

Separator-based graph embedding into multidimensional grids with small edge-congestion

メタデータ	言語: eng 出版者: 公開日: 2017-10-03 キーワード (Ja): キーワード (En): 作成者: メールアドレス: 所属:
URL	http://hdl.handle.net/2297/40596

Separator-Based Graph Embedding into Multidimensional Grids with Small Edge-Congestion

Akira Matsubayashi*

Division of Electrical Engineering and Computer Science, Kanazawa University, Kanazawa 920-1192, Japan

Abstract

We study the problem of embedding a guest graph with minimum edge-congestion into a multidimensional grid with the same size as that of the guest graph. Based on a well-known notion of graph separators, we show that an embedding with a smaller edge-congestion can be obtained if the guest graph has a smaller separator, and if the host grid has a higher but constant dimension. Specifically, we prove that any graph with N nodes, maximum node degree Δ , and with a node-separator of size s , where s is a function such that $s(n) = O(n^\alpha)$ with $0 \leq \alpha < 1$, can be embedded into a grid of a fixed dimension $d \geq 2$ with at least N nodes, with an edge-congestion of $O(\Delta)$ if $d > 1/(1 - \alpha)$, $O(\Delta \log N)$ if $d = 1/(1 - \alpha)$, and $O(\Delta N^{\alpha-1+\frac{1}{d}})$ if $d < 1/(1 - \alpha)$. This edge-congestion achieves constant ratio approximation if $d > 1/(1 - \alpha)$, and matches an existential lower bound within a constant factor if $d \leq 1/(1 - \alpha)$. Our result implies that if the guest graph has an excluded minor of a fixed size, such as a planar graph, then we can obtain an edge-congestion of $O(\Delta \log N)$ for $d = 2$ and $O(\Delta)$ for any fixed $d \geq 3$. Moreover, if the guest graph has a fixed treewidth, such as a tree, an outerplanar graph, and a series-parallel graph, then we can obtain an edge-congestion of $O(\Delta)$ for any fixed $d \geq 2$. To design our embedding algorithm, we introduce *edge-separators bounding extension*, such that in partitioning a graph into isolated nodes using edge-separators recursively, the number of outgoing edges from a subgraph to be partitioned in a recursive step is bounded. We present an algorithm to construct an edge-separator with extension of $O(\Delta n^\alpha)$ from a node-separator of size $O(n^\alpha)$.

Keywords: graph embedding, edge-congestion, grid, separator, extension

1. Introduction

A *graph embedding* of a guest graph into a host graph is to map (typically one-to-one) nodes and edges of the guest graph onto nodes and paths of the host graph, respectively, so that an edge of the guest graph is mapped onto a path connecting the images of end-nodes of the edge. The graph embedding problem is to embed a guest graph into

*Tel/Fax: +81 76 234 4837

Email address: mbayashi@t.kanazawa-u.ac.jp (Akira Matsubayashi)

a host graph with certain constraints and/or optimization criteria. This problem has applications such as efficient VLSI layout and parallel computation. I.e., the problem of efficiently laying out VLSI can be formulated as the graph embedding problem with modeling a design rule on wafers and a circuit to be laid out as host and guest graphs, respectively. Also, the problem of efficiently implementing a parallel algorithm on a message passing parallel computer system consisting of processing elements connected by an interconnection network can be formulated as the graph embedding problem with modeling the interconnection network and interprocess communication in the parallel algorithm as host and guest graphs, respectively. See for a survey, e.g., [1]. The major criteria to measure the efficiency of an embedding are dilation, node-congestion, and edge-congestion, load, and expansion, whose formal definitions are given in Sect. 2. In this paper, we consider the problem of embedding a guest graph with the unit load and minimum edge-congestion into a d -dimensional grid with $d \geq 2$ and the same size as that of the guest graph (i.e., with unit expansion). Embeddings into grids with the minimum edge-congestion are important for both VLSI layout and parallel computation. Actually, design rules on wafers in VLSI are usually modeled as 2-dimensional grids, and the minimum edge-congestion provides a lower bound on the number of layers needed to lay out a given circuit. As for parallel computation, multidimensional grid networks, including hypercubes, are popular for interconnection networks. On interconnection networks adopting circuit switching or wormhole routing, in particular, embeddings with the edge-congestion of 1 are essential to minimize the communication latency [2, 3, 4]. In addition, the setting that host and guest graphs have the same number of nodes is important for parallel computation because the processing elements are expensive resource and idling some of them is wasteful.

Previous Results

Graph embedding into grids with small edge-congestion has extensively been studied. Table 1 summarizes previous results of graph embeddings minimizing edge-congestion (and other criteria as well in some results) for various combinations of guest graphs and host grids.

VLSI layout has been studied through formulating the layout as the graph embedding into a 2-dimensional grid with objective of minimizing the grid under constrained congestion-1 routing [15]. Leiserson [16] and Valiant [17] independently proposed such embeddings based on graph separators. In particular, it was proved in [16] that any N -node graph with maximum node degree at most 4 and an edge-separator of size s , where s is a function with $s(n) = O(n^\alpha)$, can be laid out in an area of $O(N)$ if $\alpha < 1/2$, $O(N \log^2 N)$ if $\alpha = 1/2$, and $O(N^{2\alpha})$ if $\alpha > 1/2$. A separator of a graph G is a set S of either nodes or edges whose removal partitions the node set $V(G)$ of G into two subsets of roughly the same size with no edge between the subsets. The graph G is said to have a (recursive) separator of size $s(n)$ if $|S| \leq s(|V(G)|)$ and the subgraphs partitioned by S recursively have separators of size $s(n)$. Separators are important tools to design divide-and-conquer algorithms and have been extensively studied. Bhatt and Leighton [11] achieved a better layout with several nice properties including reduced dilation as well as the same or better area as that of [16] by introducing a special type of edge-separators called *bifurcators*. An approximation algorithm for VLSI layout

Table 1: Previous results of graph embeddings minimizing edge-congestion.

Guest Graph	Host Grid		Congestion	Dilation	
N : # nodes, Δ : max degree s : separator size	# nodes	dimension			
connected planar graph	N	2	NP-hard for 1	any	[5]
connected graph	$2^{\lceil \log_2 N \rceil}$	$\lceil \log_2 N \rceil$	NP-hard for 1	any	[2]
complete binary tree	$N + 1$	2	2	$O(\sqrt{N})$	[6]
complete binary tree	$N + 1$	4	1	$O(N^{1/4})$	[3]
complete binary tree	$N + O(\sqrt{N})$	2	1	$O(\sqrt{N})$	[7]
complete k -ary tree ($k \geq 3$)	$N + O(N/\sqrt{k})$	2	$\lceil k/2 \rceil + 1$	$O(\sqrt{N})$	[4]
binary tree	$2^{\lceil \log_2 N \rceil}$	$\lceil \log_2 N \rceil$	5	$\lceil \log_2 N \rceil$	[8]
2-D $h \times w$ -grid ($h \leq w$)	$h'w' \geq N^*$	2	$\lceil h/h' \rceil + 1$	$\lceil h/h' \rceil + 1$	[9]
2-D $h \times w$ -grid ($h \leq w$)	$h'w' \geq N^\dagger$	2	5	5	[9]
2-D $h \times w$ -grid ($h \leq w$)	$h'w' \geq N^\dagger$	2	4	$\geq 4h - 3$	[9]
2-D grid	$2^{\lceil \log_2 N \rceil}$	$\lceil \log_2 N \rceil$	2	3	[10]
$\Delta \leq 4, s = O(n^\alpha), \alpha < 1/2$	$O(N)$	2	1	$O(\sqrt{N}/\log N)$	[11]
$\Delta \leq 4, s = O(\sqrt{n})$	$O(N \log^2 N)$	2	1	$O(\frac{\sqrt{N} \log N}{\log \log N})$	[11]
$\Delta \leq 4, s = O(n^\alpha), \alpha > 1/2$	$O(N^{2\alpha})$	2	1	$O(N^\alpha)$	[11]
tree width t	$2^{\lceil \log_2 N \rceil}$	$\lceil \log_2 N \rceil$	$O(\Delta^4 t^3)$	$O(\log(\Delta t))$	[12]
$s = \log^{O(1)} N$	$2^{\lceil \log_2 N \rceil}$	$\lceil \log_2 N \rceil$	$\Delta^{O(1)}$	$O(\log \Delta)$	[13]
$\Delta = O(1)$	N	$d = O(1)$	$O(N^{1/d} \log N)$	$O(N^{1/d} \log N)$	[14]
$\Delta \leq 2 \lceil \log_2 N \rceil$	$2^{2 \lceil \log_2 N \rceil}$	$2 \lceil \log_2 N \rceil$	1	$2 \lceil \log_2 N \rceil$	[8]

* $h' \times w'$ -grid with $h' < h \leq w < w'$

† $h' \times w'$ -grid with $h < h' \leq w' < w$

was proposed in [18]. Separator-based graph embeddings on hypercubes were presented in [19, 20, 13]. In particular, Heun and Mayr [13] proved that any N -node graph with maximum node degree Δ and an extended edge-bisector of polylogarithmic size can be embedded into a $\lceil \log_2 N \rceil$ -dimensional cube with a dilation of $O(\log \Delta)$ and an edge-congestion of $\Delta^{O(1)}$.

A quite general embedding based on the multicommodity flow was presented by Leighton and Rao [14], who proved that any N -node bounded degree graph G can be embedded into an N -node bounded degree graph H with both dilation and edge-congestion of $O((\log N)/\alpha)$, where α is the flux of H , i.e., $\min_{U \subset V(H)} \frac{|\{(u,v) \in E(H) | u \in U, v \in V(H) \setminus U\}|}{\min\{|U|, |V(H) \setminus U|\}}$. This implies that G can be embedded into an N -node d -dimensional grid with both dilation and edge-congestion of $O(N^{1/d} \log N)$ for any fixed d .

Contributions and Technical Overview

In this paper, we improve previous graph embeddings into grids and hypercubes in terms of edge-congestion, arbitrary dimension, and minimum size of host grids. In particular, we claim that if a guest graph has a small separator, then we do not need grids with large dimension, such as hypercubes, to suppress the edge-congestion.

First, we present an embedding algorithm based on the permutation routing. The permutation routing is to construct paths connecting given pairs of source and destination nodes such that no two pairs have the same sources or the same destinations. This embedding algorithm achieves an edge-congestion as stated in the following theorem:

Theorem 1. *Any graph with N nodes and maximum node degree Δ can be embedded into a d -dimensional $\ell_1 \times \cdots \times \ell_d$ -grid ($\prod_{i=1}^d \ell_i \geq N$) with a dilation at most $2 \sum_{i=1}^d \ell_i$ and an edge-congestion at most $2\lceil\Delta/2\rceil \cdot \max_i\{\ell_i\}$.*

We prove this theorem in Sect. 4.1 by observing that for any one-to-one mapping of nodes of a guest graph G to nodes of a host graph H , routing edges of G on H can be reduced to at most $\lceil\Delta/2\rceil$ instances of permutation routing, and that the permutation routing algorithm proposed in [21] has an edge-congestion at most $2 \cdot \max_i\{\ell_i\}$. Theorem 1 achieves an edge-congestion of $2\lceil\Delta/2\rceil\lceil N^{1/d}\rceil$ if $\ell_i = \lceil N^{1/d}\rceil$ for each i . It is worth noting that this edge-congestion can slightly be improved if the host grid H is a d -dimensional cube. It is well-known that any one-to-one mapping of 2^{d+1} inputs to 2^{d+1} outputs on a d -dimensional Beneš network can be routed with the edge-congestion 1 [22]. We can easily observe that mapping the nodes in each row of the Beneš network to each node of H induces a (many-to-one) embedding with the edge-congestion 4. Because each node of H has exactly two inputs and two outputs in a row of the Beneš network, any pair of instances of permutation routing on H can be routed with an edge-congestion at most 4. At most $\lceil\Delta/2\rceil$ instances of permutation routing, obtained from any one-to-one mapping of nodes of G to nodes of H and from edges of G , can be grouped into $\lceil\lceil\Delta/2\rceil/2\rceil = \lceil\Delta/4\rceil$ pairs of instances of permutation routing. Therefore, G can be embedded into a $\lceil\log_2 N\rceil$ -dimensional cube with an edge-congestion at most $4\lceil\Delta/4\rceil$.

Second, we present an embedding algorithm based on separators that achieves an edge-congestion as stated in the following theorem:

Theorem 2. *Suppose that G is a graph with N nodes, maximum node degree Δ , and with a node-separator of size $s(n) = O(n^\alpha)$ ($0 \leq \alpha < 1$), and that M is a grid with a fixed dimension $d \geq 2$, at least N nodes, and with constant aspect ratio. Then, G can be embedded into M with a dilation of $O(dN^{1/d})$, and with an edge-congestion of $O(\Delta)$ if $d > 1/(1 - \alpha)$, $O(\Delta \log N)$ if $d = 1/(1 - \alpha)$, and $O(\Delta N^{\alpha-1+\frac{1}{d}})$ if $d < 1/(1 - \alpha)$.*

The basic idea of Theorem 2 is to partition the guest graph and the host grid using their edge-separators, embed the partitioned guest graphs into the partitioned host grids recursively, and to route cut edges of the guest graph on the host grid. We use Theorem 1 to route cut edges with a nearly minimum edge-congestion in each recursive step. However, just doing this is not sufficient for our goal. In fact, we need further techniques to suppress the total edge-congestion incurred by whole recursive steps from upper to lower levels. There are two reasons of the insufficiency.

The first reason is that recursive steps from upper to lower levels may use the same edge of the grid, which yields an edge-congestion of $\Omega(\log N)$ if we minimize the edge-congestion only in each individual recursive step. This is a crucial barrier to achieve an edge-congestion of $O(\Delta)$ for $d > 1/(1 - \alpha)$. To solve this, we divide the edge set of the grid into $\Theta(\log N)$ subsets of appropriate size and use each subset only in a constant number of recursive steps.

The second and more significant reason is that a small subgraph of the guest graph to be embedded in a lower recursive step may have nodes incident to quite a large number of edges that have been cut in upper levels, which yields a large edge-congestion.

Specifically, if such a subgraph has n nodes and x outgoing edges to the other part of the guest graph, then because a subgrid into which the subgraph is embedded has $O(dn^{1-\frac{1}{d}})$ outgoing edges, the edge-congestion is lower bounded by $x/O(dn^{1-\frac{1}{d}}) = \Omega(xn^{\frac{1}{d}-1}/d)$. A standard edge-separator aims to minimize the number of edges to be cut to partition a graph. Thus, if we recursively use such edge-separators to partition a graph into small pieces, then although the number of cut edges in each recursive step is bounded, the number of outgoing edges from a subgraph to be embedded in a lower recursive step may become extremely large compared to the number of nodes of the subgraph. Therefore, we introduce edge-separators bounding *extension*, i.e., the number of outgoing edges from a subgraph in each recursive step, and present an algorithm to construct an edge-separator with extension of $O(\Delta n^\alpha)$ from a node-separator of size $O(n^\alpha)$. We describe the algorithm for edge-separators with bounded extension in Sect. 3 and prove Theorem 2 in Sect. 4.2.

Theorem 2 achieves constant ratio approximation for a fixed $d > 1/(1-\alpha)$ because any embedding has an edge-congestion at least $\Delta/(2d)$. If $d \leq 1/(1-\alpha)$, then the edge-congestion of Theorem 2 matches an existential lower bound within a constant factor. The lower bound of $\Omega(\log N)$ for $d = 1/(1-\alpha) = 2$ and $\Delta = O(1)$ is derived from the following fact: There exists an N -node guest graph with constant degree and a node-separator of size $s(n) = O(\sqrt{n})$ whose any embedding into a 2-dimensional grid with the edge-congestion 1 requires $\Omega(N \log^2 N)$ nodes of the grid [23].¹ This implies that any embedding of the guest graph into a 2-dimensional grid with N nodes requires an edge-congestion of $\Omega(\log N)$. This is because we can easily transform an embedding into an N -node grid with an edge-congestion c into another embedding into an $O(c^2N)$ -node grid with the edge-congestion 1 by replacing each row and each column of the N -node grid with $O(c)$ rows and $O(c)$ columns, respectively.² A similar transformation for VLSI layout is described in [15].

The lower bound of $\Omega(\Delta N^{\alpha-1+\frac{1}{d}})$ for $d < 1/(1-\alpha)$ can be obtained as follows: We consider a guest graph G with N nodes and a node-separator of size $s(n) = n^\alpha$ such that each node in a cut set $U \subseteq V(G)$ with $|U| = N^\alpha$ is adjacent to every other node in G . The graph G obviously has $\Delta = N - 1$. Suppose that we arbitrarily divide $V(G)$ into two subsets of the same size. Then, at least $(|U|/2)(N-1)/2 = \Delta N^\alpha/4$ edges join nodes in one of the subsets and nodes in the other subset because at least half nodes of U are contained in one of the subsets and adjacent to all nodes in the other subset. On the other hand, we can divide a d -dimensional N -node grid into two subgrids of the same size by removing $O(N^{1-\frac{1}{d}})$ edges. Thus, any embedding of G into the grid has an edge-congestion at least $(\Delta N^\alpha/4)/O(N^{1-\frac{1}{d}}) = \Omega(\Delta N^{\alpha-1+\frac{1}{d}})$.

Theorem 2 has the following applications. It is well-known that any planar graph has a node-separator of size $s(n) = O(\sqrt{n})$ [24]. This was generalized in [25] so that any graph with an excluded minor of a fixed size has a node-separator of size $s(n) = O(\sqrt{n})$. Therefore, we obtain the following corollary:

¹Strictly, this result is proved for the VLSI layout model. However, we can easily generalize this result to the embedding model considered in this paper.

²It should be noted that the inverse transformation cannot be done in such a simple way. In fact, we do not know whether or not the inverse transformation is always possible.

Corollary 1. *Any graph with N nodes, maximum node degree Δ , and with an excluded minor of a fixed size can be embedded into a grid of a fixed dimension d with at least N nodes and constant aspect ratio, with an edge-congestion of $O(\Delta \log N)$ for $d = 2$ and $O(\Delta)$ for $d \geq 3$.*

Graphs with a fixed treewidth, such as trees, outerplanar graphs, and series-parallel graphs have a node-separator of a fixed size [26]. Therefore, we obtain the following corollary:

Corollary 2. *Any graph with N nodes, maximum node degree Δ , and with a fixed treewidth can be embedded into a grid of a fixed dimension at least 2 with at least N nodes and constant aspect ratio, with an edge-congestion of $O(\Delta)$.*

Our separator-based embedding algorithm performs in a polynomial time on the condition that a separator of the guest graph is given. Although finding a separator of minimum size is generally NP-hard [27, 28], approximation algorithms presented in [14, 29, 30] can be applied to our algorithm.

All our embedding algorithms yield a dilation of order of the diameter of the host grid. Although such a dilation is trivial when only the dilation is minimized, this is not the case when edge-congestion is minimized. As we will demonstrate in Sect. 5, in fact, there exists an N -node guest graph whose any embedding with the edge-congestion 1 into an N -node 2-dimensional grid requires a dilation of $\Theta(N)$, far from the diameter $\Theta(\sqrt{N})$. We do not know whether or not we can always achieve both a dilation of the host grid's diameter (even with a multiplicative constant factor) and constant ratio approximation for edge-congestion. This is negative if the host graph is general. As an example, suppose that H is the host graph obtained from a complete binary tree with N leaves by adding edges so that the leaves induce a $\sqrt{N} \times \sqrt{N}$ -grid. To be precise, the $N/2^i$ leaves of a subtree rooted by a node at an even distance i to the root induce a $\sqrt{N/2^i} \times \sqrt{N/2^i}$ -subgrid. If the guest graph G is an N -node complete graph, then any embedding of G into H with a dilation of the diameter $O(\log N)$ of H has an edge-congestion of $\Omega(N^2)$ because $\Omega(N^2)$ edges of G must be routed through a single node of the tree part in H to achieve such a dilation, while G can be embedded into the grid part in H with a dilation of $O(\sqrt{N})$ and an edge-congestion of $O(N^{3/2})$ using a simple row-column routing.

2. Preliminaries

For a directed or undirected graph G , $V(G)$ and $E(G)$ are the node set and edge set of G , respectively. In the rest of the paper, we call undirected graphs simply graphs. We denote the set of integers $\{i \mid 1 \leq i \leq \ell\}$ by $[\ell]$. For a d -dimensional vector $v := (x_i)_{i \in [d]}$, let $\pi_j(v) := x_j$ and $\bar{\pi}_j(v) := (x_i)_{i \in [d] \setminus \{j\}}$ for $j \in [d]$. We use π_j and $\bar{\pi}_j$ also for a set of vectors and for a directed graph whose nodes are vectors. I.e., for a set V of d -dimensional vectors, we denote $\{\pi_j(v) \mid v \in V\}$ and $\{\bar{\pi}_j(v) \mid v \in V\}$ by $\pi_j(V)$ and $\bar{\pi}_j(V)$, respectively. Moreover, for a directed graph G with $V(G) = V$, we denote the graph with the node set $\bar{\pi}_j(V(G))$ and edge multiset $\{(\bar{\pi}_j(u), \bar{\pi}_j(v)) \mid (u, v) \in E(G)\}$ by $\bar{\pi}_j(G)$. For positive integers ℓ_1, \dots, ℓ_d , the d -dimensional $\ell_1 \times \dots \times \ell_d$ -grid, denoted by $M(\ell_i)_{i \in [d]}$, is a graph with the node set $\prod_{i \in [d]} [\ell_i]$, i.e., the Cartesian product of sets




Figure 1: A routing graph R on $M(2, 3)$ and $\bar{\pi}_1(R)$ on $M(3)$.

$[\ell_1], \dots, [\ell_d]$, and edge set $\{\{u, v\} \mid \exists j \in [d] \pi_j(u) = \pi_j(v) \pm 1, \bar{\pi}_j(u) = \bar{\pi}_j(v)\}$. The *aspect ratio* of $M(\ell_i)_{i \in [d]}$ is $\max_{i, j \in [d]} \{\ell_j / \ell_i\}$. An edge $\{u, v\}$ of $M(\ell_i)_{i \in [d]}$ with $\pi_j(u) = \pi_j(v) \pm 1$ is called a *dimension- j edge*. The grid $M(\ell_i)_{i \in [d]}$ is called the *d -dimensional cube* if $\ell_i = 2$ for every $i \in [d]$.

A *routing request* on a graph H is a pair of nodes, a *source* and *target*, of H . A multiset of routing requests can be represented as a *routing graph* R with the node set $V(H)$ and directed edges joining the sources and targets of all the routing requests. It should be noted that R may have parallel edges and loops. In particular, if H is a d -dimensional grid, then $\bar{\pi}_j(R)$ is a routing graph with the multiset of edges $(\bar{\pi}_j(u), \bar{\pi}_j(v))$ for every $(u, v) \in E(R)$ on the $(d - 1)$ -dimensional grid with node set $\bar{\pi}_j(V(H))$ (cf. Fig. 1). R is called a *p - q routing graph* if the maximum outdegree and indegree of R are at most p and q , respectively. A 1-1 routing graph is also called a *permutation routing graph*. We define a *routing* of R as a mapping ρ that maps each edge $(u, v) \in E(R)$ onto a set of edges of H inducing a path connecting u and v . We denote $\rho((u, v))$ simply by $\rho(u, v)$. The *dilation* and *edge-congestion* of ρ are $\max_{e \in E(R)} |\rho(e)|$ and $\max_{e' \in E(H)} \{|e \in E(R) \mid e' \in \rho(e)\}|$, respectively.

An *embedding* $\langle \phi, \rho \rangle$ of a graph G into a graph H is a pair of mappings consisting of a mapping $\phi : V(G) \rightarrow V(H)$ and a routing ρ of an arbitrary orientation of the graph with the node set $V(H)$ and edge set $\{\{\phi(u), \phi(v)\} \mid \{u, v\} \in E(G)\}$. The *dilation* and *edge-congestion* of the embedding $\langle \phi, \rho \rangle$ are defined as the dilation and edge-congestion of ρ , respectively. The *load* and *expansion* of $\langle \phi, \rho \rangle$ is $\max_{v \in V(H)} \{|u \in V(G) \mid \phi(u) = v\}|$ and $|V(H)|/|V(G)|$, respectively.

3. Edge-Separators with Bounded Extension

The (recursive) node- and edge-separators are formally defined as follows: Let $1/2 \leq \beta < 1$ and $s(n)$ be a non-decreasing function. A graph G has a *β -node(edge, resp.)-separator of size $s(n)$* if $|V(G)| = 1$, or if G can be partitioned into two subgraphs with at most $\beta|V(G)|$ nodes ($\lceil \beta|V(G)| \rceil$ nodes, resp.) and with no edges connecting the subgraphs by removing at most $s(|V(G)|)$ nodes (edges, resp.), and the subgraphs recursively have a β -node(edge, resp.)-separator of size $s(n)$. The process of partitioning G into isolated nodes using the edge-separator repeatedly is often referred to as a *decomposition tree*. The decomposition tree \mathcal{T} is a rooted tree having a set of subgraphs




Figure 2: A $(1/2)$ -decomposition tree for $M(2,4)$ with extension 3. Dashed lines represent external edges of nodes of the decomposition tree.

of G as its node set $V(\mathcal{T})$ such that the root of \mathcal{T} is G , each non-leaf node $H \in V(\mathcal{T})$ has exactly two children obtained from H by removing the edge-separator of H , and that each leaf node of \mathcal{T} consists of a single node of G . We call \mathcal{T} a β -decomposition tree with extension $x(n)$ if it can be constructed using a β -edge-separator, and for each $H \in V(\mathcal{T})$, at most $x(|V(H)|)$ edges (called *external edges* of H in this paper) connect $V(H)$ and $V(G) \setminus V(H)$ (cf. Fig. 2).

A decomposition tree with reasonably small extension can be obtained from a node-separator as stated in the following lemma:

Lemma 1. *Any graph G with maximum node degree Δ and a β -node-separator of size $s(n) = Cn^\alpha$ ($C > 0$, $0 \leq \alpha < 1$, $1/2 \leq \beta < 1$) has a $\frac{\beta}{1-\epsilon}$ -decomposition tree with extension $x(n) = O(C\Delta n^\alpha / \epsilon)$, where $0 < \epsilon < 1 - \beta$.*

PROOF. We present an algorithm constructing a desired decomposition tree \mathcal{T} . We initially set G as the root of \mathcal{T} and construct \mathcal{T} from the root toward leaves. Assume that we have constructed \mathcal{T} up to depth (distance to the root) $i - 1 \geq 0$. For a subgraph H of G at depth $i - 1$ in \mathcal{T} , we construct children H_1 and H_2 of H as follows:

1. We inductively assume the following:
 - (a) Each node of \mathcal{T} up to depth $i - 1$ has been constructed by partitioning a subgraph of G using a β -node-separator and distributing the node-separator between the partitioned graphs. Let X_{i-1} be the set of nodes of H contained in the node-separator used for any ancestor of H in \mathcal{T} .
 - (b) All the external edges of H are incident to nodes in X_{i-1} .
 - (c) The graph H' obtained from H by removing X_{i-1} has a β -node-separator $S_i \subseteq V(H')$ of size Cn^α .

It should be noted that $X_0 = \emptyset$, and therefore, these assumptions hold if $H = G$.




Figure 3: Partition of H into H_1 and H_2 .

2. If $C|V(H')|^\alpha \leq \epsilon|V(H')|$, then we partition H' into subgraphs H'_1 and H'_2 using the node-separator S_i with $|S_i| \leq C|V(H')|^\alpha$. It follows that $|V(H'_1)| + |V(H'_2)| = |V(H')| - |S_i| \geq (1 - \epsilon)|V(H')| \geq \frac{1-\epsilon}{\beta}|V(H'_1)|$. Assume without loss of generality that $|V(H'_1)| \geq |V(H'_2)|$. Then, there exists $1/2 \leq \beta' \leq \frac{\beta}{1-\epsilon}$ with $|V(H'_1)| = \beta'(|V(H'_1)| + |V(H'_2)|)$.
3. If $C|V(H')|^\alpha > \epsilon|V(H')|$, then reset $S_i := V(H')$, and arbitrarily choose $1/2 \leq \beta' \leq \frac{\beta}{1-\epsilon}$.
4. Partition $X_{i-1} \cup S_i$ into two disjoint sets Y_1 and Y_2 such that $|Y_1| = \lceil \beta'(|X_{i-1}| + |S_i|) \rceil$ and $|Y_2| = \lfloor (1 - \beta')(|X_{i-1}| + |S_i|) \rfloor$.
5. Let H_j be the subgraph of H induced by $V(H'_j) \cup Y_j$ for $j = 1, 2$. We illustrate the construction in Fig. 3.

We first observe that H_1 and H_2 satisfy the inductive assumptions of the algorithm. For $j \in \{1, 2\}$, by inductive assumption, Y_j is the set of nodes of H_j contained in the node-separator used for an ancestor of H_j . As shown in Fig. 3, all the external edges of H_j are incident to nodes of Y_j . Moreover, the subgraph of H_j obtained by removing Y_j is H'_j , which is the subgraph of H' partitioned by the node-separator S_i of H' . Therefore, H'_j has a β -node-separator of size Cn^α .

We then estimate the numbers of nodes of H_1 and H_2 . By definition, it follows that

$$\begin{aligned} |V(H_1)| &= |V(H'_1)| + |Y_1| = \beta'(|V(H'_1)| + |V(H'_2)|) + \lceil \beta'(|X_{i-1}| + |S_i|) \rceil \\ &= \lceil \beta'|V(H)| \rceil, \text{ and} \end{aligned} \quad (1)$$

$$\begin{aligned} |V(H_2)| &= |V(H'_2)| + |Y_2| = (1 - \beta')(|V(H'_1)| + |V(H'_2)|) + \lfloor (1 - \beta')(|X_{i-1}| + |S_i|) \rfloor \\ &= \lfloor (1 - \beta')|V(H)| \rfloor \leq \beta'|V(H)|. \end{aligned} \quad (2)$$

These imply that the algorithm constructs \mathcal{T} as a β' -decomposition tree of G .

We finally prove that for $j = 1, 2$, H_j has $O(C\Delta|V(H_j)|^\alpha/\epsilon)$ external edges, implying extension $O(C\Delta n^\alpha/\epsilon)$ of \mathcal{T} . We prove this only for H_1 because the proof for H_2 is obtained with a similar argument. Because all the external edges of H_1 are incident to Y_1 , it suffices to show that $|Y_1| = O(C|V(H_1)|^\alpha/\epsilon)$. When we construct children of H_1 using the algorithm, X_i is set to Y_1 . Let $n_i := |V(H_1)|$ and n_j ($0 \leq j < i$) be the number

of nodes of the ancestor of H_1 at depth j in \mathcal{T} . Moreover, let β_j ($1 \leq j \leq i$) be β' or $1 - \beta'$ defined in Step 2 or 3 in partitioning the ancestor at depth $j - 1$. This implies that $n_j = \lceil \beta_j n_{j-1} \rceil$ as in (1) or $n_j = \lfloor \beta_j n_{j-1} \rfloor$ as in (2). Therefore,

$$n_j \leq \lceil \beta_j n_{j-1} \rceil \leq \beta_j n_{j-1} + 1 \leq n_0 \prod_{h=1}^j \beta_h + \sum_{\ell=1}^j \prod_{h=\ell+1}^j \beta_h = n_0 \prod_{h=1}^j \beta_h + O(1), \text{ and} \quad (3)$$

$$n_j \geq \lfloor \beta_j n_{j-1} \rfloor \geq \beta_j n_{j-1} - 1 \geq n_0 \prod_{h=1}^j \beta_h - \sum_{\ell=1}^j \prod_{h=\ell+1}^j \beta_h = n_0 \prod_{h=1}^j \beta_h - O(1). \quad (4)$$

Here, we have used the fact that $\sum_{\ell=1}^j \prod_{h=\ell+1}^j \beta_h \leq \sum_{\ell=1}^j (\frac{\beta}{1-\epsilon})^{j-\ell} = O(1)$. By the definition of Y_1 , we have the following recurrence of $|X_i|$:

$$|X_i| = |Y_1| = \lceil \beta_i(|X_{i-1}| + |S_i|) \rceil \leq \beta_i(|X_{i-1}| + |S_i|) + 1 \leq \sum_{j=1}^i |S_j| \prod_{h=j}^i \beta_h + \sum_{j=1}^i \prod_{h=j+1}^i \beta_h.$$

The number $|S_j|$ is less than Cn_{j-1}^α/ϵ because $|S_j| \leq C|V(H')|^\alpha \leq Cn_{j-1}^\alpha < Cn_{j-1}^\alpha/\epsilon$ if S_j is defined in Step 2, and $|S_j| = |V(H')| < C|V(H')|^\alpha/\epsilon \leq Cn_{j-1}^\alpha/\epsilon$ if S_j is defined in Step 3. Moreover, $\sum_{j=1}^i \prod_{h=j+1}^i \beta_h = O(1)$ as estimated for (3) and (4). Therefore,

$$\begin{aligned} |Y_1| &< \sum_{j=1}^i \frac{Cn_{j-1}^\alpha}{\epsilon} \prod_{h=j}^i \beta_h + O(1) \leq \sum_{j=1}^i \frac{C}{\epsilon} \left(n_0 \prod_{h=1}^{j-1} \beta_h + O(1) \right)^\alpha \prod_{h=j}^i \beta_h + O(1) \quad [\text{by (3)}] \\ &= O\left(\frac{Cn_0^\alpha}{\epsilon} \sum_{j=1}^i \prod_{h=1}^{j-1} \beta_h^\alpha \cdot \prod_{h=j}^i \beta_h \right) = O\left(\frac{Cn_0^\alpha}{\epsilon} \prod_{h=1}^i \beta_h^\alpha \cdot \sum_{j=1}^i \prod_{h=j}^i \beta_h^{1-\alpha} \right) \\ &= O\left(\frac{C}{\epsilon} \left(n_0 \prod_{h=1}^i \beta_h \right)^\alpha \cdot \sum_{j=1}^i \left(\frac{\beta}{1-\epsilon} \right)^{(1-\alpha)(i-j+1)} \right) \\ &= O\left(\frac{C}{\epsilon} (n_i + O(1))^\alpha \cdot O(1) \right) \quad [\text{by (4)}] \\ &= O\left(\frac{Cn_i^\alpha}{\epsilon} \right). \end{aligned}$$

Therefore, \mathcal{T} is a desired decomposition tree. \square

4. Embedding Algorithm

In this section, we first prove Theorem 1 by estimating the edge-congestion of the previously known permutation routing algorithm on multidimensional grids presented in [21]. We then provide an embedding algorithm based on edge-separators with bounded extension as well as the permutation routing algorithm. Combining this algorithm with Lemma 1, we prove Theorem 2.

4.1. Permutation Routing and Embedding

Any permutation routing can be used to construct a graph embedding as follows:

Lemma 2. *If any 1-1 routing graph on a host graph H can be routed with an edge-congestion at most c , then any graph with maximum node degree Δ can be embedded into H with an edge-congestion at most $c\lceil\Delta/2\rceil$.*

PROOF. Let G be a graph with maximum node degree Δ to be embedded into H . We arbitrarily choose a one-to-one mapping $\phi : V(G) \rightarrow V(H)$. Let G' be the graph with node set $\phi(V(G))$ and edge set $\{\{\phi(u), \phi(v)\} \mid \{u, v\} \in E(G)\}$. Because G' is an undirected graph with maximum node degree Δ , there is an orientation R of G' whose maximum indegree and outdegree are both at most $\lceil\Delta/2\rceil$. Such an orientation can be obtained by adding dummy edges joining nodes with odd degree so that the resulting graph has an Euler circuit, and by orienting edges along with the Euler circuit. It suffices to prove that R as a routing graph on H can be routed with an edge-congestion at most $c\lceil\Delta/2\rceil$.

We decompose R into at most $\lceil\Delta/2\rceil$ edge-disjoint 1-1 routing graphs each of which has nodes $V(R)$ and edges with the same color in an edge-coloring of R such that no two edges with the same sources or with the same targets have the same color. Such coloring can be obtained by edge-coloring the bipartite graph consisting of the source and target sets of R , i.e., two copies V^+ and V^- of $V(R)$, and edges joining $u \in V^+$ and $v \in V^-$ for all $(u, v) \in E(R)$. It should be noted that the resulting bipartite graph has node-degree at most $\lceil\Delta/2\rceil$, and hence, $\lceil\Delta/2\rceil$ colors are enough for the coloring [31]. Therefore, R can be routed on H with an edge-congestion at most $c\lceil\Delta/2\rceil$ if each of the 1-1 routing graphs can be routed with an edge-congestion at most c . \square

The algorithm of [21] routes a 1-1 routing graph R on $M := M(\ell_i)_{i \in [d]}$ as follows:

1. Color edges of R using at most ℓ_1 colors so that when we identify edges in R with corresponding edges in $\bar{\pi}_1(R)$, no two edges with the same sources or with the same targets in $\bar{\pi}_1(R)$ have the same color. This coloring can be obtained as done in the proof of Lemma 2. It should be noted that $\bar{\pi}_1(R)$ is a ℓ_1 - ℓ_1 routing graph with node set $\bar{\pi}_1(V(M))$.
2. Decompose R into edge-disjoint subgraphs R_1, \dots, R_{ℓ_1} each of which has nodes $V(R)$ and edges with the same color.
3. For each $i \in [\ell_1]$, $\bar{\pi}_1(R_i)$ is a 1-1 routing graph with node set $\bar{\pi}_1(V(M))$. Therefore, we can recursively find a routing ρ_i of $\bar{\pi}_1(R_i)$ on the $(d-1)$ -dimensional subgrid M_i induced by the nodes $\{v \in V(M) \mid \pi_1(v) = i\}$. If $d = 2$, i.e., if M_i is a path, then ρ_i simply routes each routing request of $\bar{\pi}_1(R_i)$ on the path connecting its source and target in M_i .
4. We route each $(s, t) \in E(R_i)$ on the edge set consisting of dimension-1 edges connecting s to M_i , $\rho_i(\bar{\pi}_1(s), \bar{\pi}_1(t))$, and dimension-1 edges connecting t to M_i .

We can easily observe that in this algorithm, any dimension- i edge of M is contained in at most $2\ell_i$ images of ρ . Moreover, each image of ρ contains at most $2\ell_i$ dimension- i edges. I.e., ρ has an edge-congestion of $2 \cdot \max_{i \in [d]} \{\ell_i\}$ and a dilation of $2 \sum_{i=1}^d \ell_i$. This property and Lemma 2 prove Theorem 1. With our aim of using this permutation routing algorithm to prove Theorem 2, we generalize this property as the following lemma:

Lemma 3. *Let R be a routing graph on $M := M(\ell_i)_{i \in [d]}$ with $d \geq 2$ and $\ell_h := \max_{i \in [d]} \{\ell_i\}$. If $\bar{\pi}_h(R)$ is a p - q routing graph with node set $\bar{\pi}_h(V(M))$, then R can be routed on M with a dilation at most $2 \sum_{i=1}^d \ell_i$ and an edge-congestion at most $2 \cdot \max\{p, q\}$.*

PROOF. Assume without loss of generality that $h = 1$ and $\ell_1 \geq \dots \geq \ell_d$. We prove the lemma by induction on d . If $d = 2$, then R has at most $\ell_2 \cdot \max\{p, q\}$ edges. We decompose R into ℓ_1 edge-disjoint subgraphs R_1, \dots, R_{ℓ_1} so that $\bigcup_{i=1}^{\ell_1} E(R_i) = E(R)$ and $|E(R_i)| \leq \lceil \ell_2 \cdot \max\{p, q\} / \ell_1 \rceil \leq \max\{p, q\}$ for $i \in [\ell_1]$. For each $i \in [\ell_1]$, $\bar{\pi}_1(R_i)$ can be routed on the 1-dimensional grid M_i induced by the nodes $\{v \in V(M) \mid \pi_1(v) = i\}$ with a dilation at most ℓ_2 and an edge-congestion at most $\max\{p, q\}$. The routing of R is completed by adding the dimension-1 edges connecting s to M_i and t to M_i for each $(s, t) \in E(R_i)$. Any dimension-1 edge has a congestion at most $p + q$. Moreover, at most $2\ell_1$ dimension-1 edges are added to each image of the routing. Therefore, we have the lemma for $d = 2$.

If $d \geq 3$, then we color edges of R using at most $\ell_2 \cdot \max\{p, q\}$ colors so that when we identify edges in R with corresponding edges in $\bar{\pi}_2(\bar{\pi}_1(R))$, no two edges with the same sources or with the same targets in $\bar{\pi}_2(\bar{\pi}_1(R))$ have the same color. Such coloring exists because $\bar{\pi}_2(\bar{\pi}_1(R))$ is a $\ell_2 p$ - $\ell_2 q$ routing graph with node set $\bar{\pi}_2(\bar{\pi}_1(V(M)))$. Then, we decompose R into ℓ_1 edge-disjoint subgraphs R_1, \dots, R_{ℓ_1} that have edge sets with disjoint sets of $\lceil \ell_2 \cdot \max\{p, q\} / \ell_1 \rceil \leq \max\{p, q\}$ colors. This implies that $\bar{\pi}_2(\bar{\pi}_1(R_i))$ is a $\max\{p, q\}$ - $\max\{p, q\}$ routing graph with node set $\bar{\pi}_2(\bar{\pi}_1(V(M)))$ for $i \in [\ell_1]$. By induction hypothesis, $\bar{\pi}_1(R_i)$ can be routed on the $(d - 1)$ -dimensional subgrid induced by the nodes $\{v \in V(M) \mid \pi_1(v) = i\}$ with a dilation at most $2 \sum_{i=2}^d \ell_i$ and an edge-congestion at most $2 \cdot \max\{p, q\}$. The routing of R is completed by adding dimension-1 edges as done in the case of $d = 2$, so that any dimension-1 edge has congestion at most $p + q$, and at most $2\ell_1$ dimension-1 edges are added to each image of the routing. Thus, we have routed R with a dilation at most $2 \sum_{i=1}^d \ell_i$ and an edge-congestion at most $2 \cdot \max\{p, q\}$. \square

If we do not have the assumption $\ell_h = \max_{i \in [d]} \{\ell_i\}$ in Lemma 3, then we can estimate $\lceil \ell_2 \cdot \max\{p, q\} / \ell_1 \rceil \leq \lceil \mu \cdot \max\{p, q\} \rceil$ in its proof, where μ is the aspect ratio of M . This means that $|E(R_i)| \leq \lceil \mu \cdot \max\{p, q\} \rceil$ for $d = 2$, and that $\bar{\pi}_2(\bar{\pi}_1(R_i))$ is a $\lceil \mu \cdot \max\{p, q\} \rceil$ - $\lceil \mu \cdot \max\{p, q\} \rceil$ routing graph on $\bar{\pi}_2(\bar{\pi}_1(M))$ for $d \geq 3$. Therefore, initially assuming without loss of generality that $\ell_1 = \ell_h$ and $\ell_2 \geq \dots \geq \ell_d$ in the proof, we have the following lemma:

Lemma 4. *Let R be a routing graph on $M := M(\ell_i)_{i \in [d]}$ with $d \geq 2$ and aspect ratio μ . If $\bar{\pi}_h(R)$ is a p - q routing graph with node set $\bar{\pi}_h(V(M))$ for some $h \in [d]$, then R can be routed on M with a dilation at most $2 \sum_{i=1}^d \ell_i$ and an edge-congestion at most $2 \lceil \mu \cdot \max\{p, q\} \rceil$.*

4.2. Separator-Based Embedding

The following is our core theorem:

Theorem 3. *Suppose that G is a graph with N nodes, maximum node degree Δ , and with a β -decomposition tree of extension $x(n) = Cn^\alpha$ ($C > 0$, $0 \leq \alpha < 1$, $1/2 \leq \beta < 1$),*

and that M is a grid with a dimension $d \geq 2$, at least N nodes, and with constant aspect ratio. Then, G can be embedded into M with a dilation of $O(dN^{1/d})$, and with an edge-congestion of $O(dC + d^2\Delta)$ if $d > 2/(1 - \alpha)$, $O(C/(1 - \alpha - \frac{1}{d}) + d^2\Delta)$ if $1/(1 - \alpha) < d \leq 2/(1 - \alpha)$, and $O(C(N^{\alpha-1+\frac{1}{d}} + \log N) + d^2\Delta)$ if $d \leq 1/(1 - \alpha)$,

In fact, we can obtain Theorem 2 by combining Theorem 3 with Lemma 1. If G is a graph with N nodes, maximum node degree Δ , and with a β -node-separator of size $O(n^\alpha)$, then by Lemma 1, G has a $\frac{\beta}{1-\epsilon}$ -decomposition tree with extension $O(\Delta n^\alpha / \epsilon) = O(\Delta n^\alpha)$ for any $0 < \epsilon < 1 - \beta$. By Theorem 3, therefore, G can be embedded into M with a dilation of $O(dN^{\frac{1}{d}})$, and with an edge-congestion of $O(\Delta \cdot \max\{d, 1/(1 - \alpha - \frac{1}{d})\} + d^2\Delta) = O(\Delta)$ if $d > 1/(1 - \alpha)$ is fixed, $O(\Delta(N^{\alpha-1+\frac{1}{d}} + \log N) + d^2\Delta) = O(\Delta \log N)$ if $d = 1/(1 - \alpha)$, and $O(\Delta(N^{\alpha-1+\frac{1}{d}} + \log N) + d^2\Delta) = O(\Delta N^{\alpha-1+\frac{1}{d}})$ if $d < 1/(1 - \alpha)$.

We prove Theorem 3 by constructing a desired embedding algorithm, called SBE. We first outline ideas and analysis of SBE, then specify the definition of SBE, and finally prove the correctness and the edge-congestion.

Proof Sketch

We describe a proof sketch for the case $1/(1 - \alpha) < d \leq 2/(1 - \alpha)$ since the essential part of idea appears in this case. Basically, we partition G according to its decomposition tree, recursively embed the partitioned subgraphs of G into separated subgrids of the host grid $M := M(\ell_i)_{i \in [d]}$, and route cut edges, i.e., edges removed in partitioning G . In order to avoid an edge of M being used in too many recursive steps, we route the cut edges on one of edge-disjoint subgraphs of M , called *channels*. The channel associated with a positive integer w roughly equal to $\frac{1}{2}(1 - \alpha - \frac{1}{d}) \log_2 N$ is a grid-like graph homeomorphic³ to $M' := M(\frac{\ell_1}{2^w}, \frac{\ell_2}{2^w}, \ell_3, \dots, \ell_d)$ and induced by the nodes $v \in V(M)$ with $\pi_i(v) \equiv 2^{w-1} \pmod{2^w}$ for each $i = 1, 2$. We can find the channel in M as a non-empty subgraph if $d \leq 2/(1 - \alpha)$. When we embed an n -node subgraph H of G appeared in the decomposition tree, we partially route the external edges of each child of H to the channel associated with $w \simeq \frac{1}{2}(1 - \alpha - \frac{1}{d}) \log_2 n$ and route the cut edges of H by connecting the two sets of the external edges of children of H on this channel. We here say ‘‘partially’’ in two meanings: One meaning is that external edges are viewed as half-edges just leaving a child of H and are routed halfway. The other is that an external edge leaving a node in the decomposition tree is also an external edge of some descendants and is routed step by step among recursive steps. I.e., a cut edge is routed by connecting two partially routed external edges of the children, which are recursively routed using partially routed external edges of grandchildren, and so on. Consequently, cut edges of H are routed through channels associated with integers up to $\frac{1}{2}(1 - \alpha - \frac{1}{d}) \log_2 n$ in recursive steps from base embeddings to the embedding of H .

The section of M' across a dimension i is a $(d - 1)$ -dimensional grid with node set $\bar{\pi}_i(V(M'))$. If M' is associated with $w = \frac{1}{2}(1 - \alpha - \frac{1}{d}) \log_2 n$, then the minimum size S of the section is $\min_i |\bar{\pi}_i(V(M'))| = \Omega(n^{(d-1)/d} / 2^{2w}) = \Omega(n^{1-\frac{1}{d}} / 2^{(1-\alpha-\frac{1}{d}) \log_2 n}) = \Omega(n^\alpha)$. Since H and children of H have at most Cn^α external edges, we route the external edges

³A graph X is *homeomorphic* to a graph Y if some subdivision of X , i.e., a graph obtained from X by subdividing some edges of X is isomorphic to some subdivision of Y .

(or half-edges) of each child to the channel so that at most $D := Cn^\alpha/S = O(C)$ halfway points v have the same $\bar{\pi}_i(v)$, where i minimizes $|\bar{\pi}_i(V(M'))|$. We can inductively observe that this routing can be done with an edge-congestion $O(D)$ using Lemma 3. The integer w is decremented by at least 1 in $P = O(1/(1 - \alpha - \frac{1}{d}))$ recursive steps because a guest graph is partitioned into two graphs of roughly the same size in one recursive step. Therefore, the total edge-congestion incurred by entire recursive steps is at most $O(PD) = O(C/(1 - \alpha - \frac{1}{d}))$ since channels associated with different w 's are edge-disjoint. We stop the recursive procedure at the point $\min\{\ell_i\} = \Theta(d)$, by which we can obtain a base embedding with an edge-congestion $B = O(d\Delta)$ using Lemmas 2–4. Because an edge of M can possibly be contained in $O(d)$ base embeddings, the total edge-congestion is at most $O(PD + dB) = O(C/(1 - \alpha - \frac{1}{d}) + d^2\Delta)$.

The reason of the limit $\min\{\ell_i\} = \Theta(d)$ of recursive procedure is as follows: We cannot always separate the host grid with a “flat” section due to the difference between the number of nodes of a partitioned guest graph and multiples of the size of the section. In our algorithm, therefore, we separate a host grid into two subgrids that may share a $(d-1)$ -dimensional grid as “ragged” sections. Such a $(d-1)$ -dimensional grid might be used as channels in two separated grids during $O(d)$ recursive steps in the worst case, which would yield a $2^{O(d)}$ factor in the edge-congestion. To avoid this, we actually remove any boundary of a host grid from a channel, so that two separated grids have disjoint channels. However, we might have an exponential factor again if we would continue the recursive procedure until $\min\{\ell_i\}$ is much smaller than d . For instance, if $\min\{\ell_i\} = O(1)$, then removal of the boundary for each dimension would shrink the channel exponentially, implying $S = n^\alpha/2^{O(d)}$ and hence $D = O(C2^{O(d)})$.

Definition of SBE

Suppose that G_0 and M_0 are a guest graph and a host grid, respectively, satisfying the conditions of Theorem 3. Let \mathcal{T} be a β -decomposition tree for G_0 with extension Cn^α . We define a number μ as follows:

$$\mu := \max \left\{ \text{aspect ratio of } M_0, \frac{1}{1-\beta} \left(\frac{1}{7\beta} + e\beta \right) + \frac{5}{4} \right\} > 4,$$

where e is base of the natural logarithm. It should be noted that $\mu > 4$ by $\frac{1}{1-\beta}(\frac{1}{7\beta} + e\beta) + \frac{5}{4} > 4$. We assume that any proper subgrid of M_0 has less than N nodes.

We use the following notations to define SBE formally: Let $V^w := \{v \in V(M_0) \mid \pi_i(v) \equiv 2^{w-1} \pmod{2^w} \text{ for each } i = 1, 2\}$ for an integer $w \geq 1$, and let $V^0 := V(M_0)$. For a d -dimensional subgrid M of M_0 , let $W_M^w := \{v \in V^w \cap V(M) \mid \deg_M(v) = 2d\}$, where $\deg_M(\cdot)$ is the node degree in M . The *channel of* W_M^w is the subgraph of M homeomorphic to a d -dimensional grid having W_M^w as grid points. Specifically, this graph is induced by the node set $W_M^w \cup \{s \in V(M) \mid \exists i \in \{1, 2\} \exists \{u, v\} \subseteq W_M^w \pi_i(u) < \pi_i(s) < \pi_i(v) = \pi_i(u) + 2^w, \bar{\pi}_i(u) = \bar{\pi}_i(s) = \bar{\pi}_i(v)\}$. Two channels in $M(8, 8, 4)$ are illustrated in Fig. 4. It should be noted that for any $w > w' \geq 1$, $W_M^w \cap W_M^{w'} = \emptyset$ and channels of W_M^w and $W_M^{w'}$ are edge-disjoint. The *direction of* W_M^w is a dimension $i \in [d]$ minimizing $S_M^{i,w} := |\bar{\pi}_i(W_M^w)|$. In other words, the direction is a dimension of the longest side length of a grid having grid points W_M^w . In Fig. 4, the channel for $w = 1$ has direction 1 or 2 because $S_M^{1,1} = S_M^{2,1} = 6$ and $S_M^{3,1} = 9$. A mapping




Figure 4: Channels of W_M^1 and W_M^2 for $M = M(8, 8, 4)$. Black nodes are contained in W_M^1 or W_M^2 . Dashed lines represent dimension-1 and -2 edges of the channel of W_M^2 .

$\psi : X \rightarrow W_M^w$ is said to be *uniform across dimension i* if $\psi(X)$ are uniformly distributed on $\bar{\pi}_i(W_M^w)$, i.e., $\lambda_i(\psi) := \max_{v \in \bar{\pi}_i(W_M^w)} |\{s \in X \mid \bar{\pi}_i(\psi(s)) = v\}| = \lceil |X|/S_M^{i,w} \rceil$. In Fig. 4, for example, if $\psi : [4] \rightarrow W_M^1$ maps 1, 2, 3, 4 to (3, 3, 2), (3, 5, 2), (5, 3, 3), (5, 5, 3), respectively, then ψ is uniform across dimensions both 1 and 3 but not dimension 2 because $\lambda_1(\psi) = 1 = \lceil \lceil [4] \rceil / S_M^{1,1} \rceil = \lceil 4/6 \rceil$, $\lambda_2(\psi) = 2 > \lceil \lceil [4] \rceil / S_M^{2,1} \rceil = \lceil 4/6 \rceil$, and $\lambda_3(\psi) = 1 = \lceil \lceil [4] \rceil / S_M^{3,1} \rceil = \lceil 4/9 \rceil$. We note here that for any two dimensions i and j , we can construct ψ uniform across dimensions both i and j by uniformly distributing $\psi(X)$ among nodes on a $(d - 1)$ -dimensional diagonal hyperplane between dimensions i and j in W_M^w .

Step 0—Input and Output

The formal input and output of SBE is as follows:

Algorithm SBE(G, X, M, U).

Input

- An n -node subgraph G of G_0 contained in $V(\mathcal{T})$.
- A multiset X of nodes of G incident to distinct external edges of G , i.e., a node appears in X as many times as the number of the external edges incident to the node.
- A subgrid $M = M(\ell_i)_{i \in [d]}$ of M_0 with aspect ratio at most μ , together with a set $U \subseteq V(M)$ such that $U \supseteq \{v \in V(M) \mid \deg_M(v) = 2d\}$ and $|U| = n$. Suppose that h is a dimension such that $\ell_h = \max_{i \in [d]} \{\ell_i\}$.

Output

- An embedding $\langle \phi, \rho \rangle$ of G into M such that $\phi(V(G)) = U$.




Figure 5: An example of separation a host grid M with 50 nodes in U (represented by black nodes) into a grid M_1 with $|U_1| = 20$ and a grid M_2 with $|U_2| = 30$. We choose $h = 2$ here.

- A mapping $\psi : X \rightarrow W_M^w$ uniform across the direction k of W_M^w , where $w \geq 0$ is an integer defined in Step 1.
- A routing σ of the routing graph with node set $\phi(X) \cup \psi(X)$ and edge set $\{(\phi(u), \psi(u)) \mid u \in X\}$.

Initially, we arbitrarily choose U as desired and perform $\text{SBE}(G_0, \emptyset, M_0, U)$.

Step 1—Channel Configuration

This step sets an integer w , by which we configure the channel of W_M^w to route ρ and σ . We define $w := \max\{\lfloor \frac{1}{2}((1 - \tilde{\alpha} - \frac{1}{d}) \log_2 n - \log_2 \frac{\mu}{1-\beta}) \rfloor, 0\}$, where $\tilde{\alpha} := \max\{1 - \frac{2}{d}, \alpha\}$. In other words, $w = \max\{\lfloor \frac{1}{2}(\frac{1}{d} \log_2 n - \log_2 \frac{\mu}{1-\beta}) \rfloor, 0\}$ if $d > 2/(1 - \alpha)$, $w = \max\{\lfloor \frac{1}{2}((1 - \alpha - \frac{1}{d}) \log_2 n - \log_2 \frac{\mu}{1-\beta}) \rfloor, 0\}$ if $1/(1 - \alpha) < d \leq 2/(1 - \alpha)$, and $w = 0$ if $d \leq 1/(1 - \alpha)$.

Step 2—Base Embedding

If $\ell_h \leq 2\mu d$, then SBE does not call itself recursively any longer and constructs a base embedding as follows:

1. If $X = \emptyset$, then let $\phi : V(G) \rightarrow U$ be an arbitrary one-to-one mapping. Otherwise, let Y be the set (not a multiset) of nodes incident to external edges of G . We construct a one-to-one mapping $\phi : V(G) \rightarrow U$ so that degrees of nodes in Y are uniformly distributed on $\bar{\pi}_k(U)$. Specifically, for any $v \in \bar{\pi}_k(U)$, the sum of degrees of nodes $s \in Y$ with $\bar{\pi}_k(\phi(s)) = v$ is at most $e|X|/|\bar{\pi}_k(U)| + \Delta$. We prove later in Lemma 5 that ϕ can be constructed as desired. This construction implies that if R is the routing graph to be routed by σ , then $\bar{\pi}_k(R)$ has maximum outdegree at most $e|X|/|\bar{\pi}_k(U)| + \Delta \leq e|X|/S_M^{k,w} + \Delta$.
2. Construct a mapping $\psi : X \rightarrow W_M^w$ so that ψ is uniform across dimension k , implying that $\bar{\pi}_k(R)$ has maximum indegree at most $\lceil |X|/S_M^{k,w} \rceil$.
3. Apply Lemmas 2 and 3 on M to obtain ρ .
4. If $X \neq \emptyset$, then apply Lemma 4 on M to obtain σ .
5. Return.




Figure 6: Mapping $\tilde{\psi}_j$ and routing $\tilde{\sigma}_j$ constructed in Step 4. Symbols ϕ_j , σ_j , and ψ_j are the output ϕ , σ , and ψ in the embedding of M_j , respectively. The horizontal and diagonal planes represent the uniformness of $\psi_j(X_j)$ and $\tilde{\psi}_j(X_j)$, respectively.

Step 3—Separation

Suppose $\ell_h > 2\mu d$. Let G_1 and G_2 be children of G in \mathcal{T} and have n_1 and n_2 nodes, respectively. Now M is separated into two subgrids $M_1 := M(\ell_1, \dots, \ell_{h-1}, m_1, \ell_{h+1}, \dots, \ell_d)$ and $M_2 := M(\ell_1, \dots, \ell_{h-1}, m_2, \ell_{h+1}, \dots, \ell_d)$, together with node sets $U_1 \subseteq V(M_1)$ and $U_2 \subseteq V(M_2)$ such that $m_1 + m_2 = \ell_h + 1$, $U_1 \cup U_2 = U$, $U_1 \cap U_2 = \emptyset$, and $|U_j| = n_j$ for $j = 1, 2$ (cf. Fig. 5). We here duplicate the $(d-1)$ -dimensional grid induced by $\{v \in V(M) \mid \pi_h(v) = m_1\}$ to be shared by M_1 and M_2 , so that $m_1 + m_2$ equals not ℓ_h but $\ell_h + 1$. We have to do this because M_1 and M_2 must have enough numbers of nodes in U onto which $V(G_1)$ and $V(G_2)$ can be mapped, respectively. In Fig. 5, actually, no matter how we partitioned M so that $m_1 + m_2 = \ell_h$ with $1 \leq h \leq 3$, we would have either $|V(M_1) \cap U| < 20$ or $|V(M_2) \cap U| < 30$. We will prove later in Lemma 6 that the resulting subgrids have aspect ratio at most μ .

Step 4—Recursive Embedding

This step recursively embeds G_1 and G_2 into M_1 and M_2 , respectively. We also construct a routing $\tilde{\sigma}_j$ for $j = 1, 2$, which draws the external edges of G_j to $\tilde{\psi}_j(X_j)$. Here, X_j is the multiset of nodes of G_j incident to distinct external edges of G_j , and $\tilde{\psi}_j : X_j \rightarrow W_{M_j}^w$ is a mapping uniform across dimension k , i.e., the direction of W_M^w . With this routing, in the subsequent step, we will make a routing graph with sources and targets in $\tilde{\psi}_1(X_1)$ and $\tilde{\psi}_2(X_2)$, respectively, on the channel of W_M^w to connect cut edges. Projecting this routing graph along dimension k yields a $\lceil |X_1|/S_M^{k,w} \rceil - \lceil |X_2|/S_M^{k,w} \rceil$ routing graph. Thus, the cut edges will be routed on the channel of W_M^w with an edge-congestion of $2 \cdot \max_{j=1,2} \lceil |X_j|/S_M^{k,w} \rceil$ using Lemma 3. In applying Lemma 3, we regard the underlying channel as the homeomorphic grid with grid points W_M^w . A detailed analysis for this edge-congestion will later be provided in Lemmas 8-10. We aim to suppress the edge-congestion of $\tilde{\sigma}_j$ in a similar way. Therefore, we make $\tilde{\psi}_j$ uniform across not only dimension k but also the direction k_j of $W_{M_j}^{w_j}$, where w_j is w computed for n_j in the recursive procedure, since ψ_j is made uniform across dimension k_j in the recursive procedure (cf. Fig. 6). We need to treat two more matters in constructing $\tilde{\sigma}_j$.

First, we need a channel (a grid-like graph) containing both $\psi_j(X_j) \subseteq W_{M_j}^w$ and $\tilde{\psi}_j(X_j) \subseteq W_{M_j}^{w_j}$. Just taking union of two channels of $W_{M_j}^{w_j}$ and $W_{M_j}^w$ would not suffice because $W_{M_j}^{w_j}$ and $W_{M_j}^w$ are disjoint if $w_j \neq w$. We define the channel for w_j and w as

the graph homeomorphic to a d -dimensional grid having $W_{M_j}^{w_j, w} := \{v \in V(M_0) \mid \pi_i(v) \equiv 2^{w_j-1} \text{ or } 2^{w-1} \pmod{2^w} \text{ for } i = 1, 2\} \cap \{v \in V(M_j) \mid \deg_{M_j}(v) = 2d\}$ as grid points. It should be noted that if $w > w_j > 0$, then edges contained in the channel of $W_{M_j}^{w_j, w}$ but in the channel of neither $W_{M_j}^{w_j}$ nor $W_{M_j}^w$ are uniquely determined by w_j and w and not contained any other channel in M_j except the channel of $W_{M_j}^0$.

The second matter is that the direction of $W_{M_j}^{w_j, w}$ may differ from k_j , across which ψ_j and $\tilde{\psi}_j$ are uniform. This means that just applying Lemma 3 would not guarantee a desired edge-congestion. If we applied the algorithm of Lemma 3 on the channel of $W_{M_j}^{w_j, w}$, then the algorithm would be recursively called in the non-increasing order of side lengths of a grid having $W_{M_j}^{w_j, w}$ as grid points. We here modify the order by replacing $W_{M_j}^{w_j, w}$ with $W_{M_j}^{w_j}$. The modified algorithm yields a desired edge-congestion as we will later prove in Lemma 8.

For each $j = 1, 2$, specifically, SBE performs the following procedures:

1. Call $\text{SBE}(G_j, X_j, M_j, U_j)$. Let ϕ_j, ρ_j, ψ_j , and σ_j denote the output ϕ, ρ, ψ , and σ of the recursive call, respectively.
2. Construct a mapping $\tilde{\psi}_j : X_j \rightarrow W_{M_j}^w$ uniform across dimensions both k and k_j , where k_j is the direction of $W_{M_j}^{w_j}$, and $w_j := \max\{\lfloor \frac{1}{2}((1 - \tilde{\alpha} - \frac{1}{d}) \log_2 n_j - \log_2 \frac{\mu}{1-\beta}) \rfloor, 0\}$.
3. Let $\tilde{\sigma}_j$ be a routing from $\psi_j(X_j)$ to $\tilde{\psi}_j(X_j)$ on the channel of $W_{M_j}^{w_j, w}$ obtained by using the modified algorithm of Lemma 3 in which we recursively call the algorithm in the non-increasing order of side lengths of a grid having $W_{M_j}^{w_j}$ as grid points.

Step 5—Routing Cut and External Edges

This step constructs ψ , then completes ρ and σ using $\tilde{\psi}_j$ and ψ . The routings ρ and σ are obtained simply using Lemma 3 on the channel of W_M^w since $\tilde{\psi}_j(X_j) \subseteq W_{M_j}^w \subseteq W_M^w$. The following are specific procedures of this step:

1. Construct a mapping $\psi : X \rightarrow W_M^w$ uniform across dimension k .
2. By using Lemma 3, construct $\tilde{\sigma}$ for the routing graph on the channel of W_M^w with node set $\tilde{\psi}_1(X_1) \cup \tilde{\psi}_2(X_2)$ and edge set $\{(\tilde{\psi}_1(s_1), \tilde{\psi}_2(s_2)) \mid s_1 \in X_1 \setminus X, s_2 \in X_2 \setminus X, \{s_1, s_2\} \in E(G)\} \cup \bigcup_{j=1,2} \{(\tilde{\psi}_j(s), \psi(s)) \mid s \in X_j \cap X\}$ (cf. Fig. 7). It should be noted that $\{s_1, s_2\} \in E(G)$ with $s_1 \in X_1 \setminus X$ and $s_2 \in X_2 \setminus X$ is a cut edge of G , and $s \in X_j \cap X$ is a node incident to an external edge of G .
3. Let ρ map the cut edges of G onto paths obtained by concatenating the images of $\sigma_1, \tilde{\sigma}_1, \tilde{\sigma}, \tilde{\sigma}_2$, and σ_2 . Specifically, for $\{s_1, s_2\} \in E(G)$ with $s_1 \in X_1 \setminus X$ and $s_2 \in X_2 \setminus X$, let

$$\rho(\phi(s_1), \phi(s_2)) := \bigcup_{j=1,2} (\sigma_j(\phi(s_j), \psi_j(s_j)) \cup \tilde{\sigma}_j(\psi_j(s_j), \tilde{\psi}_j(s_j))) \cup \tilde{\sigma}(\tilde{\psi}_1(s_1), \tilde{\psi}_2(s_2)).$$




Figure 7: Mapping ψ and routing $\tilde{\sigma}$ constructed in Step 5. The left, center, and right planes represent the uniformness of $\tilde{\psi}_1(X_1)$, $\psi(X)$, and $\tilde{\psi}_2(X_2)$, respectively.

4. Let σ map the external edges of G onto paths obtained by concatenating the images of σ_j , $\tilde{\sigma}_j$, $\tilde{\sigma}$. Specifically, for $s \in X_j \cap X$ ($j = 1, 2$), let

$$\sigma(\phi(s), \psi(s)) := \sigma_j(\phi(s), \psi_j(s)) \cup \tilde{\sigma}_j(\psi_j(s), \tilde{\psi}_j(s)) \cup \tilde{\sigma}(\tilde{\psi}_j(s), \psi(s)).$$

Correctness

To prove that SBE yields an output satisfying the conditions specified in Step 0, we first prove in Lemma 5 below that we can construct ϕ as desired in Step 2 for the case $X \neq \emptyset$. We then prove in Lemmas 6 and 7 below that M_j defined in Step 3 has aspect ratio at most μ , and that $W_{M_j}^w$ is non-empty as well as W_M^w . These facts guarantee that a valid input is given to a child procedure in Step 4, and that mappings $\tilde{\psi}_j$ and ψ can be constructed in Steps 4 and 5, respectively.

Lemma 5. *Suppose $X \neq \emptyset$ in Step 2. We can construct a one-to-one mapping $\phi : V(G) \rightarrow U$ such that for any $v \in \bar{\pi}_k(U)$, the sum of degrees of nodes $s \in Y$ with $\bar{\pi}_k(\phi(s)) = v$ is at most $e|X|/|\bar{\pi}_k(U)| + \Delta$, where Y is the set (not a multiset) of nodes incident to external edges of G .*

PROOF. If $U = V(M)$, then we can map Y onto U in a trivial manner so that $\max_{v \in \bar{\pi}_k(U)} \sum_{s \in Y, \bar{\pi}_k(\phi(s))=v} \deg_G(s) \leq \sum_{s \in Y} \deg_G(s) / |\bar{\pi}_k(U)| + \Delta = |X| / |\bar{\pi}_k(U)| + \Delta$. This is because this mapping can be viewed as a packing of $|Y|$ items of size at most Δ to $|\bar{\pi}_k(U)|$ bins that can contain the same number $|U| / |\bar{\pi}_k(U)| = \ell_k$ of items.

If $U \subset V(M)$, i.e., U does not contain some nodes on the boundary of M , then some of the bins cannot contain ℓ_k items. An upper bound can be obtained in the assumption that U contains no node on the boundary of M , and that we must map Y onto $\prod_{i \in [d] \setminus \{k\}} (\ell_i - 2)$ bins, in which the mapping is not one-to-one if $|Y| > \prod_{i \in [d]} (\ell_i - 2)$. Thus, we have

$$\max_{v \in \bar{\pi}_k(U)} \sum_{\substack{s \in Y \\ \bar{\pi}_k(\phi(s))=v}} \deg_G(s) \leq \frac{\sum_{s \in Y} \deg_G(s)}{\prod_{i \in [d] \setminus \{k\}} (\ell_i - 2)} + \Delta \leq \frac{|X|}{\prod_{i \in [d] \setminus \{k\}} (\ell_i - 2)} \cdot \frac{\prod_{i \in [d] \setminus \{k\}} \ell_i}{|\bar{\pi}_k(U)|} + \Delta. \quad (5)$$

We have $X \neq \emptyset$ only if the current base embedding is called by a parent procedure. This implies $\min_{i \in [d]} \ell_i > 2d$ as proved in Lemma 6 below, and therefore,

$$\frac{\prod_{i \in [d] \setminus \{k\}} \ell_i}{\prod_{i \in [d] \setminus \{k\}} (\ell_i - 2)} = \prod_{i \in [d] \setminus \{k\}} \frac{\ell_i}{\ell_i - 2} < \left(\frac{2d}{2d-2} \right)^{d-1} < e.$$

Combined with (5), we have the lemma. \square

Lemma 6. For $j = 1, 2$, M_j defined in Step 3 has aspect ratio at most μ and $\min\{\ell_1, \dots, \ell_{h-1}, m_j, \ell_{h+1}, \dots, \ell_d\} > 2d$.

PROOF. Assume without loss of generality that $m_1 \leq m_2$. Since M has aspect ratio at most μ , we have $\ell_h / \min_{i \in [d]} \ell_i \leq \mu$. Moreover, it follows that $\min_{i \in [d]} \ell_i > 2d$, for otherwise, $\ell_h \leq \mu \cdot \min_{i \in [d]} \ell_i \leq 2\mu d$, and hence, SBE entered the base step. Therefore, it suffices to prove that $m_1 > 2d$ and $\ell_h / m_1 \leq \mu$.

Because $n_j = |U_j| \geq |\{v \in V(M_j) \mid \deg_{M_j}(v) = 2d\}|$ for $j = 1, 2$, it follows that $(m_j - 2) \prod_{i \in [d] \setminus \{h\}} (\ell_i - 2) \leq n_j \leq m_j \prod_{i \in [d] \setminus \{h\}} \ell_i$, and hence,

$$m_j - 2 \leq \frac{n_j}{\prod_{i \in [d] \setminus \{h\}} (\ell_i - 2)} \leq m_j \prod_{i \in [d] \setminus \{h\}} \frac{\ell_i}{\ell_i - 2} < m_j \left(\frac{2d}{2d-2} \right)^{d-1} < em_j.$$

We have by the inequalities that

$$\frac{m_2 - 2}{n_2} \leq \frac{1}{\prod_{i \in [d] \setminus \{h\}} (\ell_i - 2)} < \frac{em_1}{n_1}. \quad (6)$$

Because $n - n_1 = n_2 \leq \lceil \beta n \rceil \leq \beta n + 1$, it follows that

$$n_1 \geq (1 - \beta)n - 1 \geq (1 - \beta) \frac{n_2 - 1}{\beta} - 1 = \frac{(1 - \beta)n_2}{\beta} - \frac{1}{\beta},$$

by which we obtain

$$n_2 \leq \frac{\beta n_1 + 1}{1 - \beta}. \quad (7)$$

Moreover, it follows that $n \geq \prod_{i \in [d]} (\ell_i - 2) > (2\mu d - 2)(2d - 2)^{d-1} \geq 8\mu - 4$, which is larger than 7μ because $\mu > 4$. Furthermore, $\mu > \frac{1}{1-\beta}(\frac{1}{7} + e\beta)$ by the definition of μ . Hence, it follows that

$$n_1 \geq (1 - \beta)n - 1 > 7\mu(1 - \beta) - 1 > 7e\beta. \quad (8)$$

Thus, by (6)–(8) and $\mu \geq \frac{1}{1-\beta}(\frac{1}{7\beta} + e\beta) + \frac{5}{4}$,

$$m_2 - 2 < \frac{em_1 n_2}{n_1} \leq \frac{em_1}{1 - \beta} \left(\beta + \frac{1}{n_1} \right) < \frac{em_1}{1 - \beta} \left(\beta + \frac{1}{7e\beta} \right) \leq \left(\mu - \frac{5}{4} \right) m_1,$$

by which we obtain $(\mu - \frac{1}{4})m_1 > m_1 + m_2 - 2 = \ell_h - 1 > 2\mu d - 1$. Therefore,

$$m_1 > \frac{2\mu d - 1}{\mu - \frac{1}{4}} = \frac{2d(\mu - \frac{1}{2d})}{\mu - \frac{1}{4}} \geq 2d, \text{ and}$$

$$\frac{\ell_h}{m_1} < \mu - \frac{1}{4} + \frac{1}{m_1} < \mu - \frac{1}{4} + \frac{1}{2d} \leq \mu.$$

\square

Lemma 7. For $j = 1, 2$, $W_{M_j}^w$ in Step 4 is non-empty.

PROOF. Suppose $M_j = M(\ell_i^j)_{i \in [d]}$ and $\ell_{\min}^j := \min_{i \in [d]} \ell_i^j$. We can observe by the definition of $W_{M_j}^w$ that if $\lfloor (\ell_i^j - 2)/2^w \rfloor > 0$ for $i = 1, 2$, and if $\ell_i^j - 2 > 0$ for $3 \leq i \leq d$, then $W_{M_j}^w$ is non-empty. Because $\ell_{\min}^j > 2d > 4$ by Lemma 6, the lemma holds if $w = 0$. Assume $w \geq 1$. Then, the lemma is implied by $\ell_{\min}^j/2^w \geq 2$. The assumption $w \geq 1$ implies that $d > 1/(1 - \alpha)$ by the definition of w , and that

$$2^w \leq \left(\frac{(1 - \beta)n^{1 - \frac{1}{d}}}{\mu} \right)^{1/2} = \left(\frac{(1 - \beta)n^{\min(1/d, 1 - \frac{1}{d})}}{\mu} \right)^{1/2} \leq \left(\frac{(1 - \beta)n^{1/d}}{\mu} \right)^{1/2}. \quad (9)$$

As estimated in (7) and (8), it follows that $n_1 \geq (1 - \beta)n - 1$ and $n > 7\mu$. It also follows that $n_2 \geq (1 - \beta)n - 1$, since $n - n_2 = n_1 \leq \lceil \beta n \rceil \leq \beta n + 1$. Moreover, $\mu > \frac{1}{1 - \beta} \cdot \frac{1}{7\beta}$. Therefore,

$$n_j \geq \left(1 - \beta - \frac{1}{n}\right)n > \left(1 - \beta - \frac{1}{7\mu}\right)n > (1 - \beta - (1 - \beta)\beta)n = (1 - \beta)^2 n. \quad (10)$$

Because M_j has aspect ratio at most μ by Lemma 6, it follows that

$$n_j^{1/d} \leq \max_{i \in [d]} \{\ell_i^j\} \leq \mu \ell_{\min}^j. \quad (11)$$

Combining (9)–(11), and by $d \geq 2$,

$$\frac{\ell_{\min}^j}{2^w} \geq \frac{n_j^{1/d}}{\mu} \left(\frac{\mu}{(1 - \beta)n^{1/d}} \right)^{1/2} > \left(\frac{(1 - \beta)^{\frac{4}{d} - 1} n^{1/d}}{\mu} \right)^{1/2} \geq \left(\frac{(1 - \beta)n^{1/d}}{\mu} \right)^{1/2} \geq 2^w \geq 2.$$

□

Edge-Congestion

We first estimate the edge-congestion of $\tilde{\sigma}_j$ and $\tilde{\sigma}$ in each recursive call of SBE. Then, we prove the total edge-congestion. In what follows, for an n -node guest graph given to SBE as input, we will use $D^w(n)$ to denote the maximum value of $\max_{i \in [d]} \lceil Cn^\alpha / S_H^{i,w} \rceil$ over all feasible d -dimensional host grids H .

Lemma 8. For $j = 1, 2$, $\tilde{\sigma}_j$ in Step 4 imposes an edge-congestion at most $2D^w(n_j)$ on the channel of $W_{M_j}^{w_j, w}$.

PROOF. As described in Step 4, in constructing $\tilde{\sigma}_j$, we recursively call the algorithm of Lemma 3 in the non-increasing order of side lengths of a grid having not $W_{M_j}^{w_j, w}$ but $W_{M_j}^{w_j}$ as grid points. We can prove through a similar argument to that of Lemma 3 that the modified algorithm achieves an edge-congestion of $2\lceil |X_j| / S_{M_j}^{k_j, w} \rceil \leq 2\lceil Cn_j^\alpha / S_{M_j}^{k_j, w} \rceil \leq 2D^w(n_j)$, noting that $W_{M_j}^{w_j} \cup W_{M_j}^w \subseteq W_{M_j}^{w_j, w}$ and $|\pi_i(W_{M_j}^{w_j})| \geq |\pi_i(W_{M_j}^w)|$ for $i \in [d]$. The following is an explicit proof.

Suppose that grids having $W_{M_j}^{w_j}$ and $W_{M_j}^{w_j, w}$ as grid points are $M' := M(\ell'_1, \dots, \ell'_d)$ and $M'' := M(\ell''_1, \dots, \ell''_d)$, respectively. I.e., $\ell'_i = |\pi_i(W_{M_j}^{w_j})|$ and $\ell''_i = |\pi_i(W_{M_j}^{w_j, w})|$ for $i \in [d]$. Because $W_{M_j}^{w_j} \subseteq W_{M_j}^{w_j, w}$, $\ell'_i \leq \ell''_i$ for each $i \in [d]$. In this proof, we assume without loss of generality that $k_j = 1$ and $\ell'_1 = \ell'_{k_j} \geq \ell'_2 \cdots \geq \ell'_d$. Moreover, we regard the routing graph R from sources $\psi_j(X_j)$ to targets $\tilde{\psi}_j(X_j)$ on the channel of $W_{M_j}^{w_j, w}$ as its corresponding routing graph on M'' . Then, $\bar{\pi}_1(R)$ is a p - q routing graph with node set $\bar{\pi}_1(V(M''))$, where $p := \lceil |X_j|/S_{M_j}^{k_j, w_j} \rceil = \lceil |X_j|/|\bar{\pi}_{k_j}(W_{M_j}^{w_j})| \rceil$ and $q := \lceil |X_j|/S_{M_j}^{k_j, w} \rceil = \lceil |X_j|/|\bar{\pi}_{k_j}(W_{M_j}^w)| \rceil$, because $\psi_j : X_j \rightarrow W_{M_j}^{w_j}$ and $\tilde{\psi}_j : X_j \rightarrow W_{M_j}^w$ are uniform across dimension $k_j = 1$. Because $w_j \leq w$, it follows that $|\pi_i(W_{M_j}^{w_j})| \geq |\pi_i(W_{M_j}^w)|$ and $|\bar{\pi}_i(W_{M_j}^{w_j})| \geq |\bar{\pi}_i(W_{M_j}^w)|$ for $i \in [d]$. This implies that $p \leq q$.

If $d = 2$, then R has $|X_j| \leq \ell'_2 p \leq \ell'_2 q$ edges. Therefore, we can decompose R into ℓ'_1 edge-disjoint subgraphs $R_1, \dots, R_{\ell'_1}$ so that $\bigcup_{i=1}^{\ell'_1} E(R_i) = E(R)$ and $|E(R_i)| \leq \lceil \ell'_2 q / \ell'_1 \rceil \leq q$ for $i \in [\ell'_1]$. Since $\ell''_1 \geq \ell'_1$, $\bar{\pi}_1(R_i)$ can be routed with an edge-congestion at most q on the 1-dimensional subgrid of M'' induced by the nodes $\{v \in V(M'') \mid \pi_1(v) = i\}$ for each $1 \leq i \leq \ell'_1 \leq \ell''_1$. Thus, we can route R on M'' with an edge-congestion at most $2q$ as in the proof of Lemma 3.

If $d \geq 3$, then since $\ell'_2 = |\pi_2(W_{M_j}^{w_j})| \geq |\pi_2(W_{M_j}^w)|$, $\bar{\pi}_2(\bar{\pi}_1(R))$ is an $\ell'_2 p$ - $\ell'_2 q$ routing graph with node set $\bar{\pi}_2(\bar{\pi}_1(V(M'')))$. Using an edge-coloring described in the proof of Lemma 3, therefore, we can decompose R into ℓ'_1 edge-disjoint subgraphs $R_1, \dots, R_{\ell'_1}$ such that $\bar{\pi}_2(\bar{\pi}_1(R_i))$ is a $\max\{p, q\}$ - $\max\{p, q\}$ routing graph with node set $\bar{\pi}_2(\bar{\pi}_1(V(M'')))$. Since $\ell''_1 \geq \ell'_1$, $\bar{\pi}_1(R_i)$ can inductively be routed with an edge-congestion at most $2 \cdot \max\{p, q\} = 2q$ on the $(d-1)$ -dimensional subgrid induced by the nodes $\{v \in V(M'') \mid \pi_1(v) = i\}$ for each $1 \leq i \leq \ell'_1 \leq \ell''_1$. Thus, we can route R on M'' with an edge-congestion at most $2q$ as in the proof of Lemma 3. \square

Lemma 9. *The routing $\tilde{\sigma}$ in Step 5 imposes an edge-congestion at most $2 \cdot \max\{2D^w(n_j), D^w(n_j) + D^w(n)\}$ on the channel of W_M^w .*

PROOF. Because $\tilde{\psi}_1, \tilde{\psi}_2$, and ψ are uniform across dimension k , it follows that $\lambda_k(\tilde{\psi}_j) = \lceil |X_j|/S_{M_j}^{k, w} \rceil \leq \lceil Cn_j^\alpha/S_{M_j}^{k, w} \rceil \leq D^w(n_j)$ for $j = 1, 2$, and that $\lambda_k(\psi) = \lceil |X|/S_M^{k, w} \rceil \leq \lceil Cn^\alpha/S_M^{k, w} \rceil \leq D^w(n)$. Therefore, if R is the routing graph for $\tilde{\sigma}$ on the channel of W_M^w , then $\bar{\pi}_k(R)$ has maximum outdegree at most $\lambda_k(\tilde{\psi}_1) + \lambda_k(\tilde{\psi}_2) \leq 2D^w(n_j)$ and maximum indegree at most $\lambda_k(\tilde{\psi}_2) + \lambda_k(\psi) \leq D^w(n_j) + D^w(n)$ as shown in Fig. 7. By Lemma 3, therefore, $\tilde{\sigma}$ has a desired edge-congestion. \square

Lemma 10. *For $j = 1, 2$, it follows that*

$$\max_{i \in [d]} \left\lfloor \frac{Cn_j^\alpha}{S_{M_j}^{i, w}} \right\rfloor = \begin{cases} O(C) & \text{if } d > 1/(1-\alpha), \\ O\left(Cn_j^{\alpha-1+\frac{1}{d}}\right) & \text{otherwise.} \end{cases}$$

PROOF. Suppose that $M_j = M(\ell_i^j)_{i \in [d]}$, $\ell_{\max}^j := \max_{i \in [d]} \{\ell_i^j\}$, and $\ell_{\min}^j := \min_{i \in [d]} \{\ell_i^j\}$. We begin with bounds of ℓ_{\max}^j and $\lfloor (\ell_i - 2)/2^w \rfloor$. Because $\ell_{\min}^j > 2d$ by Lemma 6, it follows

that

$$n_j \geq \prod_{i \in [d]} (\ell_i^j - 2) = \prod_{i \in [d]} \left(1 - \frac{2}{\ell_i^j}\right) \ell_i^j > \left(1 - \frac{1}{d}\right)^d \left(\frac{\ell_{\max}^j}{\mu}\right)^d \geq \left(\frac{\ell_{\max}^j}{2\mu}\right)^d,$$

yielding

$$\ell_{\max}^j < 2\mu n_j^{1/d}. \quad (12)$$

It follows from the proof of Lemma 7 that $\ell_{\min}^j/2^w \geq 2$. Therefore,

$$\left\lfloor \frac{\ell_i^j - 2}{2^w} \right\rfloor \geq \frac{\ell_i^j - 2^w - 1}{2^w} \geq \frac{\ell_i^j - \frac{\ell_{\min}^j}{2} - 1}{2^w} \geq \frac{\ell_i^j - 2}{2^{w+1}}. \quad (13)$$

If k' is the direction of $W_{M_j}^w$, then it follows from (12) and (13) that

$$\begin{aligned} \min_{i \in [d]} \{S_{M_j}^{i,w}\} &\geq \left(\prod_{i \in \{1,2\} \setminus \{k'\}} \left\lfloor \frac{\ell_i^j - 2}{2^w} \right\rfloor \right) \left(\prod_{i \in [d] \setminus \{1,2,k'\}} (\ell_i^j - 2) \right) \geq \frac{1}{2^{2w+2}} \prod_{i \in [d] \setminus \{k'\}} (\ell_i^j - 2) \\ &= \frac{1}{2^{2w+2}} \prod_{i \in [d] \setminus \{k'\}} \left(1 - \frac{2}{\ell_i^j}\right) \ell_i^j > \left(1 - \frac{1}{d}\right)^{d-1} \frac{\prod_{i \in [d] \setminus \{k'\}} \ell_i^j}{2^{2w+2}} > \frac{n_j / \ell_{\max}^j}{4e2^{2w}} > \frac{n_j^{1-\frac{1}{d}}}{8e\mu 2^{2w}}. \end{aligned}$$

Therefore, we have $\max_{i \in [d]} \lceil Cn_j^\alpha / S_{M_j}^{i,w} \rceil < \lceil 8e\mu 2^{2w} Cn_j^{\alpha-1+\frac{1}{d}} \rceil$. If $w = 0$, then the lemma is immediate. If $w \geq 1$, which implies $d > 1/(1-\alpha)$, then it follows from inequalities in (9) and (10) that

$$2^{2w} \leq \frac{(1-\beta)n^{\min\{1/d, 1-\alpha-\frac{1}{d}\}}}{\mu} \leq \frac{(1-\beta)n^{1-\alpha-\frac{1}{d}}}{\mu} \leq \frac{(1-\beta)^{-1+2\alpha+\frac{2}{d}} n_j^{1-\alpha-\frac{1}{d}}}{\mu}.$$

Therefore, we have $\max_{i \in [d]} \lceil Cn_j^\alpha / S_{M_j}^{i,w} \rceil < \lceil 8e(1-\beta)^{-1+2\alpha+\frac{2}{d}} C \rceil = O(C)$. \square

Through an analysis similar to that of Lemma 10, we can prove that $\lceil Cn^\alpha / S_M^{k,w} \rceil$ is $O(C)$ if $d > 1/(1-\alpha)$, and $O(Cn^{\alpha-1+\frac{1}{d}})$ otherwise. These upper bounds are independent of w and hold for any feasible guest and host graphs processed in each step of SBE unless the minimum side length of the host grid is at most $2d$. In what follows, therefore, we denote $D^w(\cdot)$ simply by $D(\cdot)$ and use these upper bounds of $D(\cdot)$ on the condition that the host grid has the minimum side length larger than $2d$.

Lemma 11. *The edge-congestion B of the base embedding in Step 2 is $O(d\Delta + C)$.*

PROOF. The edge-congestion of ρ constructed in the base embedding is at most $2\lceil \frac{\Delta}{2} \rceil \ell_h$ by Lemmas 2 and 3. If $X \neq \emptyset$, then the edge-congestion of σ is at most with an edge-congestion at most $2\lceil \mu(e|X|/S_M^{k,w} + \Delta) \rceil \leq 2\lceil \mu(eD(n) + \Delta) \rceil$. This bound is obtained from Lemma 4 and the fact that for the routing graph R to be routed by σ , $\bar{\pi}_k(R)$ has maximum outdegree at most $e|X|/S_M^{k,w} + \Delta$ by Lemma 5 and maximum indegree at most $\lceil |X|/S_M^{k,w} \rceil$ as mentioned in Step 2. Because $n^{1/d} \leq \ell_h \leq 2\mu d$ in the base embedding and $\min_{i \in [d]} \ell_i > 2d$ as described in the proof of Lemma 5, we have $B \leq 2(\lceil \frac{\Delta}{2} \rceil \ell_h + \lceil \mu(eD(n) + \Delta) \rceil) = O(d\Delta + C)$ if $d > 1/(1-\alpha)$, and $B = O(d\Delta + Cn^{\alpha-1+\frac{1}{d}}) = O(d\Delta + C(2\mu d)^{d(\alpha-1)+1}) = O(d\Delta + C)$ otherwise. \square

We now estimate the total congestion of a fixed edge r of M_0 . If r is contained in the channel of $W_{M_0}^{w,w'}$ but of neither $W_{M_0}^w$ nor $W_{M_0}^{w'}$ for a certain unique pair of $w > 0$ and $w' > 0$, then r incurs a congestion less than that in the case that it is contained in the channel of either $W_{M_0}^w$ or $W_{M_0}^{w'}$. This is because the channel of $W_{M_0}^{w,w'}$ is used only in Step 4, while the channels of $W_{M_0}^w$ and $W_{M_0}^{w'}$ are used in other steps as well. To analyze an upper bound of congestion of r , therefore, it suffices to assume that r is contained in the channel of $W_{M_0}^{w_r}$ with $w_r \geq 1$ uniquely determined by r , as well as in the channel of $W_{M_0}^0$ and some base embeddings.

Lemma 12. *The number P_1 of recursive calls of SBE that set $w \geq 1$ in Step 1, perform inductive steps (i.e., not a base embedding), and use a channel containing r is $O(d)$ if $d > 2/(1 - \alpha)$, $O(1/(1 - \alpha - \frac{1}{d}))$ if $1/(1 - \alpha) < d \leq 2/(1 - \alpha)$, and 0 otherwise.*

PROOF. Because SBE sets $w \geq 1$ only if $d > 1/(1 - \alpha)$, $P_1 = 0$ if $d \leq 1/(1 - \alpha)$. Assume $d > 1/(1 - \alpha)$. In Step 4, channels configured in separated grids are edge-disjoint because the channels do not contain boundaries of the separated grids. Therefore, there is a unique sequence of P_1 recursive calls that set $w \geq 1$, perform inductive steps, and use a channel containing r . Two consecutive calls in the sequence are a parent and its child procedures. Moreover, all but the first call (the ancestor of any other call) in the sequence set w to $w_r \geq 1$, while the first call may set $w > w_r$ and use the channel associated with w and w_r in Step 4. The number n of nodes of the guest graph in the second call in the sequence decreases to

$$n' \leq \beta^{P_1-2} \left(n - \frac{1}{1-\beta} \right) + \frac{1}{1-\beta} \leq \beta^{P_1-2} n + \frac{1}{1-\beta}$$

at the last call in the sequence. Because the last call performs inductive steps, it follows that $n' > 2\mu d > 2/(1 - \beta)$. Thus, we have $n' < \beta^{P_1-2} n + \frac{n'}{2}$, yielding $n' < 2\beta^{P_1-2} n$. Because the second and last calls set $w = w_r \geq 1$ in Step 1, it follows that

$$\left\lfloor \frac{1}{2} \left(\left(1 - \tilde{\alpha} - \frac{1}{d} \right) \log_2 n - \log_2 \frac{\mu}{1-\beta} \right) \right\rfloor = \left\lfloor \frac{1}{2} \left(\left(1 - \tilde{\alpha} - \frac{1}{d} \right) \log_2 n' - \log_2 \frac{\mu}{1-\beta} \right) \right\rfloor.$$

Removing the floors,

$$\frac{1}{2} \left(\left(1 - \tilde{\alpha} - \frac{1}{d} \right) \log_2 n - \log_2 \frac{\mu}{1-\beta} \right) < \frac{1}{2} \left(\left(1 - \tilde{\alpha} - \frac{1}{d} \right) \log_2 n' - \log_2 \frac{\mu}{1-\beta} \right) + 1.$$

Combined with the upper bound of n' obtained above,

$$\log_2 n < \log_2 n' + \frac{2}{1 - \tilde{\alpha} - \frac{1}{d}} < \log_2 (2\beta^{P_1-2} n) + \frac{2}{1 - \tilde{\alpha} - \frac{1}{d}},$$

by which we obtain

$$P_1 < \frac{1 + \frac{2}{1 - \tilde{\alpha} - \frac{1}{d}}}{\log_2 \beta^{-1}} + 2.$$

Because $\tilde{\alpha} = \max\{1 - \frac{2}{d}, \alpha\}$, P_1 is $O(d)$ if $d > 2/(1 - \alpha)$, and $O(1/(1 - \alpha - \frac{1}{d}))$ if $1/(1 - \alpha) < d \leq 2/(1 - \alpha)$. \square

Lemma 13. *The number P_0 of recursive calls of SBE that set $w = 0$ in Step 1, perform inductive steps (i.e., not a base embedding), and use a channel containing r is $O(d)$ if $d > 2/(1 - \alpha)$, $O(1/(1 - \alpha - \frac{1}{d}))$ if $1/(1 - \alpha) < d \leq 2/(1 - \alpha)$, and at most $\log_{1/\beta} N$ otherwise.*

PROOF. By an argument similar to that in the proof of Lemma 12, there exists a unique sequence of P_0 recursive calls that set $w = 0$, perform inductive steps, and use a channel containing r . Moreover, n in the first call of the sequence decreases to n' with $2\mu d < n' < 2\beta^{P_0-1}n$ at the last call in the sequence. Therefore, it follows that $P_0 < \log_{1/\beta} n - \log_{1/\beta}(\mu d) + 1 < \log_{1/\beta} n - \log_{1/\beta} 8 + 1 < \log_{1/\beta} n$. Because $n \leq N$ obviously, we have the lemma for the case $d \leq 1/(1 - \alpha)$. If $d > 1/(1 - \alpha)$, then $w = 0$ implies that

$$\left\lfloor \frac{1}{2} \left(\left(1 - \tilde{\alpha} - \frac{1}{d} \right) \log_2 n - \log_2 \frac{\mu}{1 - \beta} \right) \right\rfloor \leq 0.$$

Removing the floor,

$$\frac{1}{2} \left(\left(1 - \tilde{\alpha} - \frac{1}{d} \right) \log_2 n - \log_2 \frac{\mu}{1 - \beta} \right) < 1,$$

by which we obtain

$$\log_{1/\beta} n = \frac{\log_2 n}{\log_2 \beta^{-1}} < \frac{2 + \log_2 \frac{\mu}{1 - \beta}}{\left(1 - \tilde{\alpha} - \frac{1}{d} \right) \log_2 \beta^{-1}}.$$

Because $\tilde{\alpha} = \max\{1 - \frac{2}{d}, \alpha\}$, P_0 is $O(d)$ if $d > 2/(1 - \alpha)$, and $O(1/(1 - \alpha - \frac{1}{d}))$ if $1/(1 - \alpha) < d \leq 2/(1 - \alpha)$. \square

Lemma 14. *The edge-congestion on r is $O(dC + d^2\Delta)$ if $d > 2/(1 - \alpha)$, $O(C/(1 - \alpha - \frac{1}{d}) + d^2\Delta)$ if $1/(1 - \alpha) < d \leq 2/(1 - \alpha)$, and $O(C(N^{\alpha-1+\frac{1}{d}} + \log N) + d^2\Delta)$ otherwise.*

PROOF. The edge r is congested by $\tilde{\sigma}_j$ with j equal to either 1 or 2 and $\tilde{\sigma}$ in each recursive call performing inductive steps and using a channel containing r , and by base embeddings. The congestion on r imposed by $\tilde{\sigma}_j$ and $\tilde{\sigma}$ in the i th recursive call in the sequence obtained by concatenating the sequences of recursive calls mentioned in Lemmas 12 and 13 is at most $\max\{6D(N_{i+1}), 4D(N_{i+1}) + 2D(N_i)\} \leq 6D(N_{i+1}) + 2D(N_i)$ by Lemmas 8 and 9, where N_i is the number of nodes of a guest graph embedded in the i th recursive call in the concatenated sequence. If r is on the boundary of a host grid in some base embedding, then r can be involved in at most $2(d - 1)$ base embeddings in total. Thus, the congestion on r is at most $\sum_{i=1}^{P_1+P_0} (6D(N_{i+1}) + 4D(N_i)) + 2(d - 1)B$.

By Lemmas 10–13, this congestion is $O((P_0 + P_1)C + d(d\Delta + C)) = O(dC + d^2\Delta)$ if $d > 2/(1 - \alpha)$, and $O(C/(1 - \alpha - \frac{1}{d}) + d^2\Delta)$ if $1/(1 - \alpha) < d \leq 2/(1 - \alpha)$. If $d \leq 1/(1 - \alpha)$, then because $N_i \leq \beta N_{i-1} + 1$, implying $N_i \leq \beta^{i-1}(N - \frac{1}{1-\beta}) + \frac{1}{1-\beta} = O(\beta^{i-1}N)$, we have

$$\begin{aligned} \sum_{i=1}^{P_1+P_0} (6D(N_{i+1}) + 4D(N_i)) + 2(d - 1)B &\leq \sum_{i=1}^{\log_{1/\beta} N} O\left(C\left(\beta^{i-1}N\right)^{\alpha-1+\frac{1}{d}}\right) + O(d(d\Delta + C)) \\ &= O\left(C(N^{\alpha-1+\frac{1}{d}} + \log N) + d^2\Delta\right). \end{aligned}$$

\square




Figure 8: $G(13)$.

By Lemma 3 and (12), the dilation of SBE is at most $\sum_{i \geq 1} O(d(\beta^{i-1}N)^{1/d}) = O(dN^{1/d})$. Therefore, we have obtained Theorem 3.

5. Lower Bound on Dilation with Minimum Edge-Congestion

In this section, we demonstrate that minimizing edge-congestion may require a dilation of nearly the size of the host grid as stated in the following theorem:

Theorem 4. *There exists an N -node graph G such that any embedding of G into an N -node 2-dimensional grid with an edge-congestion 1 must have a dilation of $\Omega(N)$.*

PROOF. For an integer $\ell \geq 9$ with $\ell \bmod 4 = 1$, we define a guest graph $G(\ell)$ obtained from $M(\ell, \ell)$ by removing edges

$$\begin{aligned} & \{(i, j), (i, j+1)\} \mid 3 \leq i \leq \ell-2, i \bmod 2 = 1, 3 \leq j \leq \ell-3\} \\ & \cup \{(i, 3), (i+1, 3)\} \mid 5 \leq i \leq \ell-3, i \bmod 4 \in \{1, 2\}\} \\ & \cup \{(i, \ell-2), (i+1, \ell-2)\} \mid 3 \leq i \leq \ell-5, i \bmod 4 \in \{3, 0\}\} \end{aligned}$$

and by adding an edge joining $e := \{(3, 3), (\ell-2, \ell-2)\}$. We illustrate $G(13)$ in Fig. 8. The graph $G(\ell)$ can be embedded into $M(\ell, \ell)$ with the edge-congestion 1 with an identity mapping for nodes and routing e on the edges removed from $M(\ell, \ell)$ to obtain $G(\ell)$. This embedding clearly has a dilation of $\Theta(\ell^2)$. We prove that if $G(\ell)$ can be embedded with the edge-congestion 1 into $M := M(\ell_1, \ell_2)$ with $\ell_1 \leq \ell \leq \ell_2$ and $\ell_1 \ell_2 = \ell^2$, then $\ell_1 = \ell_2 = \ell$ and such an embedding $\langle \phi, \rho \rangle$ is unique within rotation and/or reflection. Our proof is based on the following observation:

Observation 1. *If ρ maps k edges of $G(\ell)$ on k paths h out of which ends at a node v of M , and the other $k - h$ of which pass through v , then*

$$\deg_M(v) \geq \deg_{G(\ell)}(\phi^{-1}(v)) + 2(k - h).$$

It should be noted that because $G(\ell)$ and M have exactly the same number of nodes, there exists a node $\phi^{-1}(v)$ of $G(\ell)$ for every node v of M . We actually use this observation in different forms.

Observation 2. *If ρ maps an edge of $G(\ell)$ on a path containing a node v of M with degree 3, then this edge is incident to $\phi^{-1}(v)$.*

Observation 3. *If ρ maps two edges of $G(\ell)$ on two paths containing a node v of M with degree 4, then at least one of these edges is incident to $\phi^{-1}(v)$.*

Observations 2 and 3 are implied by Observation 1 because $G(\ell)$ has no node with degree less than 2, and therefore, for $k \geq \lfloor \deg_M(v)/2 \rfloor$,

$$h \geq \frac{\deg_{G(\ell)}(\phi^{-1}(v))}{2} + k - \frac{\deg_M(v)}{2} \geq 1 + \left\lfloor \frac{\deg_M(v)}{2} \right\rfloor - \frac{\deg_M(v)}{2} \geq \frac{1}{2},$$

implying $h \geq 1$.

We first identify nodes of $G(\ell)$ mapped onto the boundary of M . Because $G(\ell)$ has no node with degree less than 2, the node $\phi^{-1}((1, 1))$ has degree 2. Two edges of $G(\ell)$ incident to $\phi^{-1}((1, 1))$ must be routed on nodes $(1, 2)$ and $(2, 1)$ of M with degree 3. By Observation 2, therefore, these edges are incident to $\phi^{-1}((1, 2))$ and $\phi^{-1}((2, 1))$. Because only four corner nodes of $G(\ell)$, i.e., $(1, 1)$, $(1, \ell)$, $(\ell, 1)$, and (ℓ, ℓ) have degree 2 and are incident to a node with degree 3, we may assume without loss of generality that $\phi^{-1}((1, 1)) = (1, 1)$, $\phi^{-1}((1, 2)) = (1, 2)$, and $\phi^{-1}((2, 1)) = (2, 1)$. Repeating a similar argument, we can identify $\phi^{-1}((i, 1)) = (i, 1)$ for $3 \leq i \leq \ell_1$. This implies that if $\ell_1 < \ell$, then $\deg_M((\ell_1, 1)) = 2$ and $\deg_{G(\ell)}((\ell_1, 1)) = 3$, yielding an edge-congestion more than 1. Hence, we obtain $\ell_1 = \ell_2 = \ell$. As a consequence, $\phi((i, 1)) = (i, 1)$, $\phi((i, \ell)) = (i, \ell)$, $\phi((1, i)) = (1, i)$, and $\phi((\ell, i)) = (\ell, i)$ for $1 \leq i \leq \ell$.

We then identify nodes of $G(\ell)$ mapped onto nodes on one row and one column inside the boundary of M . Because $\phi((1, 2)) = (1, 2)$ and $\phi((2, 1)) = (2, 1)$, two edges of $G(\ell)$ incident to $(1, 2)$ and $(2, 1)$ must be routed on the node $(2, 2)$ of M . By Observation 3, therefore, at least one of these two edges of $G(\ell)$ is incident to $\phi^{-1}((2, 2))$. Thus, we can identify $\phi^{-1}((2, 2)) = (2, 2)$ because all the other nodes of $G(\ell)$ adjacent to $(1, 2)$ or $(2, 1)$, i.e., $(1, 1)$, $(1, 3)$, and $(3, 1)$ have already been identified to be mapped to other positions. Repeating a similar argument, we obtain $\phi((i, 2)) = (i, 2)$, $\phi((i, \ell - 1)) = (i, \ell - 1)$, $\phi((2, i)) = (2, i)$, and $\phi((\ell - 1, i)) = (\ell - 1, i)$ for $2 \leq i \leq \ell - 1$.

Because $\phi((2, 3)) = (2, 3)$ and $\phi((3, 2)) = (3, 2)$, we can identify $\phi((3, 3)) = (3, 3)$ as done for $\phi((2, 2)) = (2, 2)$, and similarly, $\phi((i, 3)) = (i, 3)$ for $i \in \{3, 4, 5, \ell - 2\}$ and $\phi((i, \ell - 2)) = (i, \ell - 2)$ for $i \in \{3, \ell - 4, \ell - 3, \ell - 2\}$.

Now we identify the routing of e . Two edges of $G(\ell)$ incident to $\phi^{-1}((3, 3)) = (3, 3)$ and $\phi^{-1}((2, 4)) = (2, 4)$ must be routed on the node $(3, 4)$ of M . By Observation 3, therefore, we can identify $\phi^{-1}((3, 4)) = (3, 4)$ and $\rho(e)$ passing through $(3, 4)$ because

all the other nodes of $G(\ell)$ adjacent to $(3, 3)$ or $(2, 4)$, including $(\ell-2, \ell-2)$, have already been identified to be mapped to other positions. With this fact, either e or an edge of $G(\ell)$ incident to $\phi^{-1}((3, 4)) = (3, 4)$, and an edge of $G(\ell)$ incident to $\phi^{-1}((4, 3)) = (4, 3)$ must be routed on the node $(4, 4)$ of M . By Observation 3 again, we can identify $\phi^{-1}((4, 4)) = (4, 4)$ because all the other nodes of $G(\ell)$ adjacent to $(3, 3)$, $(3, 4)$, or $(4, 3)$ have already been identified to be mapped to other positions. This implies that $\rho(e)$ passes through $(3, 4)$ toward $(3, 5)$. Repeating a similar argument, we obtain $\phi((3, i)) = (3, i)$, $\phi((4, i)) = (4, i)$, and $\rho(e)$ passes through $(3, i)$ toward $(3, i+1)$ for $4 \leq i \leq \ell-3$.

The path $\rho(e)$ passes through $(3, \ell-2)$ toward $(4, \ell-2)$ because it cannot go toward other directions. Then, $\rho(e)$ passes through $(4, \ell-2)$ and $(5, \ell-2)$ toward $(5, \ell-3)$ with fixing $\phi^{-1}((4, \ell-2)) = (4, \ell-2)$, $\phi^{-1}((5, \ell-2)) = (5, \ell-2)$, and $\phi^{-1}((6, \ell-2)) = (6, \ell-2)$ as similarly discussed above. We can also identify $\phi^{-1}((7, \ell-2)) = (7, \ell-2)$ since $\phi^{-1}((6, \ell-2)) = (6, \ell-2)$ and $\phi^{-1}((7, \ell-1)) = (7, \ell-1)$.

At this point we have obtained the situation for $\rho(e)$ leaving from $(5, \ell-2)$ toward $(5, \ell-3)$, together with identified nodes of $G(\ell)$ mapped onto the 4th row, $(i, \ell-2)$ for $i \in \{5, 6, 7\}$, and onto $(5, 3)$. Continuing this process until $\rho(e)$ arrives at $(\ell-2, \ell-2)$, we conclude that the embedding $\langle \phi, \rho \rangle$ is unique. \square

6. Concluding Remarks

An open question is to improve the approximation ratio for $d \leq 1/(1-\alpha)$. A main defect of SBE in approximation for $d \leq 1/(1-\alpha)$ is the use of an edge of the host grid in $\Theta(\log N)$ recursive steps, yielding a gap of $\Theta(\log N)$ factor to the optimal edge-congestion in the worst case. Another open question is to improve the dilation. In this connection, the author suspects that there is a general trade-off between edge-congestion and dilation, such as existence of guest graphs whose any embedding into a grid does not allow constant ratio approximation for both dilation and edge-congestion.

An analogous fact to Theorem 4 for hypercubes can also be proved using the existence of an induced path of length $\Theta(N)$ in an N -node hypercube [32].

References

- [1] A. L. Rosenberg, L. S. Heath, Graph Separators, with Applications, Kluwer Academic/Plenum Publishers, 2001.
- [2] Y. M. Kim, T.-H. Lai, The complexity of congestion-1 embedding in a hypercube, *J. Algorithms* 12 (1991) 246–280.
- [3] S.-K. Lee, H.-A. Choi, Link-disjoint embedding of complete binary trees in meshes, *Networks* 30 (1997) 283–292.
- [4] J. Trdlička, P. Tvrđík, Embedding complete k -ary trees into k -square 2-D meshes with optimal edge congestion, *Parallel Comput.* 26 (2000) 783–790.
- [5] M. Formann, F. Wagner, The VLSI layout problem in various embedding models, in: Proc. 16th International Workshop on Graph-Theoretic Concepts in Computer Science, volume 484 of *Lecture Notes in Computer Science*, Springer, 1991, pp. 130–139.

- [6] P. Zienicke, Embeddings of treelike graphs into 2-dimensional meshes, in: Proc. 16th International Workshop on Graph-Theoretic Concepts in Computer Science, volume 484 of *Lecture Notes in Computer Science*, Springer, 1991, pp. 182–192.
- [7] D. Gordon, Efficient embeddings of binary trees in VLSI arrays, *IEEE Trans. Comput.* C-36 (1987) 1009–1018.
- [8] A. Matsubayashi, S. Ueno, Small congestion embedding of graphs into hypercubes, *Networks* 33 (1999) 71–77.
- [9] M. Röttger, U.-P. Schroeder, Efficient embeddings of grids into grids, *Discrete Appl. Math.* 108 (2001) 143–173.
- [10] M. Röttger, U.-P. Schroeder, Embedding 2-dimensional grids into optimal hypercubes with edge-congestion 1 or 2, *Parallel Process. Lett.* 8 (1998) 231–242.
- [11] S. N. Bhatt, F. T. Leighton, A framework for solving VLSI graph layout problems, *J. Comput. Syst. Sci.* 28 (1984) 300–343.
- [12] V. Heun, E. W. Mayr, Embedding graphs with bounded treewidth into their optimal hypercubes, *J. Algorithms* 43 (2002) 17–50.
- [13] V. Heun, E. W. Mayr, Efficient embeddings into hypercube-like topologies, *The Computer Journal* 46 (2003) 632–644.
- [14] T. Leighton, S. Rao, Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms, *J. ACM* 46 (1999) 787–832.
- [15] J. D. Ullman, *Computational Aspects of VLSI*, Computer Science Press, 1984.
- [16] C. E. Leiserson, *Area Efficient VLSI Computation*, MIT Press, 1983.
- [17] L. G. Valiant, Universality considerations in VLSI circuits, *IEEE Trans. Comput.* C-30 (1981) 135–140.
- [18] G. Even, S. Guha, B. Schieber, Improved approximations of crossings in graph drawings and VLSI layout areas, *SIAM J. Comput.* 32 (2003) 231–252.
- [19] S. N. Bhatt, F. R. K. Chung, F. T. Leighton, A. L. Rosenberg, Efficient embeddings of trees in hypercubes, *SIAM J. Comput.* 21 (1992) 151–162.
- [20] B. Obrenić, An approach to emulating separable graphs, *Math. Systems Theory* 27 (1994) 41–63.
- [21] M. Baumslag, F. Annexstein, A unified framework for off-line permutation routing in parallel networks, *Math. Systems Theory* 24 (1991) 233–251.
- [22] V. E. Beneš, *Mathematical Theory of Connecting Networks and Telephone Traffic*, Academic Press, 1965.
- [23] F. T. Leighton, New lower bound techniques for VLSI, *Math. Systems Theory* 17 (1984) 47–70.

- [24] R. J. Lipton, R. E. Tarjan, A separator theorem for planar graphs, *SIAM J. Appl. Math.* 36 (1979) 177–189.
- [25] N. Alon, P. Seymour, R. Thomas, A separator theorem for nonplanar graphs, *J. Amer. Math. Soc.* 3 (1990) 801–808.
- [26] T. Kloks, *Treewidth—Computations and Approximations*, volume 842 of *Lecture Notes in Computer Science*, Springer, 1994.
- [27] M. R. Garey, D. S. Johnson, L. Stockmeyer, Some simplified NP-complete graph problems, *Theoret. Comput. Sci.* 1 (1976) 237–267.
- [28] T. N. Bui, C. Jones, Finding good approximate vertex and edge partitions is NP-hard, *Inform. Process. Lett.* 42 (1992) 153–159.
- [29] U. Feige, M. Hajiaghayi, J. R. Lee, Improved approximation algorithms for minimum weight vertex separators, *SIAM J. Comput.* 38 (2008) 629–657.
- [30] S. Arora, S. Rao, U. Vazirani, Expander flows, geometric embeddings and graph partitioning, *J. ACM* 56 (2009) article 5.
- [31] D. König, Über Graphen und ihre Anwendung auf Determinantentheorie und Mengenlehre, *Math. Annalen* 77 (1916) 453–465.
- [32] H. L. Abbott, M. Katchalski, On the snake in the box problem, *J. Combin. Theory Ser. B* 45 (1988) 13–24.