-simplex graph* $\kappa_2(G)$ of a graph G , which is the subgraph of $\kappa(G)$ induced by the vertices u_K of $\kappa(G)$ corresponding to cliques K with at most 2 vertices. Alternatively, the 2-simplex graph of G may be formed by subdividing each edge of G and adding a new vertex u_\emptyset adjacent to each vertex that existed prior to the subdivision. An example is given in Figure 3. When G has no triangle, then $\kappa_2(G) = \kappa(G)$. 2-simplex graphs were used in [33] to establish a close connection between the recognition complexity of triangle-free graphs and of median graphs.

Finally, *computing an embedding* of G into Q_d (or Γ_d) means to attach to each vertex v of G a tuple $\beta(v)$ that is a vertex of Q_d such that β provides an isometric embedding.

3 Combinatorial aspects

3.1 The general case

Proposition 3.1 *Let G be a connected graph. Then $\text{fdim}(G)$ is finite if and only if $\text{idim}(G)$ is finite. Moreover,*

$$\text{idim}(G) \leq \text{fdim}(G) \leq 2 \text{idim}(G) - 1.$$

Proof. Let $f = \text{fdim}(G) < \infty$, so that G isometrically embeds into Γ_f . By Theorem 2.1, Γ_f isometrically embeds into Q_f , hence G isometrically embeds into Q_f . The Fibonacci strings with which Γ_f was derived may be used directly as the coordinates of an isometric embedding. Consequently $\text{idim}(G) \leq f = \text{fdim}(G)$.

Conversely, let $k = \text{idim}(G) < \infty$ and consider G isometrically embedded into Q_k . To each vertex $u = u^{(1)}u^{(2)} \dots u^{(k-1)}u^{(k)}$ of G (embedded into Q_k) assign the vertex $\tilde{u} = u^{(1)}0u^{(2)}0 \dots u^{(k-1)}0u^{(k)}$. Clearly, $\tilde{u}^{(i)} \cdot \tilde{u}^{(i+1)} = 0$ for any $i \in [2k - 2]$. Therefore, we can consider \tilde{u} as a vertex of Γ_{2k-1} . Let \tilde{G} be the subgraph of Γ_{2k-1} induced by the vertices \tilde{u} , $u \in V(G)$. Since Γ_{2k-1} is isometric in Q_{2k-1} (invoking Theorem 2.1 again), it readily follows that \tilde{G} is isometric in Γ_{2k-1} . We conclude that $\text{fdim}(G) \leq 2k - 1 = 2 \text{idim}(G) - 1$. \square

It is now clear that we only need to study the Fibonacci dimension of partial cubes. Using the lattice dimension $\text{idim}(G)$ we will further improve in Proposition 3.7 the upper bound on $\text{fdim}(G)$, and provide an alternative lower bound in Proposition 3.8.

Let G be a partial cube with $\text{idim}(G) = k$. In order to obtain an expression for $\text{fdim}(G)$ in terms of $\text{idim}(G)$ we construct the graph $X(G)$ as follows. The nodes of $X(G)$ are the semicubes $W_{(i,\chi)}$, $(i, \chi) \in [k] \times \{0, 1\}$, of G , semicubes $W_{(i,\chi)}$ and $W_{(j,\chi')}$ being adjacent if $i \neq j$ and $W_{(i,\chi)} \cap W_{(j,\chi')} = \emptyset$. Note that $X(G)$ is very close to the complement of the Eppstein's semicube graph $\text{Sc}(G)$.

A path P of $X(G)$ with the property that $|P \cap \{W_{(i,0)}, W_{(i,1)}\}| \leq 1$ for each complementary pair of semicubes $W_{(i,0)}, W_{(i,1)}$, will be called a *coordinating path*. A set of paths \mathcal{P} of $X(G)$ will be called a *system of coordinating paths* provided that any $P \in \mathcal{P}$ is a coordinating path and for each complementary pair of semicubes $W_{(i,0)}, W_{(i,1)}$ there is exactly one $P \in \mathcal{P}$ such that $|P \cap \{W_{(i,0)}, W_{(i,1)}\}| = 1$.

Lemma 3.2 *Let G be a partial cube and let \mathcal{P} be a system of coordinating paths of $X(G)$. Then there is an isometric embedding of G into $\Gamma_{f'}$, where $f' = \text{idim}(G) + |\mathcal{P}| - 1$.*

Proof. Let $k = \text{idim}(G)$, let $p = |\mathcal{P}|$, and let $\mathcal{P} = \{P_1, \dots, P_p\}$ be the given system of coordinating paths of $X(G)$. Let

$$\begin{aligned} P_1 &: W_{(a_1,\chi_1)} \rightarrow W_{(a_2,\chi_2)} \rightarrow \dots \rightarrow W_{(a_{i_1},\chi_{i_1})} \\ P_2 &: W_{(a_{i_1+1},\chi_{i_1+1})} \rightarrow W_{(a_{i_1+2},\chi_{i_1+2})} \rightarrow \dots \rightarrow W_{(a_{i_2},\chi_{i_2})} \\ &\vdots \\ P_p &: W_{(a_{i_{p-1}+1},\chi_{i_{p-1}+1})} \rightarrow W_{(a_{i_{p-1}+2},\chi_{i_{p-1}+2})} \rightarrow \dots \rightarrow W_{(a_{i_p},\chi_{i_p})}. \end{aligned}$$

As the paths meet exactly one of the complementary semicubes exactly once, $i_p = k$. More precisely, there is a bijection $\phi : \{a_1, a_2, \dots, a_{i_p}\} \rightarrow [k]$ such that if $\phi(a_i) = j$ then either $W_{(a_i,\chi_i)} = W_{(j,0)}$ or $W_{(a_i,\chi_i)} = W_{(j,1)}$ holds.

For any vertex u of G and any $i \in [k]$ set

$$\bar{u}^{(i)} = \begin{cases} 1 & \text{if } u \in W_{(a_i,\chi_i)}; \\ 0 & \text{otherwise.} \end{cases}$$

Assigning the k -tuple

$$u = \bar{u}^{(1)}\bar{u}^{(2)} \dots \bar{u}^{(i_p)}$$

to any vertex u of G yields the canonical isometric embedding of G into $Q_k = Q_{i_p}$. Now assign to u the following d -tuple:

$$\bar{u}^{(1)} \dots \bar{u}^{(i_1)} 0 \bar{u}^{(i_1+1)} \dots \bar{u}^{(i_2)} 0 \dots 0 \bar{u}^{(i_{p-1}+1)} \dots \bar{u}^{(i_p)}.$$

In this way, G is embedded into $Q_{f'}$, where $f' = k + p - 1$. Moreover, the embedding is clearly still isometric. Because $W_{(a_i,\chi_i)} \cap W_{(a_{i+1},\chi_{i+1})} = \emptyset$ provided that $W_{(a_i,\chi_i)}$ and $W_{(a_{i+1},\chi_{i+1})}$ are connected by an edge of some path P_j , the labeling of u is a Fibonacci string. Hence we have described an isometric embedding of G into $\Gamma_{f'}$. \square

Let $p(X(G))$ be the minimum size of a system of coordinating paths of $X(G)$. Then:

Theorem 3.3 *Let G be a partial cube. Then*

$$\text{fdim}(G) = \text{idim}(G) + p(X(G)) - 1.$$

Proof. Let $p = p(X(G))$, $k = \text{idim}(G)$, and $f = \text{fdim}(G)$. It readily follows from Lemma 3.2 and the definition of $p(X(G))$ that $f \leq k + p - 1$.

Consider now G isometrically embedded into Γ_f . For $u \in V(G)$ let $u^{(1)} \dots u^{(f)}$ be the embedded vertex. Let $1 \leq i_1 < i_2 < \dots < i_r \leq f$ be the indices for which all the vertices of G are labeled 0. That is, $u^{(i_j)} = 0$ holds for any $u \in V(G)$ and any i_j , $1 \leq j \leq r$. Then

$$\beta(u) = u^{(1)} \dots u^{(i_1-1)} u^{(i_1+1)} \dots u^{(i_2-1)} u^{(i_2+1)} \dots u^{(i_{r-1}-1)} u^{(i_{r-1}+1)} \dots u^{(i_r)}$$

is an isometric embedding into Q_{f-r} .

We next assert that for any coordinate i of the $(f-r)$ -tuples β , $Y_i = \{\beta^{(i)}(u) \mid u \in V(G)\} = \{0, 1\}$. Note first that $Y_i \neq \{1\}$ because otherwise the i th coordinate could be removed and hence we would isometrically embed G into Γ_{f-1} . On the other hand $Y_i \neq \{0\}$ since we have removed all such coordinates in the construction of β . Hence the assertion. However, this implies that f is an irredundant embedding and therefore

$$k = \text{idim}(G) = f - r.$$

For a given coordinate ℓ of β , set $W_\ell = \{u \in V(G) \mid \beta^{(\ell)}(u) = 1\}$. Then W_ℓ is a semicube. Moreover, because β is obtained from Fibonacci strings, the paths

$$\begin{aligned} W_1 &\rightarrow W_2 \rightarrow \dots \rightarrow W_{i_1-1}, \\ W_{i_1+1} &\rightarrow W_{i_1+2} \rightarrow \dots \rightarrow W_{i_2-1}, \\ &\vdots \\ W_{i_{r-1}+1} &\rightarrow W_{i_{r-1}+2} \rightarrow \dots \rightarrow W_{i_r}, \end{aligned}$$

form a system of coordinating paths with $r+1$ paths. Consequently, $r+1 \geq p$ and hence

$$k = f - r \leq f - p + 1.$$

We conclude that $f \geq k + p - 1$ which completes the proof. \square

Note that Proposition 3.1 also follows easily from Theorem 3.3.

3.2 Particular cases

It is interesting to ask which partial cubes have extremal Fibonacci dimension. Interestingly, it is difficult to characterize the partial cubes whose Fibonacci dimension is as small as possible compared to their isometric dimension; see Section 4.1. However, there is a neat characterization for the opposite case, the partial cubes whose Fibonacci dimension is as large as possible, which we provide next. Afterwards we establish the Fibonacci dimension of the Cartesian product of graphs and the Fibonacci dimension of trees.

The *crossing graph* $G^\#$ of a partial cube G has the Θ -classes of G as its nodes, where two nodes of $G^\#$ are joined by an edge whenever they cross as Θ -classes in G ; see [36]. More precisely, if $W_{(a,0)}, W_{(a,1)}$ and $W_{(b,0)}, W_{(b,1)}$ are pairs of complementary semicubes corresponding to Θ -classes E and F , then E and F cross if each semicube has a nonempty intersection with the semicubes from the other pair; that is, it holds that $W_{(a,0)} \cap W_{(b,0)}$, $W_{(a,0)} \cap W_{(b,1)}$, $W_{(a,1)} \cap W_{(b,0)}$, and $W_{(a,1)} \cap W_{(b,1)}$ are nonempty.

Corollary 3.4 *Let G be a partial cube with $\text{idim}(G) = k$. Then $\text{fdim}(G) = 2k - 1$ if and only if $G^\# = K_k$.*

Proof. By Theorem 3.3, $\text{fdim}(G) = 2k - 1$ if and only if $p(X(G)) = k$. This holds if and only if $X(G)$ has no edges which is in turn true if and only if for any distinct $i, j \in [k]$ the semicubes $W_{(i,0)}$ and $W_{(i,1)}$ nontrivially intersect $W_{(j,0)}$ and $W_{(j,1)}$. But this is true if and only if the corresponding Θ -classes cross. \square

A characterization of complete crossing graphs in terms of the expansion procedure is given in [36]: $G^\#$ is complete if and only if G can be obtained from K_1 by a sequence of all-color expansions. We also note that among median graphs only hypercubes have complete crossing graphs [39].

Corollary 3.5 *For any partial cubes G and H , $\text{fdim}(G \square H) = \text{fdim}(G) + \text{fdim}(H) + 1$.*

Proof. It is easy to infer that $X(G \square H)$ is isomorphic to $X(G) \cup X(H)$. Therefore, $p(X(G \square H)) = p(X(G)) + p(X(H))$. Since it is well-known that

$$\text{idim}(G \square H) = \text{idim}(G) + \text{idim}(H)$$

we have

$$\begin{aligned} \text{fdim}(G \square H) &= \text{idim}(G \square H) + p(X(G \square H)) - 1 \\ &= \text{idim}(G) + \text{idim}(H) + p(X(G) \cup X(H)) - 1 \\ &= \text{idim}(G) + \text{idim}(H) + p(X(G)) + p(X(H)) - 1 \\ &= (\text{idim}(G) + p(X(G)) - 1) + (\text{idim}(H) + p(X(H))) \\ &= \text{fdim}(G) + \text{fdim}(H) + 1, \end{aligned}$$

where for the first equality Theorem 3.3 is applied. \square

Corollary 3.6 *For any tree T , $\text{fdim}(T) = \text{idim}(T) = |E(T)|$.*

Proof. Let $n = |V(T)|$. It is well-known that $\text{idim}(T) = |E(T)| = n - 1$, and that each edge e of T constitutes a Θ -class [26] (cf. [32, Corollary 3.4.]). This means that each edge $e \in E(T)$ defines a pair of complementary semicubes: each semicube is the set of vertices in one of the two subtrees of $T - e$.




Figure 4: A tree with a longest path P marked with thicker edges. In the proof of Corollary 3.6, P_1 would be the path between r and v , P_2 would be the path between r and u , the labeling of the edges of P_1 corresponds to a proper enumeration, and the nodes marked with squares correspond to the semicube W_{e_5} .

Let P be a longest path in the tree T . We split P at a vertex r into two subpaths P_1, P_2 , such that P_1 and P_2 have the same length (if P has an even number of edges), or differ by one edge (if P has an odd number of edges). Without loss of generality, let us assume that P_1 is not strictly shorter than P_2 . Therefore $|E(P_1)| = |E(P_2)|$ if $|E(P)|$ is even and $|E(P_1)| = 1 + |E(P_2)|$ if $|E(P)|$ is odd. See Figure 4.

It may be convenient to visualize T as rooted at r . We further define the *level* of an edge xy of T as the minimum of $d_T(r, x), d_T(r, y)$. For any edge e of T , let W_e denote the subset of vertices in the subtree $T - e$ that does not contain the vertex r . As noted before, W_e is a semicube, and hence a node of $X(T)$, for any $e \in E(T)$.

Let e_1, e_2, \dots, e_{n-1} be an enumeration of the edges of T with the following properties: (a) any edge at level i is listed before any edge at level $i + 1$, and (b) the edge of P_1 at level i is the first edge at level i in the enumeration. Consider the sequence of semicubes $W_{e_1}, W_{e_2}, \dots, W_{e_{n-1}}$. If the edges e_i and e_{i+1} are at the same level, then clearly $W_{e_i} \cap W_{e_{i+1}} = \emptyset$. If e_i and e_{i+1} are not at the same level, then e_{i+1} must be an edge on P_1 while e_i cannot be an edge on P_1 . Therefore we also have $W_{e_i} \cap W_{e_{i+1}} = \emptyset$ in this case. This means that $W_{e_1} \rightarrow W_{e_2} \rightarrow \dots \rightarrow W_{e_{n-1}}$ is a path in $X(G)$, and furthermore forms a system of coordinating paths because it visits each complementary pair of semicubes exactly once. We conclude that $p(X(T)) = 1$, and thus $\text{fdim}(T) = \text{idim}(T)$ by Theorem 3.3. \square

3.3 Relation to lattice dimension

Using the lattice dimension $\text{l dim}(G)$, we can provide upper and lower bounds on the Fibonacci dimension $\text{fdim}(G)$. The first bound improves upon Proposition 3.1.

Proposition 3.7 *Let G be a partial cube. Then $\text{fdim} \leq \text{idim}(G) + \text{l dim}(G) - 1$.*

Proof. For any integers a, b with $a \leq b$, we use $P_{(a,b)}$ to denote the subgraph of \mathbb{Z}^1 induced by vertices $a, a + 1, \dots, b - 1, b$. Hence $P_{(a,b)}$ is a path on $b - a + 1$ vertices and $\text{fdim}(P_{(a,b)}) = b - a$ by Corollary 3.6.




Figure 5: A lattice embedding of Γ_4 .

Let $\ell = \text{ldim}(G)$ and consider an isometric embedding β of G into \mathbb{Z}^ℓ . For each coordinate $i \in [\ell]$, let $a_i = \min\{\beta^{(i)}(v) \mid v \in V(G)\}$ and let $b_i = \max\{\beta^{(i)}(v) \mid v \in V(G)\}$. It is shown in [15, Lemma 1] that $\sum_i (b_i - a_i)$ is precisely $\text{idim}(G)$. By the choice of a_i, b_i , the embedding β is also an isometric embedding of G into the Cartesian product $P_{(a_1, b_1)} \square P_{(a_2, b_2)} \square \cdots \square P_{(a_\ell, b_\ell)}$, and therefore

$$\text{fdim}(G) \leq \text{fdim}(P_{(a_1, b_1)} \square P_{(a_2, b_2)} \square \cdots \square P_{(a_\ell, b_\ell)}).$$

Since Corollary 3.5 implies

$$\begin{aligned} \text{fdim}(P_{(a_1, b_1)} \square P_{(a_2, b_2)} \square \cdots \square P_{(a_\ell, b_\ell)}) &= \left(\sum_{i=1}^{\ell} \text{fdim}(P_{(a_i, b_i)}) \right) + (\ell - 1) \\ &= \left(\sum_{i=1}^{\ell} (b_i - a_i) \right) + (\ell - 1) \\ &= \text{idim}(G) + \text{ldim}(G) - 1, \end{aligned}$$

we conclude that $\text{fdim}(G) \leq \text{idim}(G) + \text{ldim}(G) - 1$. □

Proposition 3.8 *Let G be a partial cube. Then $\text{ldim}(G) \leq \lceil \text{fdim}(G)/2 \rceil$.*

Proof. Consider the Fibonacci cube Γ_f for $f \geq 3$, and let u^* denote the last $f - 2$ entries of each tuple $u \in V(\Gamma_f)$. Define an embedding β of Γ_f into $\mathbb{Z}^1 \square \Gamma_{f-2}$ by

$$\beta(u) = \begin{cases} (0, u^*) & \text{if } u = 01u^*; \\ (1, u^*) & \text{if } u = 00u^*; \\ (2, u^*) & \text{if } u = 10u^*. \end{cases}$$

It is straightforward to see that β is an isometric embedding. Using induction on the Fibonacci dimension, with base cases $\text{ldim}(\Gamma_1) = \text{ldim}(\Gamma_2) = 1$, we obtain

$$\text{ldim}(\Gamma_f) \leq 1 + \text{ldim}(\Gamma_{f-2}) \leq 1 + \lceil (f - 2)/2 \rceil = \lceil f/2 \rceil.$$

If a partial cube isometrically embeds into Γ_f , we then have $\text{ldim}(G) \leq \text{ldim}(\Gamma_f) \leq \lceil f/2 \rceil$, and the result follows. \square

For graphs with low lattice dimension, we may determine the Fibonacci dimension exactly:

Proposition 3.9 *Suppose that $\text{ldim}(G) = 2$. Then $\text{fdim}(G) = \text{idim}(G) + i$, where $i = 1$ when G is isomorphic to the Cartesian product of two paths and $i = 0$ otherwise.*

Proof. When G is isomorphic to the product of two paths, the result follows from Corollaries 3.5 and 3.6. Otherwise, G is a proper subgraph of $P_1 \square P_2$, where P_1 and P_2 are two paths with total length equal to $\text{idim}(G)$. Among the four corner vertices of $P_1 \square P_2$ determined by pairs of endpoints of the two paths, at least one corner must be absent in G if G is to be a proper subgraph of the product of paths; we may assume without loss of generality that this missing corner corresponds to the last vertex of P_1 and the first vertex of P_2 .

We may embed P_1 isometrically into a Fibonacci cube (following Corollary 3.6) using the coordinates

$$101010\dots, 001010\dots, 000010\dots, \dots, \dots 010000, \dots 010100, \dots 010101$$

when P_1 has even length, or with a similar pattern when P_1 has odd length. That is, we start with an alternating sequence of zeros and ones, remove the ones one at a time, and then add ones one at a time to end with the opposite alternating sequence of ones and zeros. This pattern can be chosen in such a way that the final coordinate is zero for all vertices of P_1 except for its the last vertex. Similarly, we may embed P_2 isometrically into a set of Fibonacci strings in such a way that the initial coordinate is zero except in the first vertex of P_2 . Concatenating these two representations of positions in P_1 and P_2 produces an irredundant isometric embedding of G into a Fibonacci cube. \square

4 Algorithmic aspects

4.1 Bad news

We show that it is NP-complete to decide whether the isometric and Fibonacci dimensions of a given graph are the same. Furthermore, we show that it is NP-hard to approximate the Fibonacci dimension within $(741/740) - \epsilon$, for any constant $\epsilon > 0$.

Let G be a graph with n vertices. We assume for simplicity that $V(G) = [n]$, and use a, b to refer to the vertices of G . Let \bar{G} be the complementary graph of G .

Lemma 4.1 *Let H be either the simplex graph $\kappa(G)$ or the 2-simplex graph $\kappa_2(G)$. Then H is a partial cube with $\text{idim}(H) = n$.*

Proof. Consider the embedding $\beta : H \rightarrow Q_n$ given as follows:




Figure 6: Isometric embedding of $\kappa(G)$ for the graph G of Figure 3, left.

- for u_\emptyset we set $\beta(u_\emptyset) = u^{(1)} \dots u^{(n)}$ with $u^{(i)} = 0$ for all $i \in [n]$;
- for each $a \in [n]$ we set $\beta(u_a) = u^{(1)} \dots u^{(n)}$ with $u^{(a)} = 1$ and $u^{(i)} = 0$ for all $i \in [n] \setminus \{a\}$;
- for each node u_K of H , we set $\beta(u_K) = \sum_{a \in K} \beta(u_a)$.

See Figure 6 for an example when $H = \kappa(G)$. It is straightforward to see that β is an isometric embedding of H into Q_n , and hence H is a partial cube. Moreover, β is irredundant: $\beta^{(i)}(u_a)$ is nonzero if and only if $a \neq i$. Since there is an irredundant isometric embedding of a graph H into Q_k if and only if $\text{idim}(H) = k$, it follows that $\text{idim}(H) = n$. \square

In fact, stronger result that Lemma 4.1 was proved in [2] for $\kappa(G)$ and in [33] for $\kappa_2(G)$: H is a median graph.

Lemma 4.2 *Let H be either the simplex graph $\kappa(G)$ or the 2-simplex graph $\kappa_2(G)$. There is a set \mathbb{W} of semicubes of H with the following properties:*

- Each node in \mathbb{W} has degree zero in $X(H)$.
- Each pair of complementary semicubes of H has a node in \mathbb{W} .
- $X(H) - \mathbb{W}$ is isomorphic to \bar{G} .

Proof. We will use the isometric embedding β given in the proof of Lemma 4.1. For any $a \in [n]$ we then have the semicubes

$$\begin{aligned} W_{(a,0)} &= \{u_K \in V(H) \mid \beta^{(a)}(u_K) = 0\} \\ &= \{u_K \in V(H) \mid a \text{ is not a vertex in } K\}, \end{aligned}$$

and

$$\begin{aligned} W_{(a,1)} &= \{u_K \in V(H) \mid \beta^{(a)}(u_K) = 1\} \\ &= \{u_K \in V(H) \mid a \text{ is a vertex in } K\}. \end{aligned}$$

Let us now consider the graph $X(H)$. See Figure 7 for an example. The node set of $X(H)$ is $W_{(a,\chi)}$, $(a, \chi) \in [n] \times \{0, 1\}$. For the edge set, we have the following properties:

- There is no edge between $W_{(a,0)}$ and $W_{(b,0)}$ because $u_\emptyset \in W_{(a,0)} \cap W_{(b,0)}$.
- There is no edge between $W_{(a,0)}$ and $W_{(b,1)}$ because $u_b \in W_{(a,0)} \cap W_{(b,1)}$.
- There is an edge between $W_{(a,1)}$ and $W_{(b,1)}$ if and only if $ab \notin E(G)$. Indeed, there is a vertex u_K of H in $W_{(a,1)} \cap W_{(b,1)}$ if and only if a and b are vertices in the clique K , which happens precisely when ab is an edge of G . Therefore $W_{(a,1)} \cap W_{(b,1)} \neq \emptyset$ if and only if $ab \in E(G)$.

It follows that each node $W_{(a,0)}$, $a \in [n]$, has degree zero in $X(H)$. Therefore, the subfamily of nodes $\mathbb{W} = \{W_{(a,0)} \mid a \in [n]\}$ of $X(H)$ satisfies properties (a) and (b) in the lemma. The graph $X(G) - \mathbb{W}$ contains only the nodes $W_{(a,1)}$, $a \in [n]$. Since there is an edge between $W_{(a,1)}$ and $W_{(b,1)}$ if and only if $ab \notin E(G)$, the mapping $a \mapsto W_{(a,1)}$ is an isomorphism between \bar{G} , and property (c) follows. \square

Let (1, 2)-TSP denote the (metric) symmetric Traveling Salesman Problem in which all distances are either 1 or 2. The (1, 2)-TSP problem is NP-hard. Furthermore, Engebretsen and Karpinski [14] have shown that it is NP-hard to approximate the (1, 2)-TSP within $(741/740) - \varepsilon$ for every constant $\varepsilon > 0$. On the positive side, Berman and Karpinski [6] have given an $(8/7)$ -approximation algorithm for (1, 2)-TSP.

Any graph G naturally defines an instance I_G of (1, 2)-TSP, where the points of the metric space are the vertices of G , and the distance between two points is 1 if there is an edge between them in G , and 2 otherwise. Let $\ell(I_G)$ denote the length of the optimal tour for an instance I_G of (1, 2)-TSP.

For later use, it will be convenient to exchange now the roles of G and its complementary graph \bar{G} .

Lemma 4.3 *Let H be either the simplex graph $\kappa(\bar{G})$ or the 2-simplex graph $\kappa_2(\bar{G})$. The graph G has a Hamiltonian path if and only if $\text{fdim}(H) = n$. If G does not have a Hamiltonian path, then $\text{fdim}(H) = \ell(I_G) - 1$.*

Proof. Consider the set of nodes \mathbb{W} in $X(H)$ given by Lemma 4.2. Since each node of \mathbb{W} has degree zero in $X(H)$ and \mathbb{W} contains one semicube from each pair of complementary semicubes of H , we can just disregard the nodes \mathbb{W} for finding the value $p(X(H))$. When we disregard the nodes \mathbb{W} , we obtain $X(H) - \mathbb{W}$, which is isomorphic to G because of property (c) in Lemma 4.2. (Recall we exchanged the roles of G and \bar{G} .) It follows that $p(X(H))$ is the minimum number of vertex-disjoint paths that are needed to cover each vertex of G .




Figure 7: The graph $X(\kappa_2(G))$ for the graph G of Figure 3, left.

The graph G has a Hamiltonian path if and only if $p(X(H)) = 1$. Using Theorem 3.3 and Lemma 4.1, this is equivalent to

$$\text{fdim}(H) = \text{idim}(H) + p(X(H)) - 1 = \text{idim}(H) = n.$$

If G does not have a Hamiltonian path, then $\ell(I_G)$ is $|V(G)| = n$ plus the minimum number of vertex-disjoint paths that are needed to cover each vertex of G . Thus $\ell(I_G) = n + p(X(H))$. Using Theorem 3.3 and Lemma 4.1 we conclude that

$$\text{fdim}(H) = \text{idim}(H) + p(X(H)) - 1 = n + p(X(H)) - 1 = \ell(I_G) - 1.$$

□

We next show that computing the Fibonacci dimension, or even to approximate it, is NP-hard.

Theorem 4.4 *It is NP-complete to decide whether $\text{idim}(H) = \text{fdim}(H)$ for a given graph H .*

Proof. Note that $\text{idim}(H)$ can be computed in polynomial time [1, 17, 31]. Therefore, an explicit isometric embedding of H into $\Gamma_{\text{idim}(H)}$ would be enough to check in polynomial time that $\text{idim}(H) = \text{fdim}(H)$. It follows that the problem is in the class NP.

To show hardness, consider the graph $H = \kappa_2(\bar{G})$. It is clear that H can be constructed in polynomial time for any given graph G ; this is not necessarily true for $\kappa(\bar{G})$ if \bar{G} has large cliques. Lemma 4.3 implies that $\text{idim}(H) = \text{fdim}(H)$ if and only if G has a Hamiltonian path. Since deciding whether a graph has a Hamiltonian path is NP-complete [23], it is NP-hard to decide whether $\text{idim}(H) = \text{fdim}(H)$. \square

Theorem 4.5 *It is NP-hard to approximate the Fibonacci dimension of a graph within $(741/740) - \varepsilon$ for every constant $\varepsilon > 0$.*

Proof. Assume that there is a constant $\varepsilon > 0$ and a polynomial time algorithm APPROXFIB that, for any input graph H , computes a value $f'(H)$ such that

$$\text{fdim}(H) \leq f'(H) \leq \left(\frac{741}{740} - \varepsilon\right) \text{fdim}(H).$$

Given any graph G with n vertices, we can apply algorithm APPROXFIB to the graph $H = \kappa_2(\bar{G})$ to obtain a value f' that satisfies

$$\text{fdim}(H) \leq f' \leq \left(\frac{741}{740} - \varepsilon\right) \text{fdim}(H). \quad (1)$$

Consider the value $\ell' = f' + 1$ as an approximation to $\ell(I_G)$.

From Lemma 4.3 it follows that

$$\ell(I_G) - 1 \leq \text{fdim}(H) \leq \ell(I_G). \quad (2)$$

(There is the special case when G has a Hamiltonian cycle because then $\ell(I_G) = n = \text{fdim}(H)$.) Combining inequalities (1) and (2) we obtain

$$\ell(I_G) \leq \text{fdim}(H) + 1 \leq f' + 1 = \ell',$$

and

$$\begin{aligned} \ell' &= f' + 1 \\ &\leq \left(\frac{741}{740} - \varepsilon\right) \text{fdim}(H) + 1 \\ &\leq \left(\frac{741}{740} - \varepsilon\right) \ell(I_G) + 1 \\ &= \left(\frac{741}{740} + \frac{1}{\ell(I_G)} - \varepsilon\right) \ell(I_G) \\ &\leq \left(\frac{741}{740} + \frac{1}{n} - \varepsilon\right) \ell(I_G). \end{aligned}$$

Since $2/\varepsilon$ is a constant, we may assume that G has more than $2/\varepsilon$ vertices, and thus

$$\ell(I_G) \leq \ell' \leq \left(\frac{741}{740} + \frac{1}{n} - \varepsilon\right) \ell(I_G) \leq \left(\frac{741}{740} - \frac{\varepsilon}{2}\right) \ell(I_G).$$

We then conclude that ℓ' can be computed in polynomial time and approximates the value $\ell(I_G)$ within $(741/740) - (\varepsilon/2)$. However, Engebretsen and Karpinski [14] have shown that it is NP-hard to approximate the (1, 2)-TSP within $(741/740) - \delta$ for every constant $\delta > 0$. Therefore, it is also NP-hard to approximate the Fibonacci dimension of a graph within $(741/740) - \varepsilon$ for every constant $\varepsilon > 0$. \square

4.2 Good news

We first provide an exact algorithm to compute $\text{fdim}(G)$ whose running time is exponential in $\text{idim}(G)$. We then provide a $(3/2)$ -approximation algorithm for arbitrary graphs, and better approximation algorithms specialized to simplex graphs.

We assume that our input is a partial cube G with n vertices and also that we are given an embedding β of G into Q_k , where $k = \text{idim}(G)$. Such embedding can be constructed in $O(n^2)$ time [17]¹. We first describe how to construct $X(G)$ and then give an algorithmic counterpart of Lemma 3.2.

Lemma 4.6 *The graph $X(G)$ can be computed in $O(k^2n)$ time.*

Proof. Each semicube $W_{(i,\chi)}$ is identified by a pair $(i, \chi) \in [k] \times \{0, 1\}$. We first construct the complete graph on the node set $\{(i, \chi) \in [k] \times \{0, 1\}\}$ and then, for each vertex $v \in V(G)$, the edges $(i, \beta^{(i)}(v))(j, \beta^{(j)}(v))$ are removed for all distinct $i, j \in [k]$. The resulting graph is (isomorphic to) $X(G)$. Using any standard data structure for graphs, each edge can be deleted in constant time. For each of the n vertices of G , we thus spend $O(k^2)$ time, for a total of $O(k^2n)$ time. \square

Lemma 4.7 *Assume we are given a system of p coordinating paths of $X(G)$. Then we can compute in $O(kn)$ time an isometric embedding of G into $\Gamma_{f'}$, where $f' = k + p - 1$.*

Proof. The proof given in Lemma 3.2 is constructive and can be implemented in $O(n(k + p)) = O(kn)$ time. \square

From $X(G)$ it is possible to compute $p(X(G))$, and thus $\text{fdim}(G)$, in roughly $O(k!)$ time by trying all permutations of the indices $[k]$ to obtain systems of coordinating paths of $X(G)$. We next improve this to a dependency that is exponential in k .

Proposition 4.8 *Given a partial cube G with n vertices and an isometric embedding $G \rightarrow Q_k$, where $k = \text{idim}(G)$, we can compute in $O(2^k k^2 + k^2 n)$ time an isometric embedding of G in Γ_f , where $f = \text{fdim}(G)$.*

Proof. Firstly, we construct the graph $X(G)$ using Lemma 4.6 in $O(k^2n)$ time. Secondly, we find in $O(2^k k^2)$ time a system of coordinating paths of $X(G)$ with minimum size using dynamic programming, as described below. Finally, we use Lemma 4.7 to construct the embedding in $O(kn)$ time. We only have to describe the second step.

We compute the value $p(X(G))$ using dynamic programming across subsets of pairs of complementary semicubes, as follows. Our approach is essentially the same as a standard one for TSP [5, 27]. For any subset of indices $I \subseteq [k]$, let $X_I(G)$ denote the subgraph of $X(G)$ induced by nodes $W_{(i,0)}, W_{(i,1)}$, $i \in I$. For any triple $(I, j, \chi) \in 2^{[n]} \times [n] \times \{0, 1\}$ with $j \in I$, let $\pi(I, j, \chi)$ denote the minimum number of paths in a system of coordinating

¹The algorithm in [17] assumes the word-RAM model of computation. Without bit-manipulation, there are algorithms [1, 31] taking $O(n^2 \log n)$ time.

paths for $X_I(G)$, with the property that $W_{(j,\chi)}$ is an end-node of some coordinating path. That is, $\pi(I, j, \chi)$ is the minimum number of paths in $X(G)$ that visit each pair of complementary semicubes $W_{(i,0)}, W_{(i,1)}$, $i \in I$, exactly once, it has one path ending at node $W_{(j,\chi)}$, and does not visit any semicube $W_{(i,*)}$ for $i \notin I$.

It is clear that for any $j \in [k]$ it holds $\pi(\{j\}, j, 0) = \pi(\{j\}, j, 1) = 1$. For subsets I with more than one index there are two cases to distinguish, depending on whether the paths defining $\pi(I, j, \chi)$ have $W_{(j,\chi)}$ as an isolated node or not. Therefore it holds

$$\pi(I, j, \chi) = \min \begin{cases} 1 + \min_{(j', \chi') \in (I \setminus \{j\}) \times \{0,1\}} \pi(I \setminus \{j\}, j', \chi') \\ (j', \chi') \in (I \setminus \{j\}) \times \{0,1\} \\ \text{s.t.} \\ W_{(j,\chi)} W_{(j',\chi')} \in E(X(G)) \end{cases}$$

Note that $\pi(I, j, \chi)$ only depends on values $\pi(I', j', \chi')$ with $|I'| = |I| - 1$. Therefore, we can compute all values $\pi(I, j, \chi)$ by considering them for increasing values of $|I|$, at a cost of $O(k)$ per value. Since there are at most $|2^{[k]}| \cdot |[k]| \cdot 2 = O(2^k k)$ tuples (I, j, χ) to consider, we can compute in $O(2^k k^2)$ time the values $\pi(I, j, \chi)$ for all tuples $(I, j, \chi) \in 2^{[n]} \times [k] \times \{0,1\}$ with $j \in I$. Finally, it holds that

$$p(X(G)) = \min_{(j,\chi) \in [n] \times \{0,1\}} \pi([n], j, \chi),$$

and hence we can recover $p(X(G))$ in $O(k)$ time. To obtain the actual system of coordinating paths, we only need to augment each entry (I, j, χ) with a list of the paths that define $\pi(I, j, \chi)$. \square

We next move onto approximation algorithms. Note that the value $\text{idim}(G) + \text{ldim}(G) - 1$ is a $(3/2)$ -approximation to the value $\text{fdim}(G)$ because of Propositions 3.7 and 3.8. Moreover, the proofs of Propositions 3.7 and 3.8 are constructive, and therefore we can use isometric embeddings of G into $Q_{\text{idim}(G)}$ and into $\mathbb{Z}^{\text{ldim}(G)}$ to construct an isometric embedding of G into $\Gamma_{\text{idim}(G) + \text{ldim}(G) - 1}$. Since an isometric embedding of G into $\mathbb{Z}^{\text{ldim}(G)}$ can be computed in polynomial time [16], we can then compute in polynomial time an embedding of G into $\Gamma_{f'}$ for $f' \leq (3/2)\text{fdim}(G)$. We next give an alternative algorithm with the same performance (time and approximation factor) that does not make the detour through finding an isometric embedding into $\mathbb{Z}^{\text{ldim}(G)}$.

Theorem 4.9 *Given a partial cube G with n vertices and an isometric embedding $G \rightarrow Q_k$, where $k = \text{idim}(G)$, we can compute in $O(k^2 n)$ time an isometric embedding of G into $\Gamma_{f'}$, where $f' \leq (3/2)\text{fdim}(G)$.*

Proof. We first describe the algorithm, then derive its running time, and finally discuss the bound on the dimension of the computed embedding.

The algorithm is as follows. Firstly, construct the graph $X(G)$. Secondly, construct the graph $Y(G)$ obtained from $X(G)$ by identifying each pair of complementary semicubes into a single node. Hence, $Y(G)$ has k nodes. Thirdly, construct a matching M_Y in $Y(G)$ of maximum cardinality. Let M_X denote the matching in $X(G)$ that corresponds to M_Y . We can then regard each edge of M_X as a path in $X(G)$ that passes through two nodes. Let $P_1, \dots, P_{|M_X|}$ denote these paths. There are precisely $k - 2|M_X|$ pairs of complementary semicubes that are not adjacent to M_X . For each of those pairs, we make a path consisting of a single semicube of the pair. This gives a family of $|M_X| + (k - 2|M_X|) = k - |M_X|$ paths that form a system of coordinating paths of $X(G)$. Finally, we compute the embedding into $\Gamma_{f'}$ given by Lemma 4.7, where

$$f' = k + (k - |M_X|) - 1 = 2k - |M_X| - 1. \quad (3)$$

This finishes the description of the algorithm. Clearly, this algorithm computes a valid embedding of G into $\Gamma_{f'}$.

To derive its running time, note that $X(G)$ is constructed in $O(k^2n)$ time using Lemma 4.6. We can then construct $Y(G)$ by identifying the nodes $(i, 0)$ and $(i, 1)$ of $X(G)$ for all $i \in [d]$. Finding a maximum matching M_Y in $Y(G)$ takes $O(k^{5/2})$ time using [40] because $Y(G)$ has k nodes and $O(k^2)$ edges. From M_Y we can recover the matching M_X in $X(G)$, and construct the embedding using Lemma 4.7 in $O(kn)$ time. We conclude that the algorithm takes $O(k^{5/2} + k^2n)$ time, which is $O(k^2n)$ because $k \leq n$.

It remains to bound f' . Let $p = p(X(G))$ and let P_1, P_2, \dots, P_p be a system of coordinating paths of $X(G)$. Consider these paths in $Y(G)$, and let E_P denote the set of edges appearing in P_1, P_2, \dots, P_p . It holds that

$$p = k - |E_P|. \quad (4)$$

Taking each other edge in the path P_i , we see that P_i contains a matching with $\lceil |E(P_i)|/2 \rceil$ edges. Thus the paths P_1, \dots, P_p in $Y(G)$ contain a matching with at least $|E_P|/2$ edges. We conclude that

$$|M_X| = |M_Y| \geq \frac{|E_P|}{2}. \quad (5)$$

Combining equations (3)-(5) we obtain

$$\begin{aligned} f' &= 2k - |M_X| - 1 \\ &\leq 2k - \frac{|E_P|}{2} - 1 \\ &= 2k - \frac{k - p}{2} - 1 \\ &= \frac{3}{2}(k + p - 1) - p + \frac{1}{2} \\ &\leq \frac{3}{2}(\text{fdim}(G)), \end{aligned}$$

where in the last step we have used Theorem 3.3. □

We now turn our attention to approximation algorithms for simplex graphs.

Theorem 4.10 *Let $\varepsilon \in (0, 1)$ be a constant. Given a simplex graph H with n vertices and an isometric embedding $H \rightarrow Q_k$, where $k = \text{idim}(H)$, we can compute in $O(n^{2+1/\varepsilon})$ time an isometric embedding of H in $\Gamma_{f'}$, where $f' \leq (1 + \varepsilon)\text{fdim}(H)$.*

Proof. Since H is a simplex graph, then $H = \kappa(G)$ for some graph G . It holds $k = \text{idim}(H) = |V(G)|$ because of Lemma 4.1. The graph $X(H)$ can be constructed in $O(k^2n)$ time by Lemma 4.6. We can construct the set of nodes \mathbb{W} of Lemma 4.2 by placing in \mathbb{W} , for each $i \in [k]$, either $W_{(i,0)}$ or $W_{(i,1)}$, whichever has degree zero in $X(G)$. By property (c) of Lemma 4.2 it holds that $X(H) - \mathbb{W}$ and \bar{G} are isomorphic.

We construct a system of coordinating paths of $X(H)$ with the following greedy procedure. We start with a system of k coordinating paths $\mathcal{P} = \{P_1, \dots, P_k\}$, where each path consists of a node from $V(X(H)) \setminus \mathbb{W}$. Then, we repeat the following step as many times as possible: if there are paths $P, P' \in \mathcal{P}$ that can be joined by adding an edge between two of its extreme nodes, we do so, and replace in \mathcal{P} the paths P, P' by the new path. This step may be repeated at most $k - 1$ times, in which case we end up with a single path in \mathcal{P} . Since each repetition of the step takes $O(k^2)$ time, the whole procedure needs $O(k^3)$ time. Let \mathcal{P} denote the resulting system of coordinating paths.

We distinguish two cases depending on the size of \mathcal{P} . If $|\mathcal{P}| \leq \varepsilon k$, then we can use Lemma 4.7 with \mathcal{P} to construct an embedding of H into $\Gamma_{f'}$, where $f' \leq k + |\mathcal{P}| - 1 \leq (1 + \varepsilon)k \leq (1 + \varepsilon)\text{fdim}(H)$. If $|\mathcal{P}| > \varepsilon k$, then by selecting an extreme node in each path of \mathcal{P} we obtain in $X(G) - \mathbb{W}$ an independent set of nodes of cardinality at least εk . Since $X(G) - \mathbb{W}$ and \bar{G} are isomorphic, this means that \bar{G} has an independent set with at least εk vertices, and hence G has a clique K with at least εk vertices. Since the input graph H is the simplex graph of G , we conclude that H has at least $2^{|K|} \geq 2^{\varepsilon k}$ vertices; that is $n \geq 2^{\varepsilon k}$. Using Proposition 4.8, we can then compute an isometric embedding of H into $\Gamma_{\text{fdim}(H)}$ in

$$O(2^k k^2 + k^2 n) = O\left((2^{\varepsilon k})^{1/\varepsilon} n^2 + n^3\right) = O\left(n^{1/\varepsilon} n^2 + n^3\right) = O\left(n^{2+1/\varepsilon}\right)$$

time. In either case, we obtain in $O(n^{2+1/\varepsilon})$ time an isometric embedding of H into $\Gamma_{f'}$, where $f' \leq (1 + \varepsilon)\text{fdim}(H)$. \square

The proof of this result is closely related to (and was the inspiration for) related results on approximation algorithms that find a solution to one of two given graph optimization problems (such as, in this case, the maximum clique or the minimum 1-2 TSP) with a better approximation ratio than either problem could be approximated alone; see [18].

Note that the last result takes time polynomial in the size of the given simplex graph. However, we could consider that the simplex graph $\kappa(G)$ is described by G , and give G as the input. The following result approximates the Fibonacci dimension of $\kappa(G)$ for a given graph G in time that is polynomial in the size of G . Note that we cannot compute an explicit isometric embedding of $\kappa(G)$ in polynomial time of because the size of $\kappa(G)$ may be exponential in G .

Theorem 4.11 *Let $\varepsilon \in (0, 1)$ be a constant. Given a graph G with n vertices, we can compute in polynomial time a value f' such that $\text{fdim}(\kappa(G)) \leq f' \leq (\frac{8}{7} + \varepsilon)\text{fdim}(\kappa(G))$.*

Proof. If G has less than $1/\varepsilon$ vertices, which is a constant, we can then compute $\text{fdim}(\kappa(G))$ exactly in constant time. Let us assume henceforth that G has at least $1/\varepsilon$ vertices.

From Lemma 4.3 it follows that $\text{fdim}(\kappa(G)) = \ell(I_{\bar{G}}) - 1$ if \bar{G} does not have a Hamiltonian cycle, and $\text{fdim}(\kappa(G)) = n = \ell(I_{\bar{G}})$ if \bar{G} has a Hamiltonian cycle. Berman and Karpinski [6] describe a polynomial-time $(8/7)$ -approximation algorithm for $(1, 2)$ -TSP. Let ℓ' be the $(8/7)$ -approximation to the value $\ell(I_{\bar{G}})$ returned by their algorithm. If $\ell' = n$, then \bar{G} has a Hamiltonian cycle, and $\text{fdim}(\kappa(G)) = n$. If $\ell' > n$, consider $f' = \ell' - 1$ as an approximation to $\text{fdim}(\kappa(G))$. For f' we have the upper bound

$$f' = \ell' - 1 \leq \frac{8}{7}\ell(I_{\bar{G}}) - 1 \leq \frac{8}{7}(\ell(I_{\bar{G}}) - 1) + \frac{1}{7} \leq \frac{8}{7}\text{fdim}(\kappa(G)) + \frac{1}{7},$$

which using that $n\varepsilon \geq 1$ and $n \leq \text{fdim}(\kappa(G))$ leads to

$$f' \leq \frac{8}{7}\text{fdim}(\kappa(G)) + \frac{\varepsilon n}{7} \leq \frac{8}{7}\text{fdim}(\kappa(G)) + \varepsilon\text{fdim}(\kappa(G)) \leq (\frac{8}{7} + \varepsilon)\text{fdim}(\kappa(G)).$$

On the other hand, for f' we also have the lower bound

$$\text{fdim}(\kappa(G)) = \max\{n, \ell(I_{\bar{G}}) - 1\} \leq \max\{n, \ell' - 1\} = \ell' - 1 = f'.$$

The result follows. □

By Lemma 4.3 it holds that $\text{fdim}(\kappa(G)) = \text{fdim}(\kappa_2(G))$, and hence the same result holds for the 2-simplex graph $\kappa_2(G)$.

Acknowledgments

Work of D. Eppstein was supported in part by NSF grant 0830403 and by the Office of Naval Research under grant N00014-08-1-1015. Work of S. Cabello and S. Klavžar was supported in part by the Slovenian Research Agency, program P1-0297. We also would like to thank David Johnson and Thore Husfeldt for clarifying the history of using dynamic programming for TSP.

References

- [1] F. Aurenhammer and J. Hagauer. Recognizing binary Hamming graphs in $O(n^2 \log n)$ time. *Math. Systems Theory*, 28(5):387–395, 1995.
- [2] H.-J. Bandelt and M. van de Vel. Embedding topological median algebras in products of dendrons. *Proc. London Math. Soc. (3)*, 58(3):439–453, 1989.
- [3] L. Beaudou, S. Gravier, and K. Meslem. Isometric embeddings of subdivided complete graphs in the hypercube. *SIAM J. Discrete Math.*, 22(5):1226–1238, 2008.
- [4] I. Beck. Partial orders and the Fibonacci numbers. *Fibonacci Quart.*, 28(2):172–174, 1990.

- [5] R. Bellman. Combinatorial processes and dynamic programming. In *Proc. Sympos. Appl. Math., Vol. 10*, pages 217–249. American Mathematical Society, Providence, R.I., 1960.
- [6] P. Berman and M. Karpinski. 8/7-approximation algorithm for (1,2)-TSP. In *Proc. 17th ACM-SIAM Symp. Discrete Algorithms, SODA '06*, pages 641–648, New York, NY, USA, 2006. ACM.
- [7] G. Birkhoff and S. A. Kiss. A ternary operation in distributive lattices. *Bull. Amer. Math. Soc.*, 53:749–752, 1947.
- [8] B. Brešar and S. Klavžar. Θ -graceful labelings of partial cubes. *Discrete Math.*, 306(13):1264–1271, 2006.
- [9] V. Chepoi. Isometric subgraphs of Hamming graphs and d -convexity. *Kibernetika (Kiev)*, 1:6–9, 15, 133, 1988.
- [10] B. Cong, S. Q. Zheng, and S. Sharma. On simulations of linear arrays, rings and 2D meshes on Fibonacci cube networks. In *Proc. 7th Int. Parallel Processing Symposium*, pages 748–751, 1993.
- [11] E. Dedó, D. Torri, and N. Zagaglia Salvi. The observability of the Fibonacci and the Lucas cubes. *Discrete Math.*, 255(1-3):55–63, 2002.
- [12] M. M. Deza and M. Laurent. *Geometry of Cuts and Metrics*, volume 15 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, 1997.
- [13] D. Ž. Djoković. Distance-preserving subgraphs of hypercubes. *J. Combinatorial Theory Ser. B*, 14:263–267, 1973.
- [14] L. Engebretsen and M. Karpinski. TSP with bounded metrics. *J. Comput. Syst. Sci.*, 72(4):509–546, 2006.
- [15] D. Eppstein. The lattice dimension of a graph. *European J. Combin.*, 26(5):585–592, 2005.
- [16] D. Eppstein. Cubic partial cubes from simplicial arrangements. *Electron. J. Combin.*, 13(1):Research Paper 79, 14 pp. (electronic), 2006.
- [17] D. Eppstein. Recognizing partial cubes in quadratic time. In *Proc. 19th ACM-SIAM Symp. Discrete Algorithms, SODA '08*, pages 1258–1266, Philadelphia, PA, USA, 2008. Society for Industrial and Applied Mathematics.
- [18] D. Eppstein. Paired approximation problems and incompatible inapproximabilities. In *Proc. 21st ACM-SIAM Symp. Discrete Algorithms, SODA '10*, pages 1258–1266, Philadelphia, PA, USA, 2010. Society for Industrial and Applied Mathematics.
- [19] D. Eppstein, J.-C. Falmagne, and S. Ovchinnikov. *Media Theory*. Springer-Verlag, Berlin, 2008.
- [20] S. L. Fitzpatrick and R. J. Nowakowski. The strong isometric dimension of finite reflexive graphs. *Discuss. Math. Graph Theory*, 20(1):23–38, 2000.
- [21] D. Fronček, J. Jerebic, S. Klavžar, and P. Kovář. Strong isometric dimension, biclique coverings, and Sperner's theorem. *Combin. Probab. Comput.*, 16(2):271–275, 2007.

- [22] E. R. Gansner. On the lattice of order ideals of an up-down poset. *Discrete Math.*, 39(2):113–122, 1982.
- [23] M. R. Garey and D. S. Johnson. *Computers and Intractability*. W. H. Freeman and Co., San Francisco, Calif., 1979. A Guide to the Theory of NP-completeness, A Series of Books in the Mathematical Sciences.
- [24] B. Gong and S. Q. Zheng. Near-optimal embeddings of trees into Fibonacci cubes. In *Proc. 28th Southeastern Symp. on System Theory*, pages 421–425, 1996.
- [25] M. Gorše Pihler and J. Žerovnik. Partial cubes are distance graphs. *Discrete Math.*, 308(5-6):820–826, 2008.
- [26] R. L. Graham and P. M. Winkler. On isometric embeddings of graphs. *Trans. Amer. Math. Soc.*, 288(2):527–536, 1985.
- [27] M. Held and R. M. Karp. A dynamic programming approach to sequencing problems. *J. Soc. Indust. Appl. Math.*, 10:196–210, 1962.
- [28] H. Höft and M. Höft. A Fibonacci sequence of distributive lattices. *Fibonacci Quart.*, 23(3):232–237, 1985.
- [29] W.-J. Hsu. Fibonacci cubes—a new interconnection technology. *IEEE Trans. Parallel Distrib. Syst.*, 4(1):3–12, 1993.
- [30] W.-J. Hsu, C. V. Page, and J.-S. Liu. Fibonacci cubes—a class of self-similar graphs. *Fibonacci Quart.*, 31(1):65–72, 1993.
- [31] W. Imrich and S. Klavžar. A simple $O(mn)$ algorithm for recognizing Hamming graphs. *Bull. Inst. Combin. Appl.*, 9:45–56, 1993.
- [32] W. Imrich and S. Klavžar. *Product Graphs: Structure and Recognition*. Wiley-Interscience Series in Discrete Mathematics and Optimization. Wiley-Interscience, New York, 2000.
- [33] W. Imrich, S. Klavžar, and H. M. Mulder. Median graphs and triangle-free graphs. *SIAM J. Discrete Math.*, 12(1):111–118 (electronic), 1999.
- [34] S. Klavžar. On median nature and enumerative properties of Fibonacci-like cubes. *Discrete Math.*, 299(1-3):145–153, 2005.
- [35] S. Klavžar and M. Kovše. On semicube graphs. *European J. Combin.*, 30(1):5–10, 2009.
- [36] S. Klavžar and H. M. Mulder. Partial cubes and crossing graphs. *SIAM J. Discrete Math.*, 15(2):235–251, 2002.
- [37] S. Klavžar and S. Shpectorov. Tribes of cubic partial cubes. *Discrete Math. Theor. Comput. Sci.*, 9(1):273–291, 2007.
- [38] S. Klavžar and P. Žigert. Fibonacci cubes are the resonance graphs of Fibonaccenes. *Fibonacci Quart.*, 43(3):269–276, 2005.
- [39] F. R. McMorris, H. M. Mulder, and F. S. Roberts. The median procedure on median graphs. *Discrete Appl. Math.*, 84(1-3):165–181, 1998.

- [40] S. Micali and V. V. Vazirani. An $O(\sqrt{|V|}|E|)$ Algorithm for Finding Maximum Matching in General Graphs. In *Proc. 21st Annu. Symp. on Foundations of Computer Science, FOCS' 80*, pages 17–27, 1980.
- [41] E. Munarini and N. Salvi Zagaglia. Structural and enumerative properties of the Fibonacci cubes. *Discrete Math.*, 255(1-3):317–324, 2002.
- [42] S. Ovchinnikov. Partial cubes: structures, characterizations, and constructions. *Discrete Math.*, 308(23):5597–5621, 2008.
- [43] N. Polat. Netlike partial cubes I. General properties. *Discrete Math.*, 307(22):2704–2722, 2007.
- [44] N. Polat. Netlike partial cubes. IV. Fixed finite subgraph theorems. *European J. Combin.*, 30(5):1194–1204, 2009.
- [45] A. Taranenko and A. Vesel. Fast recognition of Fibonacci cubes. *Algorithmica*, 49(2):81–93, 2007.
- [46] P. M. Winkler. Isometric embedding in products of complete graphs. *Discrete Appl. Math.*, 7(2):221–225, 1984.