

Search for the end of a path in the d -dimensional grid and in other graphs

Dániel Gerbner* Balázs Keszegh† Dömötör Pálvölgyi‡
Günter Rote§ Gábor Wiener¶

October 11, 2015

Keywords: Separator, graph, search, grid.

AMS 2010 Mathematics Subject Classification: 90B40, 05C85.

Abstract

We consider the worst-case query complexity of some variants of certain **PPAD**-complete search problems. Suppose we are given a graph G and a vertex $s \in V(G)$. We denote the directed graph obtained from G by directing all edges in both directions by G' . D is a directed subgraph of G' which is unknown to us, except that it consists of vertex-disjoint directed paths and cycles and one of the paths originates in s . Our goal is to find an endvertex of a path by using as few queries as possible. A query specifies a vertex $v \in V(G)$, and the answer is the set of the edges of D incident to v , together with their directions.

We also show lower bounds for the special case when D consists of a single path. Our proofs use the theory of graph separators. Finally, we consider the case when the graph G is a grid graph. In this case, using the connection with separators, we give asymptotically tight bounds as a function of the size of the grid, if the dimension of the grid is considered as

*Hungarian Academy of Sciences, Alfréd Rényi Institute of Mathematics, P.O.B. 127, Budapest H-1364, Hungary. gerbner.daniel@renyi.mta.hu. Supported by the Hungarian National Science Fund (OTKA), grant PD 109537

†Hungarian Academy of Sciences, Alfréd Rényi Institute of Mathematics, P.O.B. 127, Budapest H-1364, Hungary. keszegh@renyi.hu. Supported by the Hungarian Scientific Research Fund (OTKA), grant PD 108406, NN 102029 (EUROGIGA project GraDR 10-EuroGIGA-OP-003), NK 78439, by the János Bolyai Research Scholarship of the Hungarian Academy of Sciences, and DAAD.

‡Institute of Mathematics, Eötvös University, Pázmány Péter sétány 1/C, H-1117, Budapest, Hungary. domotorp@gmail.com. Supported by the Hungarian Scientific Research Fund (OTKA), grant PD 104386 and NN 102029 (EUROGIGA project GraDR 10-EuroGIGA-OP-003), and the János Bolyai Research Scholarship of the Hungarian Academy of Sciences.

§Freie Universität Berlin, Institut für Informatik, Takustraße 9, 14195 Berlin, Germany. rote@inf.fu-berlin.de. Supported by the ESF EUROCORES programme EuroGIGA-VORONOI, Deutsche Forschungsgemeinschaft (DFG): RO 2338/5-1.

¶Department of Computer Science and Information Theory, Budapest University of Technology and Economics, Műgyetem rkp. 3., H-1111, Budapest, Hungary. wiener@cs.bme.hu. Supported by the Hungarian Scientific Research Fund (OTKA), grant 108947, and the János Bolyai Research Scholarship of the Hungarian Academy of Sciences.

42 fixed. In order to do this, we prove a separator theorem about grid graphs,
43 which is interesting on its own right.

44 1 Introduction

45 This paper deals with the following search problem. We are given a simple, undi-
46 rected, connected graph G and a vertex $s \in V(G)$. We denote the directed graph
47 obtained from G by directing all edges in both directions by G' . Let D be a directed
48 subgraph of G' , which is the vertex-disjoint union of a directed path starting at s
49 and possibly some other directed paths and cycles. D is unknown to us, and our
50 goal is to identify an endvertex of a directed path. We may *query* a vertex v , and
51 as an answer, we learn the edges of D incident to v together with their directions.
52 In particular, if the answer is only one incoming edge, then we know that v is an
53 endvertex. We analyze the minimum number of queries that are necessary in the
54 worst case.

55 We give lower bounds in the more restrictive model where we know D is one
56 directed path. Note that if instead of looking for an endvertex, we look for an
57 ending or a starting vertex of a path (different from s), then this model still gives
58 a lower bound for this easier problem. In Section 4 we mention some additional
59 models.

60 Denote by $h(G)$ the minimum number of queries needed to find an endvertex
61 in the worst case for any $s \in G$. If we know that D is one directed path, denote
62 this quantity by $h_P(G)$.

63 **Biseparators and multiseparators.** To state some of our results we need to
64 define separators of graphs. This notion can be defined in two different ways and
65 both definitions are widely used. Here we distinguish between the two definitions.

66 **Definition 1.1.** 1. Given a graph $G = (V, E)$, a subset $S \subseteq V$ is called an
67 α -biseparator of G if $V \setminus S$ can be divided into two parts, A and B , such that
68 there are no edges between A and B , and both have cardinality at most $\alpha|V|$.

69 2. Given a graph $G = (V, E)$, a subset $S \subseteq V$ is called an α -multiseparator of
70 G if every connected component of $V \setminus S$ has cardinality at most $\alpha|V|$.

71 Note that A or B in the definition of a biseparator can be empty: we do not
72 require $V \setminus S$ to be disconnected. Small biseparators make sense only for $\alpha \geq 1/2$.

73 Given these definitions, when we write *separator*, it can mean either a bisepara-
74 tor or a multiseparator, as in many cases it makes no difference. In the literature,
75 the notation $f(n)$ -separator can also be found, where $f(n)$ is an upper bound on
76 the cardinality of S in terms of the number n of vertices. In this paper it is more
77 straightforward to fix α and then look for the smallest α -separator. Therefore, we
78 let $s_\alpha^{\text{bi}}(G)$ be the minimum cardinality of an α -biseparator in G and $s_\alpha^m(G)$ be the
79 minimum cardinality of an α -multiseparator in G .

80 It follows from the definitions that every α -biseparator is an α -multiseparator,
81 and thus $s_\alpha^{\text{bi}}(G) \geq s_\alpha^m(G)$. In many cases they are of the same order of magnitude.
82 In particular, if we have a bound $s_\alpha^m(G) \leq O(n^c)$ for a class of graphs which is
83 closed under taking subgraphs for some $c < 1$ and for *some arbitrary* $\alpha < 1$, we

84 get the same asymptotic bound on $s_{1/2}^{\text{bi}}(G)$, by iteratively separating one of the
 85 components. However, there are cases when multiseparators are much smaller
 86 than biseparators. For example, if G consists of three disjoint cliques of equal size,
 87 all connected to a degree-three vertex, then $s_{1/2}^m(G) = 1$ but $s_{1/2}^{\text{bi}}(G) = \lceil n/6 \rceil$. For
 88 any tree, $s_{1/2}^m(G) = 1$ but it is not hard to show that for a complete ternary tree,
 89 $s_{1/2}^{\text{bi}}(G) = \Theta(\log n)$, see Appendix A. Finally, if we consider a class of graphs closed
 90 under taking subgraphs, by repeatedly refining the separation, then it is obvious
 91 that $s_\alpha^m(G)$ and $s_{\alpha'}^m(G)$ have the same order of magnitude for any two constants α
 92 and α' .

93 **Results.** Our main result establishes a connection between the biseparators and
 94 the search complexity for general graphs.

95 **Theorem 1.2.** *For any connected graph G with at least 2 vertices, we have*
 96 $s_{1/2}^{\text{bi}}(G) \leq h_P(G) \leq h(G)$.

97 We can prove an upper bound of the same order of magnitude, if every subgraph
 98 has small multiseparators. Note that when bounding $h(G)$, $s^{\text{bi}}(G)$, the larger of
 99 the separators, gives the lower bound and $s^m(G)$, the smaller one, gives the almost
 100 matching upper bound, which implies that indeed for a large class of graphs $s^{\text{bi}}(G)$
 101 and $s^m(G)$ have the same order of magnitude.

102 **Theorem 1.3.** *Let $0 < \alpha, \beta < 1$ be constants, let f be a monotone function, and*
 103 *let G be a graph such that any subgraph H of G has an α -multiseparator of size at*
 104 *most $f(|V(H)|)$. If $f(\alpha x) \leq \beta f(x)$ for all $x > 0$, then*

$$105 \quad h_P(G) \leq h(G) \leq \frac{f(|V(G)|)}{1 - \beta}.$$

106 The condition on f could be interpreted as having “at least polynomial growth”.
 107 The condition is fulfilled by the function $f(x) = \text{const} \cdot x^c$ if and only if $c \geq \log_\alpha \beta$.
 108 To put it differently, if α and $c > 0$ are given, the theorem applies with $\beta := \alpha^c$.

109 We also study the search problem for the special case of grid graphs.

110 **Definition 1.4.** *Let d be a positive integer and (n_1, \dots, n_d) a sequence of posi-*
 111 *tive integers. The d -dimensional grid graph of side length (n_1, \dots, n_d) , denoted by*
 112 $G_d(n_1, \dots, n_d)$, *has vertex set $\times_i \{0, 1, 2, \dots, n_i - 1\}$, and there is an edge between*
 113 *two vertices if and only if they differ in exactly one coordinate and the difference*
 114 *is 1. If $n_1 = n_2 = \dots = n_d$, then we simply write $G_d(n)$.*

115 We estimate the search complexity of grid graphs as follows.

116 **Theorem 1.5.** $\Omega(n^{d-1}/\sqrt{d}) \leq h_P(G_d(n)) \leq h(G_d(n)) \leq O(n^{d-1})$.

117 As a tool, we will prove a bound on the cardinality of separators of grid graphs,
 118 using classic results from the theory of vertex isoperimetric problems and cube
 119 slicing.

120 **Theorem 1.6.** *The smallest 1/2-biseparator of the grid graph $G_d(n)$ has cardinal-*
 121 *ity $s^{\text{bi}}(G_d(n)) = \Theta(n^{d-1}/\sqrt{d})$.*

122 We note that when considering grid graphs, one could also study the related
 123 problem that the path starting at s is monotone, i.e., if u and v are on the path and
 124 $u \leq v$ (according to the usual partial order of the vectors), then the edge between
 125 u and v (if it exists) is directed towards v . In this case the needed number of
 126 queries reduces dramatically. Indeed, the trivial algorithm which follows the path
 127 uses at most dn queries. In two dimensions we could improve slightly this upper
 128 bound, yet there is a more significant improvement by Xiaoming Sun (personal
 129 communication), who proved that $8n/5$ queries are enough in two dimensions.
 130 From below, at least $n - 2$ queries are needed regardless of d [6, Lemma 6]. This
 131 problem resembles the pyramid-path search problem (but it is not exactly the
 132 same), where also a lower bound of n is proved for the two-dimensional case [4].

133 **Motivation.** Hirsch, Papadimitriou and Vavasis [6] have proved worst-case lower
 134 bounds for finding Brouwer fixed points for algorithms using only function evalua-
 135 tion. They showed a lower bound that is exponential in the dimension, disproving
 136 the conjecture that Scarf’s algorithm is polynomial. In our language, they have
 137 (implicitly) proved that $h(G_d(n)) = \Omega(n^{d-2}/d^2)$ [6, Lemma 16]. Our Theorem 1.5
 138 is an improvement of their result, although we do not use the continuous setting
 139 but rather focus only on the discretization of the problem.

140 Later, Papadimitriou [10] considered similar complexity search problems in
 141 great detail and defined corresponding complexity classes **PPA**, **PPAD**, etc. In
 142 his model, an exponential-size graph is given by a *succinct* representation, i.e., by
 143 the description of a Turing-machine T . The vertices of the graph correspond to
 144 binary sequences of length n and if we input such a sequence to T , it outputs all
 145 the neighbors of the corresponding vertex in polynomial time (thus the degrees are
 146 bounded by a polynomial). Therefore, in his model, instead of considering query
 147 cost, one can work with the classical running time of the algorithm that gets T as
 148 input. If the algorithm uses T as a black box, we get back the query-cost model.

149 Papadimitriou considered the problem when the maximum degree of the graph
 150 is 2, i.e., it consists of vertex disjoint paths and cycles and we are also given, as part
 151 of the input, a degree-one vertex, s , and our goal is to output another degree-one
 152 vertex. This search problem is denoted by **LEAF**, and the complexity class **PPA**
 153 is defined such that **LEAF** is complete for **PPA**. (**PPA** stands for “Polynomial
 154 Parity Argument”.)

155 Papadimitriou introduced another variant, where the underlying graph is di-
 156 rected (T outputs both the in- and out-neighbors of its input in this case), the in-
 157 and out-degree of every vertex is at most one, and we are given a starting vertex
 158 s with in-degree zero and out-degree one. Therefore, the resulting digraph is the
 159 vertex-disjoint union of a directed path starting at s and possibly some other di-
 160 rected paths and cycles, exactly like in the problem that we study. Here our goal
 161 can be either to output an in-degree one, out-degree zero vertex (called **LEAFDS**
 162 problem) or an in-degree plus out-degree equals one vertex (called **LEAFD** prob-
 163 lem), which means the end of a path, just like in the problem we study. Thus, the
 164 query-cost of **LEAFD** is exactly $h(K_{2^n})$.

165 The complexity classes for which the problems **LEAFDS** and **LEAFD** are
 166 complete are denoted, respectively, by **PPADS** and **PPAD**. It is easy to see that
 167 **PPAD** is contained in both **PPA** and **PPADS**, while an oracle separation is

known for the two latter classes [1]. Nowadays **PPAD** enjoys huge popularity, as several problems, among them finding an ϵ -approximate Nash-equilibrium, turned out to be **PPAD**-complete. This is why this paper focuses on $h(G)$, the query-cost version of **PPAD**, though most of our results would also hold for the other variants.

An extensive list of **PPAD**-complete problems can be found on Wikipedia.

2 Upper bounds

Observation 2.1. *Suppose that the connected components of $G \setminus S$ are Y_1, \dots, Y_k . If every vertex of S has been queried, we know a Y_i which contains an endvertex (or that an endvertex is in S , hence already identified).*

Proof. The answers clearly show how many times we enter and leave S from each component Y_i . If we enter a component Y_i more times than we leave it, then Y_i must contain an endvertex. If there is no such component, the component containing s must contain an endvertex. \square

This simple observation is crucial for our upper bounds and it does not hold if the answers would contain only the edges leaving the queried vertex.

Proof of Theorem 1.3. Let us choose an α -multiseparator S_1 with $|S_1| \leq f(|V(G)|)$ which cuts G into parts Y_1, \dots, Y_k , and query all vertices of S_1 . By Observation 2.1 we know a part Y_j which contains an endvertex. Let G_1 be G restricted to Y_j and choose an α -multiseparator S_2 of size at most $f(|V(G_1)|)$, which cuts G_1 into parts Z_1, \dots, Z_l .

Then $S_1 \cup S_2$ is a separator of G , which cuts it into parts $Y_1, \dots, Y_{j-1}, Y_{j+1}, \dots, Y_k, Z_1, \dots, Z_l$. Thus, by again using Observation 2.1 after asking every vertex of $S_1 \cup S_2$ we know which part Z_i contains an endvertex.

After this we can continue the same way, defining G_2 and asking S_3 , defining G_3 and asking S_4 and so on, until an endvertex is in some S_i . As $|V(G_j)| \leq \alpha|V(G_{j-1})|$ for any j , one can easily see that $|V(G_j)| \leq \alpha^j|V|$. By the assumptions on f , $f(|S_j|) \leq f(|V(G_{j-1})|) \leq f(\alpha^{j-1}|V|) \leq \beta^{j-1}f(|V|)$. Altogether at most $\sum_{j=1}^{\infty} \beta^{j-1}f(|V|) \leq f(|V|)/(1 - \beta)$ questions were asked. \square

A celebrated theorem of Lipton and Tarjan [7] states that planar graphs have $2/3$ -separators of size at most $\sqrt{8} \cdot \sqrt{|V|}$. Thus we have the following corollary.

Corollary 2.2. *If G is planar, then $h(G) = O(\sqrt{|V|})$.*

Now, let us look at d -dimensional grid graphs. Miller, Teng and Vavasis [8] introduced the so-called overlap graphs for every d and proved that every member G of the class has separator of size $O(|V(G)|^{(d-1)/d})$. They mention that any subset of the d -dimensional infinite grid graph belongs to the class of overlap graphs. The polynomial function $f(x) = cx^{(d-1)/d}$ satisfies the assumption of Theorem 1.3. Since $|V(G_d(n))| = n^d$, this implies that $h(G) = O(n^{d-1})$. Here we show that the multiplicative constant is less than 3.

Theorem 2.3. $h(G_d(n)) \leq (2 + \frac{1}{2^{d-1}-1})n^{d-1}$.

208 *Proof.* We follow the proof of Theorem 1.3, but the cuts we use are always axis-
 209 aligned hyperplanes, which cut the current part into two smaller grid graphs.
 210 More precisely, for any i let $j \equiv i \pmod{d}$, $0 \leq j \leq d-1$; now S_i is a hyperplane
 211 perpendicular to the j^{th} coordinate axis, and it cuts G_{i-1} into two parts of size
 212 at most $|V(G_{i-1})|/2$. One can easily see that this is possible and $|S_{i+1}| \leq |S_i|/2$,
 213 except if $j = 0$, in which case $|S_{i+1}| \leq |S_i|$. This means that there are at most

$$214 \quad n^{d-1}(1 + 1/2 + 1/4 + \dots + 1/2^{d-1})(1 + 1/2^{d-1} + 1/2^{2(d-1)} + \dots)$$

$$215 \quad \leq n^{d-1}(2 - 1/2^{d-1}) \frac{1}{1 - 1/2^{d-1}} = n^{d-1} \left(2 + \frac{1}{2^{d-1} - 1} \right)$$

216 queries. □

217 3 Lower bounds

218 Before proving Theorem 1.6 which claims that any $1/2$ -separator in the grid graph
 219 $G_d(n)$ has cardinality $\Omega(n^{d-1}/\sqrt{d})$, we present a slightly weaker result, as it has a
 220 short proof not using results from the theory of isoperimetric problems.

221 **Claim 3.1.** *Any α -multiseparator in the grid graph $G_d(n)$ has cardinality at least*
 222 *$(1 - \alpha)n^{d-1}/d$ for $\alpha \geq 1/2$.*

223 *Proof.* We use induction on d . The claim is trivial for $d = 1$. Let us denote by S
 224 an α -multiseparator.

225 Let us choose an arbitrary axis, and denote by \mathcal{L} the n^{d-1} parallel lines in
 226 the grid which go in that direction. Let $\mathcal{L}' \subset \mathcal{L}$ be the set of those lines which
 227 intersect S . Note that every other element of \mathcal{L} contains vertices only from one
 228 component of $G \setminus S$. If $|\mathcal{L}'| \geq (1 - \alpha)n^{d-1}/d$, then we are done. Hence we can
 229 suppose $|\mathcal{L}'| < (1 - \alpha)n^{d-1}/d$.

230 Elements of \mathcal{L}' cover less than $(1 - \alpha)n^d/d$ points, hence for any component C
 231 of $G \setminus S$, the other components together contain at least $((1 - \alpha)d - (1 - \alpha))n^d/d$
 232 vertices, which are not covered by elements of \mathcal{L}' . This means that there are at
 233 least $(1 - \alpha)(d - 1)n^{d-1}/d$ elements of \mathcal{L} which contain only vertices not in C .
 234 Now consider a hyperplane in the grid, orthogonal to the direction of the lines of
 235 \mathcal{L} , and denote by \mathcal{H} the vertices of $G_d(n)$ that belong to the hyperplane. Clearly,
 236 \mathcal{H} contains at least $(1 - \alpha)(d - 1)n^{d-1}/d$ elements not in C , hence $S \cap \mathcal{H}$ is an
 237 α' -multiseparator of \mathcal{H} (with $\alpha' := 1 - (1 - \alpha)(d - 1)/d$) and so we can apply
 238 induction on each of these $(d - 1)$ -dimensional hyperplanes.

239 By induction, there are at least $(1 - \alpha)(d - 1)n^{d-2}/d(d - 1)$ elements of S
 240 in every such hyperplane, which gives at least $n(1 - \alpha)n^{d-2}/d = (1 - \alpha)n^{d-1}/d$
 241 elements in total. □

242 Before proving the stronger version of this result, we need to introduce some
 243 notations and results.

244 Let A be an arbitrary set of vertices. The set of vertices that are not in A ,
 245 but are connected to some vertex of A is called the *boundary* of A , denoted by
 246 ∂A . Following the notations of Bollobás and Leader [2], we define an order on the
 247 vertices, the simplicial order, by setting $x < y$ if $\sum x_i < \sum y_i$, or $\sum x_i = \sum y_i$

and for some j we have $x_j > y_j$ and $x_i = y_i$ for all $i < j$. This coincides with the lexicographic order according to the vector $(\sum x_i, -x_1, -x_2, \dots, -x_n)$.

Theorem 3.2 (Bollobás and Leader [2]). *In $G_d(n)$, among sets of vertices of a given size, the initial segment of the simplicial order has the smallest boundary.*

The special case $n = 2$, i.e., the hypercube, was previously treated by Harper [5], while the unbounded case of $n = \infty$ was solved by Wang and Wang [13]. We note that in the paper of Bollobás and Leader the definition of boundary is different: they also include A in ∂A .

We will also need some results about the volume of slices of a cube, i.e., intersections of the cube with specific hyperplanes. For a contemporary approach to this area we refer to [14]. In the next theorem $H^d(t)$ denotes the following set in the d -dimensional unit cube I^d : $H^d(t) = \{x \in I^d \mid \sum x_i = t\}$; Vol_i denotes the i -dimensional volume of some set of dimension i .

Theorem 3.3 ([12, 14]). $\lim_{d \rightarrow \infty} \text{Vol}_{d-1}(H^d(d/2 + s\sqrt{d})) = \sqrt{\frac{6}{\pi}} e^{-6s^2}$, for each fixed s .

Let L_k denote the k -th layer of $G_d(n)$: the set of all vertices in $G_d(n)$ whose coordinates sum to k . The layer range from 0 to $(n-1)d$. We define the size of the “middle-most” layers $Z_{n,d}$ by

$$Z_{n,d} := \begin{cases} |L_{((n-1)d-1)/2}| = |L_{((n-1)d+1)/2}|, & \text{for } (n-1)d \text{ odd,} \\ \min\{|L_{(n-1)d/2-1}|, |L_{(n-1)d/2}|, |L_{(n-1)d/2+1}|\}, & \text{for } (n-1)d \text{ even.} \end{cases}$$

$$Z_{n,d}^{\max} := \begin{cases} |L_{((n-1)d-1)/2}| = |L_{((n-1)d+1)/2}| = Z_{n,d}, & \text{for } (n-1)d \text{ odd,} \\ |L_{(n-1)d/2}|, & \text{for } (n-1)d \text{ even.} \end{cases}$$

In the even case, we actually know that the middle level $L_{(n-1)d/2}$ is the largest of the three levels in the definition of $Z_{n,d}$, as the levels decrease symmetrically in size from the middle to the ends [3]. From discretizing the above theorem, one can obtain the following bound on $Z_{n,d}$. Its proof can be found in Appendix B.

Corollary 3.4. *For every d , there exists a constant C_d such that*

$$Z_{n,d} = C_d/\sqrt{d} \cdot n^{d-1} \pm O(n^{d-2}) \text{ and}$$

$$Z_{n,d}^{\max} = C_d/\sqrt{d} \cdot n^{d-1} \pm O(n^{d-2}).$$

$C_d \rightarrow \sqrt{6/\pi}$ as $d \rightarrow \infty$.

Now we are ready to prove Theorem 1.6.

Proof of Theorem 1.6. We start with the lower bound. Let us denote by S a 1/2-biseparator which separates the vertex set A and B (such that $V = A \cup B \cup S$). If $|S| \geq Z_{n,d}$ we are done. Thus we suppose that $|S| < Z_{n,d}$. Denote by A' the vertex set of size $|A|$ which is an initial segment of the simplicial order. By Theorem 3.2 we know that $|S| \geq |\partial A| \geq |\partial A'|$.

By the definition of the simplicial order, $\partial A'$ is contained in the union of two successive layers k and $k+1$: $\partial A' = P_1 \cup P_2$, where $P_1 \subseteq L_k$ and $P_2 \subseteq L_{k+1}$. First

we claim that k must be very close to the middlemost layer. More precisely, if nd is odd, we can assume $k = \frac{nd-1}{2}$, and if nd is even, we can assume $k = \frac{nd}{2} - 1$ or $k = \frac{nd}{2}$.

We treat only the odd case, the even case being similar. First, we show that A' must reach at least level $k = \frac{nd-1}{2}$. If A' were disjoint from L_k , we would get

$$|A| + |S| = |A'| + |S| < |A'| + Z_{n,d} = |A' \cup L_k| \leq n^2/2,$$

since the last set contains only vertices in the lower half of the levels. This contradicts the requirement fact that $A \cup S$ must cover at least half of the vertices. Secondly, if A' would contain vertices of level $k+1$, it would contain more than the levels $0, 1, \dots, k$ which make up half of all vertices. This is again a contradiction to the $1/2$ -biseparator property.

By the definition of $Z_{n,d}$, we have now established that each of the two central layers L_k and L_{k+1} contains at least $Z_{n,d}$ points. To conclude the proof, we show that the separator $\partial A'$ which is contained in the two layers L_k and L_{k+1} must have size at least $Z_{n,d} - O(n^{d-2})$. If a vertex $v = (x_1, \dots, x_d)$ of L_{k+1} is not in P_2 , then the adjacent vertex v^- defined by $v^- = (x_1, \dots, x_{d-1}, x_d - 1)$ must be in P_1 unless it is not a point of the grid $G(n, d)$ (i.e., $x_d = 0$):

$$(L_{k+1} \setminus P_2)^- \cap G(n, d) \subseteq P_1$$

Since the number of vertices of L_{k+1} for which $x_d = 0$ is $O(n^{d-2})$, we obtain

$$|L_{k+1}| - |P_2| - O(n^{d-2}) \leq |P_1|,$$


from which the bound $|\partial A'| = |P_1| + |P_2| \geq Z_{n,d} - O(n^{d-2})$ follows.

For the upper bound, we simply take the central layer $L_{\lfloor (n-1)d/2 \rfloor}$ of size $Z_{n,d}^{\max}$ as a biseparator. \square


Now we are ready to prove Theorem 1.2, that $s_{1/2}^{\text{bi}}(G) \leq h_P(G)$.

Proof of Theorem 1.2. We will use an adversary argument for the lower bound on the number of queries. The adversary will try to answer the queries in such a way that the discovery of the endvertex by the searcher is delayed as much as possible. The adversary need not choose a path D in advance, but it is required that the answers remain consistent with *some* path.

Let Q denote the vertices that have been queried so far in the search. We will show that the adversary can achieve that after the other end of the path is found, Q becomes a $1/2$ -biseparator. The adversary maintains a component C of $V - Q$, see Figure 1. C is the set of vertices which can possibly be the endvertex of the path. (The adversary will follow a greedy strategy of keeping this set as large as possible.) In addition to C , the adversary maintains a path P between s and some vertex $p \in C$, which will be part of the final path and for which $P \cap C = \{p\}$. The remaining components of $V - Q$ are partitioned into two sets $V \setminus (Q \cup C) = A \cup B$ such that both A and B contain at most $|V|/2$ vertices and there are no edges between A and B . Thus we always have a partition into four disjoint sets $V = Q \cup A \cup B \cup C$. The adversary can reveal all these data to the searcher as free additional information. Initially, $C = V$, $p = s$ and $Q = A = B = \emptyset$.



326 Figure 1: A schematic drawing of the situation maintained by the adversary. The
 327 queried vertices, Q , are marked by squares.



328 Figure 2: Updating the set C after a query q

329 The strategy is the following. If the queried vertex q is in Q , the adversary
 330 repeats the previous answer for this vertex. If $q \in P \setminus \{p\}$, the adversary answers
 331 by reporting the ingoing and outgoing edge of P at that vertex. If $q \notin C \cup P$, then
 332 the answer is that “the path does not pass through this vertex.” In these cases,
 333 no new information is revealed to the searcher. The vertex p , the set C , and the
 334 path P remain unchanged; the only change is that q is moved from $A \cup B$ to Q .

335 Let us now look at the case $q \in C$. Let $C \setminus \{q\} = D_1 \cup D_2 \cup \dots \cup D_m$ be the
 336 partition of $C \setminus \{q\}$ into $m \geq 1$ connected components. The adversary chooses a
 337 largest component D_j , and will answer in such a way that the new set C becomes
 338 $C^{\text{new}} = D_j$.

339 Therefore, if C^{new} contains p , the answer is again “the path does not pass
 340 through this vertex,” see Figure 2a. The current endpoint p and the path P are
 341 unchanged. If C^{new} does not contain p (including the case $q = p$), then choose
 342 $p^{\text{new}} \in C^{\text{new}}$ to be a neighbor of q , see Figure 2b. As q was a possible endpoint of
 343 the path before this step, there is a path P^{new} from p to q which lies in $C \setminus C^{\text{new}}$.
 344 The adversary uses P^{new} and the edge qp^{new} to extend the path P to a longer path
 345 P^{new} . (This is the only case when the path is updated.) The adversary reports
 346 the last arc of P^{new} as the ingoing arc at q and qp^{new} as the outgoing arc.

347 To maintain the invariant that $|A|, |B| \leq |V|/2$, we go through the components
 348 $D_i \neq C^{\text{new}}$ and add them either to A or to B (to eventually obtain A^{new} and
 349 B^{new}), whichever is smaller. If, for example, $|A| \leq |B|$, then $|A| + |D_i| \leq |B| +$
 350 $|C^{\text{new}}| \leq |V|/2$ as $A, D_i, B, C^{\text{new}}$ are disjoint subsets of V . Therefore, the invariant

351 is maintained.

352 The searcher can only identify t , the end of the path, when $|C|$ becomes 1. By
 353 assumption, the graph G has at least two vertices and is connected, and therefore
 354 $Q \neq \emptyset$. Thus, at this point,

$$355 \quad \min\{|A|, |B|\} \leq |V \setminus (Q \cup C)|/2 \leq (|V| - 1 - 1)/2 = |V|/2 - 1.$$

356 We can now add the singleton set $C = \{t\}$ to the smaller of A and B without
 357 exceeding the size bound $|V|/2$. The set Q of queried vertices forms thus a $1/2$ -
 358 biseparator. \square

359 **Corollary 3.5.** $h_P(G_d(n)) = \Omega(n^{d-1}/\sqrt{d})$. \square

360 Theorem 1.5 summarizes the above results. The lower and upper bounds are
 361 quite close. Specifically, if we consider d as fixed, then the theorem gives exact
 362 asymptotics in n for the needed number of queries.

363 4 Concluding Remarks

364 Here we mention three more variants of the problem.

365 In the first variant, we consider any directed subgraph of G' and a vertex s
 366 with larger out-degree than in-degree. In this version there is a vertex with higher
 367 in-degree than out-degree, our goal is to find such a vertex. All of our algorithms
 368 work in this case, and obviously the same lower bounds hold.

369 In the second variant, D consists of directed paths and cycles, but we also
 370 assume that they cover every vertex. This is a special case of our model, hence
 371 the upper bounds hold. However, a lower bound similar to Theorem 1.2 is not
 372 plausible, as there are graphs that have only big separators, yet there are only a
 373 few valid choices for D . For example if G contains a vertex of degree one, different
 374 from the source, then this vertex must be the endvertex. But in case of grid graphs
 375 we can show that the additional assumption on D does not make the problem much
 376 easier.

377 Denote by $h_U(G)$ the minimum number of queries needed to find an endvertex
 378 in the worst-case for any $s \in G$. Now we show how to give a lower bound for
 379 $h_U(G_d(n))$. Let us suppose we are given an $r_1 \times r_2 \times r_3 \times \dots \times r_d$ grid graph G .
 380 Then let $G^{4,4}$ denote the $4r_1 \times 4r_2 \times r_3 \times \dots \times r_d$ grid graph.

381 **Theorem 4.1.** *Let G be a grid graph. Then $h_P(G) \leq h_U(G^{4,4})$.*

382 The proof of this theorem can be found in Appendix C.

383 One can easily see that if 4 divides n and G is the $n/4 \times n/4 \times n \times \dots \times n$
 384 grid graph, then $G_d(n) = G^{4,4}$. We need a lower bound on the size of separators
 385 in G . It is easy to see that if we replace every vertex of G by 16 vertices to get
 386 $G_d(n)$, an α -separator is replaced by an α -separator, hence the same lower bound
 387 of $\Omega(n^{d-1}/\sqrt{d})$, divided by 16, holds for G .

388 **Corollary 4.2.** $\Omega(n^{d-1}/\sqrt{d}) \leq h_U(G_d(n)) \leq O(n^{d-1})$.

389 In the third variant, D is undirected. Our goal is to find another endvertex
390 and the answer to the query is the at most two incident edges. Obviously, this is a
391 harder problem than the directed variant. Hence our lower bounds hold, and one
392 can easily modify our proofs to get the same upper bounds as well. For example,
393 in Observation 2.1, the endvertex is in the component Y_i which is connected to S
394 by an odd number of edges, counting an extra edge for the component of s .

395 Finally, a straightforward application of our proofs gives the asymptotics to a
396 question recently asked on MathOverflow [9], which is the following. Given a path
397 P_1 from the bottom-left vertex of an $n \times n$ grid to its top-right vertex, and another
398 path P_2 from its top-left vertex to its bottom-right vertex, how many queries are
399 needed to find a vertex contained in both paths? The proofs of Theorems 1.2 and
400 2.3 can be adapted to show that $\Theta(n)$ queries are necessary and sufficient.

401 Acknowledgment

402 We would like to thank our anonymous referee for the remarks that improved the
403 presentation of the paper.

404 References

- 405 [1] P. Beame, S.A. Cook, J. Edmonds, R. Impagliazzo, T. Pitassi, The Relative
406 Complexity of NP Search Problems, *J. Comput. Syst. Sci.* **57(1)**, 3–19 (1998).
- 407 [2] B. Bollobás, I. Leader, Compressions and Isoperimetric Inequalities, *Journal*
408 *of Combinatorial Theory, Series A* **56** (1991), 47–62.
- 409 [3] N. G. de Bruijn, Ca. van Ebbenhorst Tengbergen, and D. Kruyswijk, On the
410 set of divisors of a number, *Nieuw Arch. Wiskunde* **(2) 23** (1951), 191–193.
- 411 [4] D. Gerbner, B. Keszegh: Path-search in the pyramid and in other graphs,
412 *Journal of Statistical Theory and Practice*, **6(2)** (2012), 303–314.
- 413 [5] L. H. Harper, Optimal numberings and isoperimetric problems on graphs, *J.*
414 *Combin. Theory* **1** (1966), 385–393.
- 415 [6] D. Hirsch, C. H. Papadimitriou, S. A. Vavasis: Exponential Lower Bounds
416 for Finding Brouwer Fixed Points, *Journal Of Complexity* **5** (1989), 379–416.
- 417 [7] R. J. Lipton, R. E. Tarjan, A separator theorem for planar graphs, *SIAM*
418 *Journal on Applied Mathematics* **36 (2)** (1979), 177–189.
- 419 [8] G.L. Miller, S. Teng, S.A. Vavasis: A Unified Geometric Approach to Graph
420 Separators, 32nd Annual Symposium on Foundations of Computer Science
421 Proceedings (1991), 538–547.
- 422 [9] mathoverflow.net/q/185003, A problem on chains of squares – can one find
423 an easy combinatorial proof?
- 424 [10] C. Papadimitriou, On the complexity of the parity argument and other inef-
425 ficient proofs of existence. *J. Comput. System Sci.* **48** (1994), 498–532.

- 426 [11] R.P. Stanley, Enumerative Combinatorics. Vol. 1, Corrected reprint of the
 427 1986 original. Cambridge Studies in Advanced Mathematics, 49, Cambridge
 428 University Press, Cambridge (1997)
- 429 [12] G. Pólya, Berechnung eines bestimmten Integrals, Math. Ann. **74** (1913) 204–
 430 212.
- 431 [13] D.-L. Wang and P. Wang, Discrete isoperimetric problems, SIAM J. Appl.
 432 Math. **32** (1977), 860–870.
- 433 [14] Y. Xu, R. Wang, Asymptotic properties of B -splines, Eulerian numbers and
 434 cube slicing, Journal of Computational and Applied Mathematics **236(5)**
 435 (2011), 988–995.

436 A Biseparators for Ternary Trees

437 We show that a rooted ternary tree with $k + 1$ complete levels has $s_{1/2}^{\text{bi}}(G) = \Theta(k)$.
 438 Any root-to-leaf path is a 1/2-biseparator, establishing the upper bound. Let us
 439 turn to the lower bound. A complete ternary tree of height h has $n = (3^{h+1} - 1)/2$
 440 vertices. It is convenient to give each vertex a “weight” of 2. The total weight of
 441 the tree becomes $2n = 3^{k+1} - 1$, which is very near to a power of 3. In ternary
 442 notation, $2n = (22 \dots 2)_3$ with k twos, and the ideal weight for the halves of the
 443 biseparator is $2n/2 = n = (11 \dots 1)_3$.

444 After removing a separating set, any union of components of the complement
 445 can be represented as a sum and difference of subtrees. Here, by a subtree we
 446 mean a node together with all its descendents. If the separator has s nodes, we
 447 must be able to group the resulting components into a set that has between $n/2 - s$
 448 and $n/2$ nodes, i.e., weight between $n - 2s$ and n . Each separator node creates at
 449 most four new subtrees from which the sum and difference can be formed: its own
 450 subtree and the three children subtrees. (These latter ones exist only if the node
 451 was not a leaf.) So with s separating nodes, we get $1 + 4s$ subtrees from which to
 452 form the sum and difference. Each tree has a weight of the form $3^h - 1$.

453 If we take a sum and difference of $L \leq 4s + 1$ subtrees we must fulfill the
 454 inequality

$$455 \quad n - 2s \leq \sum_{i=1}^L (\pm(3^{h_i} - 1)) \leq n,$$

456 which implies

$$457 \quad n - 2s - L \leq \sum_{i=1}^L (\pm 3^{h_i}) \leq n + L$$

458 and

$$459 \quad n - 6s - 1 \leq \sum_{i=1}^L (\pm 3^{h_i}) \leq n + 4s + 1.$$

460 For any number p in the range $n - 6s - 1 \leq p \leq n + 4s + 1$, the ternary representation
 461 starts with at least $k - 1 - \lceil \log_3(6s + 1) \rceil$ ones. On the other hand, one easily sees
 462 by induction that a sum and difference of L powers of 3 has at most L ones in its

463 ternary representation. We thus get the relation $4s+1 \geq L \geq k-1 - \lceil \log_3(6s+1) \rceil$,
 464 from which $s \geq \Omega(k)$ follows. \square

465 **B Proof of Corollary 3.4**

466 We show that for any fixed $\delta \geq 0$ (and then by symmetry for every $\delta < 0$ too),
 467 whenever $(n-1)d/2 + \delta$ is an integer,

$$468 \quad |L_{(n-1)d/2+\delta}| = C_d/\sqrt{d} \cdot n^{d-1} \pm O(n^{d-2}).$$

469 We define $C_d = \text{Vol}_{d-1} H^d(d/2)$, i.e., the volume of the middle slice of the unit
 470 hypercube. Setting $s = 0$ in Theorem 3.3 establishes the convergence of C_d to
 471 $\sqrt{6/\pi}$.

472 The layer L_k , for $k = (n-1)d/2 + \delta$, is a discrete version of a slice of a cube.
 473 If we fix the first $d-1$ coordinates, then there is at most one vertex in L_k that
 474 has these first $d-1$ coordinates. Thus $|L_k| = |L'_k|$, where L'_k is the projection of
 475 L_k along the last axis.

476 To estimate the size of L'_k (and thus of L_k) take first the middle slice $H^d(d/2)$
 477 of the continuous unit cube and project it to the first $d-1$ coordinates, yielding
 478 the polytope $H^d(d/2)'$. As the normal vector of the slice is $(1, 1, \dots, 1)$, projecting
 479 it to the hyperplane orthogonal to the last axis scales the volume by a factor of
 480 $1/\sqrt{d}$:

$$481 \quad \text{Vol}_{d-1} H^d(d/2)' = \text{Vol}_{d-1} H^d(d/2)/\sqrt{d}.$$

482 Now let $H^d(d/2)'' = nH^d(d/2)'$, i.e., we blow up $H^d(d/2)'$ by a factor n . Let
 483 M be the set of grid points in this $H^d(d/2)''$. As for fixed d , $H^d(d/2)''$ is a factor- n
 484 blow up of some fixed $(d-1)$ -dimensional convex polytope, the difference between
 485 its volume and the number of grid points in it is $O(n^{d-2})$ (this follows basically
 486 from the definition of the volume, for details see e.g., [11, Proposition 4.6.13]),
 487 thus

$$488 \quad |M| = n^{d-1} \text{Vol}_{d-1} H^d(d/2)' + O(n^{d-2}) =$$

$$489 \quad = n^{d-1} \text{Vol}_{d-1} H^d(d/2)/\sqrt{d} + O(n^{d-2}) = C_d/\sqrt{d} \cdot n^{d-1} + O(n^{d-2}).$$

490 Now we are left to show that $|L'_k| = |M| + O(n^{d-2})$. For that it is enough to
 491 show that $|L'_k \setminus M|$ and $|M \setminus L'_k|$ are $O(n^{d-2})$. For all of these points the sum of
 492 the $d-1$ coordinates is equal to $(n-1)d/2 + i$ (resp. $(n-1)d/2 - n + i$) for some
 493 $0 < i \leq \delta$. This is $O(n^{d-2})$ points for every i , altogether $2\delta O(n^{d-2}) = O(n^{d-2})$
 494 points, which finishes the proof. \square

495 **C Proof of Theorem 4.1**

496 Suppose we are given a grid graph G and an Algorithm A which finds t in $G^{4,4}$
 497 in case one path and some cycles cover every vertex. We show an Algorithm B
 498 which finds the endvertex in G in case there is only a directed path. We can
 499 naturally identify every vertex of G with a 4×4 grid in $G^{4,4}$: the vertex $v =$
 500 (i_1, \dots, i_d) corresponds to the axis-parallel 4×4 rectangle (we call it a block) $B(v)$

501 having 16 vertices, whose two opposite corners are $(4i_1 - 3, 4i_2 - 3, i_3, \dots, i_d)$ and
502 $(4i_1, 4i_2, i_3, \dots, i_d)$. We call $(4i_1 - 3, 4i_2 - 3, i_3, \dots, i_d)$ and $(4i_1, 4i_2, i_3, \dots, i_d)$ the *even*
503 corners and the two other corners $(4i_1 - 3, 4i_2, i_3, \dots, i_d)$ and $(4i_1, 4i_2 - 3, i_3, \dots, i_d)$
504 the *odd* corners.

505 Consider a directed path P in G . We call a system of a directed path and some
506 directed cycles in $G^{4,4}$ *good* if they cover every vertex and the path goes through
507 exactly those blocks which correspond to the vertices of P , in the same order.

508 Now we construct good systems. If a vertex $v \in V(G)$ is not on the path, we
509 cover the corresponding block by a cycle. In case of a vertex $v = (i_1, \dots, i_d)$ on
510 the path in G , the directed path arrives at the corresponding block $B(v)$ in some
511 corner $p_1(v)$, and goes straight to a neighboring corner $p_2(v)$, where it leaves. The
512 remaining vertices form a 4×3 rectangle, which can be covered by a cycle. Finally,
513 when v is the very last vertex on the path, we define $p_1(v)$ similarly, and cover the
514 remaining vertices by a path starting in $p_1(v)$.

515 Our good systems will satisfy an additional property. If, for a vertex $v =$
516 $(i_1, \dots, i_d) \in G$, the coordinate sum $\sum_{j=3}^d i_j$ is even, then the first vertex $p_1(v)$ of
517 the path in the corresponding block is an even corner, and the last vertex $p_2(v)$
518 is an odd corner. In case $\sum_{j=3}^d i_j$ is odd, it is the other way round. Note that if
519 it is true for $B(s)$, it has to be true for every other block as well. Indeed, when
520 the path leaves a block at, for example, an odd corner, it either moves in one of
521 the first two dimensions (then it arrives to an even corner, and $\sum_{j=3}^d i_j$ does not
522 change), or in another dimension (then it arrives to an odd corner, but the parity
523 of $\sum_{j=3}^d i_j$ changes).

524 Note that these properties do not uniquely determine the system. We will
525 incrementally determine the graph as queries arrive.

526 Now we are ready to define Algorithm B. At every step we call Algorithm A,
527 and then answer such a way that at the end we get a good system. If Algorithm
528 A would query a vertex v in $G^{4,4}$, Algorithm B queries the corresponding vertex v'
529 in G instead (i.e., the vertex v' with $v \in B(v')$). Using the answer for this query,
530 we choose all the edges incident to vertices of $B(v')$ and answer to Algorithm A
531 according to this. If v' has been asked before, we have already determined the
532 edges in $B(v')$, and answer accordingly. Suppose that v' has not been queried
533 before. In case the answer is that v' is not on the path, choose an arbitrary cycle
534 covering the vertices of the corresponding block $B(v')$ and answer according to the
535 edges incident to v .

536 In case the answer gives two arcs uv' and $v'w$, we have to choose the entering
537 vertex $p_1(v')$ and the exit vertex $p_2(v')$. We will discuss this choice below. This
538 choice will define 5 edges on the path and a cycle of length 12. One edge connects
539 the blocks corresponding to u and v , leaving the last vertex of the path in $B(u)$
540 and arriving at the first vertex of the path in $B(v')$, i.e., this edge is $p_2(u)p_1(v')$.
541 Similarly we add the edge $p_2(v')p_1(w)$. We also add the three edges which connect
542 $p_1(v')$ and $p_2(v')$. Finally we cover the remaining 12 vertices with a cycle.

543 We still have to tell which one of the two possible first vertices we use as $p_1(v')$,
544 and similarly for the possible last vertices. If $p_2(u)$ has already been determined,
545 this fixes the choice of $p_1(v')$ as the vertex adjacent to it. If uv' is parallel to one of
546 the first two axes, this also reduces the choice of the corner $p_1(v')$ to one possibility.
547 Otherwise we pick $p_1(v')$ arbitrarily among the two choices. The exiting vertex

548 $p_2(v')$ is determined analogously.

549 Even if Algorithm A would know all answers in $B(v')$, it does not give more
550 information than what Algorithm B knows after asking v' . Algorithm A does not
551 finish before Algorithm B finds the end vertex, thus Algorithm A needs at least
552 as many queries as Algorithm B (on the respective graphs), which finishes the
553 proof. \square