SHALLOW GRATES

by
Maria M. Klawe

Technical Report 92-3
January 1992

Department of Computer Science
University of British Columbia
Rm 333 - 6356 Agricultural Road
Vancouver, B.C.
CANADA V6T 1Z2

SHALLOW GRATES

Maria M. Klawe
Department of Computer Science
The University of British Columbia
Vancouver, B.C., Canada V6T 1Z2

Abstract

This note proves the existence of acyclic directed graphs of logarithmic depth, such that a superlinear number of input-output pairs remain connected after the removal of any sufficiently small linearly sized subset of the vertices. The technique can be used to prove the analogous, and asymptotically optimal, result for graphs of arbitrary depth, generalizing Schnitger's grate construction for graphs of large depth. Interest in this question relates to efforts to use graph theoretic methods to prove circuit complexity lower bounds for algebraic problems such as matrix multiplication. In particular, it establishes the optimality of Valiant's depth reduction technique as a method of reducing the number of connected input-output pairs. The proof uses Schnitger's grate construction, but also involves a lemma on expanding graphs which may be of independent interest.

1. Introduction

Connectivity properties of graphs have been extensively studied in theoretical computer science because of their importance in a variety of areas. One such area involves efforts to use connectivity properties of computation graphs to prove lower bounds. One early attempt was by Valiant in [V], where he showed that computation graphs for matrix/vector products are grates. An acyclic directed graph is said to be a **grate** if a quadratic number of input-output pairs remain connected after the removal of any sufficiently small linearly sized subset of the vertices. If grates could be shown to require a superlinear number of edges, this would yield a superlinear lower bound on the size of circuits for matrix/vector products. In [S], however, Schnitger gave a construction of grates with a linear number of edges.

Paul, Pippenger, Szemeredi and Trotter [PPST] were able to prove that the computation graphs of one dimensional Turing machines do not possess a grate-like property. Specifically they showed that if n is the number of vertices in the computation graph, it is possible to remove $o(n)$ edges such that each vertex has $o(n)$ ancestors in the remaining graph. As a consequence, they proved that linear-time bounded nondeterministic Turing machines are more powerful than linear-time bounded deterministic Turing machines. Recently Razborov and Wigderson suggested the problem of trying to prove that for any acyclic directed graph of logarithmic depth, there is a small linearly sized vertex subset whose removal leaves at most a linear number of input/output pairs connected. Their hope was to obtain a lower bound for the size of logarithmic depth circuits which compute matrix/vector products by combining this with recent advances in matrix rigidity.

This note shows that Schnitger's grate construction can be used to construct acyclic directed graphs of logarithmic depth, with the property that a superlinear number of input-output pairs remain connected after the removal of any sufficiently small linearly sized subset of the vertices, thus settling the Razborov/Wigderson question in a negative fashion. The technique yields the analogous, and asymptotically optimal, result for graphs of arbitrary depth, generalizing Schnitger's result for grates. The proof requires a new result on connectivity properties of paths of expanding graphs which may have applications elsewhere. In addition, the result shows that in the worst case, using Valiant's depth reduction techniques [V] is an almost optimal method for reducing the number of connected input-output pairs. This was known, as a consequence of Schnitger's results, for graphs of large (i.e. almost linear) depth, but is new for graphs of shallower depth.

Sections 2 and 3 describe the construction for logarithmic depth. Section 3 concludes with the generalization to arbitrary depth, and a comparison of the result with Valiant's technique for reducing the number of connected input/output pairs via depth reduction. For clarity of exposition, floors and ceilings are omitted throughout.

2. The Construction

An **input (output)** in an acyclic directed graph is a vertex with in-degree (out-degree) equal to zero. The depth of an acyclic directed graph is the length of the longest directed path joining an input to an output. We say two vertices are connected if they are joined by a directed path. We wish to construct a family of graphs, $\{G_m\}$ with the following properties.

There exist positive constants α, β, Δ such that for m sufficiently large we have

2.1 G_m is an acyclic directed graph of degree Δ with m vertices

2.2 G_m has depth $\leq \beta \log m$.

2.3 If any subset of at most αm vertices is removed from G_m, at least $m(1+\alpha)^{(\log m)^{2/3}}$ input/output pairs remain connected.

By hanging an extra input and output on every vertex in G_m (see Figure 1), we see that we can replace 2.3 by the property:

2.3′ If any subset of at most αm vertices is removed from G_m, at least $m(1+\alpha)^{(\log m)^{2/3}}$ pairs of vertices remain connected.

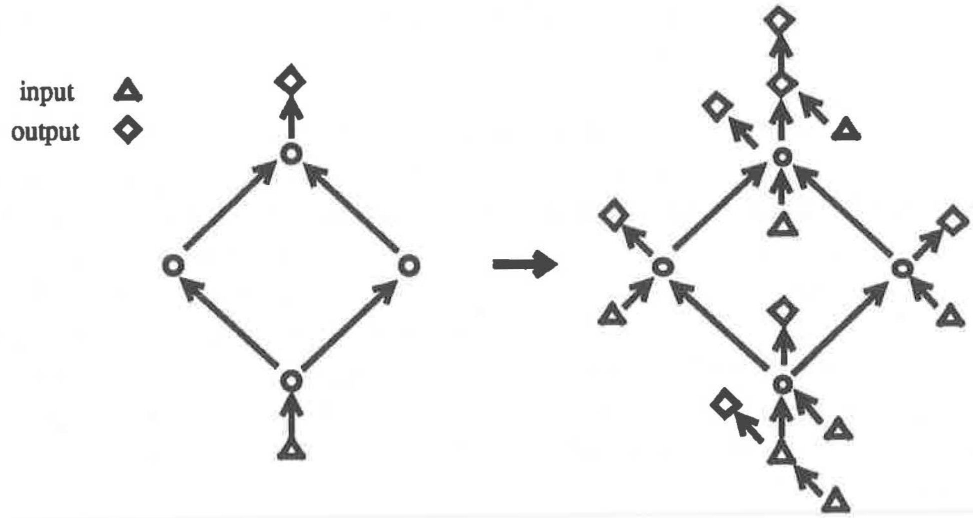

Figure 1.

We will call a family of graphs satisfying properties 2.1, 2.2 and 2.3′ a family of **weak grates**.

In [S], Schnitger constructs a family of graphs $\{H_n\}$ with the following properties.
There exists a positive constant c such that

2.4 H_n is an acyclic directed graph of bounded degree with n vertices

2.5 If any subset of at most cn vertices is removed from H_n, the remaining graph contains at least $cn^{1/3}$ vertex-disjoint paths of length $n^{2/3}$.

Before giving the details of the construction, we first sketch the underlying intuition. A natural approach to converting a Schnitger graph into a weak grate is to take $n = \log m$, and to create G_m by replacing each vertex v of H_n by a set S_v of m/n new vertices. We will call each such S_v a **supervertex.** We also replace each directed edge (u, v) in H_n with a bounded degree directed bipartite graph from S_u to S_v. We will call this bipartite graph a **superedge.** Given that we are trying to create a graph with strong connectivity properties, expanding graphs are obvious choices for superedges (see the beginning of the next section for a definition of expanding graphs).

Now consider the effect of removing a set T of αm vertices from G_m. If T happened to be contained in the union of cn supervertices, we could delete all those supervertices and the remaining graph would contain $cn^{1/3}$ disjoint 'superpaths' of supervertices of length at least $n^{2/3}$. Now, if the superedges are expanding graphs, it is easy to prove (see lemma 3.1) that there is a positive constant δ such that at least half the vertices in a superpath will be connected to at least $(1+\delta)^{n^{2/3}/2}$ other vertices. Since there are cm vertices contained in the superpaths, we would thus have at least $cm(1+\delta)^{n^{2/3}/2}/4$ connected pairs, and would thus have satisfied 2.3′ for an appropriate choice of α.

Considering the case of an arbitrary set T, we call a vertex v in H_n **bad** if $|S_v \cap T| > m/100n$. We say that a supervertex S_v is **robust** if $|S_v \cap T| \leq m/100n$. Let T' be the set of bad vertices in H_n. Obviously $|T'| \leq 100\alpha n$. Setting $\alpha < c/100$, we have at least $cn^{1/3}$ disjoint superpaths of robust supervertices of length at least $n^{2/3}$ in the graph remaining after T is removed from G_m. Unfortunately the remaining bipartite graph between two adjacent robust supervertices is no longer likely to be an expanding graph. Thus to complete the proof, we must prove that by choosing the right type of expanding graphs for the superedges, we can guarantee the following. After the removal of T, a fixed fraction of the vertices in each directed path of robust supervertices will form a directed path of expanding graphs. This result is contained in the next section, together with more precise definitions and details.

3. Paths of expanding graphs

If A is a subset of a graph, let $\Gamma(A)$ denote the set of vertices adjacent to some vertex in A. Similarly if F is a subset of the edges in the graph, $\Gamma_F(A)$ denotes the set of vertices adjacent to A in the subgraph induced by F. Expanding graphs are graphs satisfying some sort of condition which gives a lower bound on $|\Gamma(A)|$ for all subsets A which are not too large. The definition we use here is that a directed bipartite graph on equally sized vertex sets X and Y is **δ-expanding** for some positive constant δ, if for any subset $|A|$ of at most half the vertices in X we have $|\Gamma(A)| \geq (1+\delta)|A|$. Expanding graphs have been used throughout theoretical computer science for a wide variety of purposes, including proving lower bounds, asymptotically optimal constructions of graphs with connectivity properties (such as superconcentrators and grates), algorithms, simulations, and sorting networks (see [AKLLW,AKS,AM,EGS,FP,GG,L,P,Pi,S,UW] for example). The existence of expanding graphs with bounded degree can generally be proved in a fairly straightforward manner by counting methods [EGS, P]. In the most common cases explicit constructions are also known [GG, LPS]. In this paper we need a stronger version of expansion. This variant was first used by Upfal and Wigderson in [UW]. Essentially, we want a bounded degree bipartite graph that will remain an expanding graph when we restrict it to an arbitrary subset of edges, so long as each vertex still has high enough degree.

Before discussing the stronger notion of expansion, we formally define what we mean by a path of expanding graphs, and show that it yields the desired number of connected pairs of vertices. A (j, k, δ) **path of expanding graphs** is a graph $G = (V, E)$ such that V is the disjoint union of $j + 1$ sets of size k, $V_0, \ldots, V_j$, and such that E is the disjoint union of j sets, $E_1, \ldots, E_j$, where for each i the graph $(V_{i-1} \cup V_i, E_i)$ is a directed bipartite δ-expanding graph from V_{i-1} to V_i.

Lemma 3.1. If G is a (j, k, δ) path of expanding graphs where $(1+\delta)^{j/2} \leq k/2$, then the number of connected pairs is at least $jk(1+\delta)^{j/2}/2$.

Proof. For $0 < h \leq j - i$, induction on h shows that each vertex in V_i is connected to a set of $(1+\delta)^h$ vertices in V_{i+h}. The lemma follows immediately from applying this observation with $h = j/2$ to the vertices in $V_0 \cup \ldots \cup V_{j/2}$. ∎

We say that a d-regular directed bipartite graph, B_k, on k vertex subsets X and Y is a **strong expander** if it has the following property. If A is any subset of X with $|A| \leq k/2$, and if F is any set of edges of B_k such that each vertex in A is adjacent to at least $d/2$ edges in F, then $|\Gamma_F(A)| \geq (10/9)|A|$. The following lemma is easily proved by standard counting arguments; a similar result is used in [UW].

Lemma 3.2. There is a positive constant d such that for every k, there is a d-regular bipartite graph, B_k, on k vertex subsets X and Y which is a strong expander.

We now complete the definition of the directed graph G_m by specifying that each superedge is a copy of the strong expander B_k where $k = m/n$. More precisely, for each directed edge (u, v) of H_n the edges in G_m connecting S_u to S_v form a directed graph isomorphic to B_k, and if (u, v) is not an edge of H_n then there are no edges in G_m between the vertices in S_u and S_v. Now we must show that for any small enough set T of vertices to be removed from G_m, each path of robust supervertices in G_m will contain a large enough subgraph (i.e. containing a fixed fraction of the vertices in supervertices in the path) which remains a path of expanding graphs after the removal of T. We will denote the graph obtained by removing T from G_m by $G_m \backslash T$. The next lemma provides the basic property of adjacent robust supervertices that we will apply inductively in 3.4 to get the desired result.

Lemma 3.3. Let T be a set of less than $cm/100$ vertices of G_m, let S_u, S_v be robust supervertices such that (u, v) is an edge of H_n, and let $k = m/n$. Then for any subset V_v of S_v such that $V_v \cap T = \emptyset$ and $|V_v| = 9k/10$, there is a subset V_u of S_u with $V_u \cap T = \emptyset$ and $|V_u| = 9k/10$, and such that the induced subgraph of G_m on $V_u \cup V_v$ is (10/9)-expanding.

Proof. Let A be the subset of S_u consisting of all vertices which are adjacent to at least $d/2$ vertices in V_v. Then we have $|A| \geq 91k/100$. To see this consider the set $B = S_u \backslash A$ and the set F of edges joining vertices in B to vertices in $S_v \backslash V_v$. It suffices to show that $|B| \leq 9k/100$. Each vertex in B is adjacent to at least $d/2$ edges in F. Thus for any subset B' of B with $|B'| \leq k/2$ we have $|\Gamma_F(B')| \geq 10|B'|/9$ by the definition of strong expander. Combining this with the facts that $\Gamma_F(B') \subset S_v \backslash V_v$ and $|S_v \backslash V_v| = k/10$, we obtain $|B'| \leq 9k/100$. Thus B has no subsets of size greater than $9k/100$ but less than or equal to $k/2$, and hence B itself must have size at most $9k/100$. Now let V_u be any subset of A such that $V_u \cap T = \emptyset$ and $|V_u| = 9k/10$. Such a subset exists since $|A \cap T| \leq k/100$ by the robustness of S_u. Finally, it is easy to see that the induced subgraph of G_m on $V_u \cup V_v$ is (10/9)-expanding. ∎

Corollary 3.4. Let T be a set of less than $cm/100$ vertices of G_m, let $S_{v_0}, \ldots, S_{v_j}$ be robust supervertices which form a superpath in G_m, and let $k = m/n$. Then for $i = 0, \ldots, j$ there is a subset V_{v_i} of S_{v_i} with $|V_{v_i}| = 9k/10$, such that the induced graph on $V_{v_0} \cup \ldots \cup V_{v_j}$ in $G_m \backslash T$ is a path of (10/9)-expanding graphs.

Proof. The proof follows easily from Lemma 3.3 by induction on j. Specifically, for $j = 1$ we let V_{v_1} be any subset of $S_{v_1} \backslash T$ containing $9k/10$ vertices, and let V_{v_0} be the set V_u obtained by applying 3.3 with $(u, v) = (v_0, v_1)$. Now assuming $j > 1$ and that the result holds for $j - 1$, we have the existence of the sets $V_{v_1}, \ldots, V_{v_j}$ with the desired properties. As before, taking V_{v_0} to be the set V_u obtained by applying 3.3 with $(u, v) = (v_0, v_1)$ completes the proof.∎

As discussed in section 2, combining 3.1 and 3.4 with the property of H_n establishes the desired result.

Theorem 3.5. $\{G_m\}$ is a family of weak grates.

The only factor affecting the depth of G_m in the above construction is the choice of n in relation to m. By choosing $n = D(m)$ one obtains the following result.

Theorem 3.6. There exist positive constants α, β, Δ such that for any function D(m) satisfying $1 \leq D(m) \leq m$, for each m sufficiently large there is an acyclic directed graph, G_m, of degree Δ with m vertices, such that G_m has depth $\leq \beta D(m)$, and such that if any subset of at most αm vertices are removed from G_m, at least $m(1+\alpha)^{D(m)^{2/3}}$ input/output pairs remain connected.

In closing, it is interesting to compare this result with Valiant's technique for reducing the number of connected pairs by reducing the depth of the graph. Valiant proved in [V] that for any bounded degree acyclic directed graph G with m vertices and depth $D(m)$, and any positive integer k, it is possible to reduce the depth to $D(m)/2^k$ by removing $O(mk/\log D(m))$ vertices of G. (In fact Valiant proves a more general result for graphs of unbounded degree in terms of the number of edges that need to be removed to reduce the depth.) Valiant's result implies that there is a constant α such that after the removal of αm vertices the depth is at most $D(m)^{2/3}$. Thus the number of input/output pairs which remain connected is $O(m(\Delta)^{D(m)^{2/3}})$ where Δ is the maximum out-degree of vertices in G, showing that the result in Theorem 3.6 is optimal up to multiplicative constants in the exponent.

Acknowledgments

I would like to thank Nick Pippenger, Sasha Razborov, Bob Tarjan and Avi Wigderson for helpful conversations on this and other topics. The research in this paper is partially supported by an NSERC Operating Grant.

REFERENCES

[AKLLW] A. Aggarwal, M. Klawe, D. Lichtenstein, N. Linial, and A. Wigderson, "A lower bound on the area of permutation layouts", Algorithmica 6(1991), pp. 241-255.

[AKS] M. Ajtai, J. Komlos, and E. Szemeredi, "An $O(n \log n)$ size sorting network", Combinatorica 3 (1983), pp. 1-19.
[AM] N. Alon and V. Milman, "λ_1, isoperimetric inequalities for graphs, and superconcentrators", J. Comb. Th. Ser. B. 38 (1985), pp. 73-88.
[EGS] P. Erdos, R. L. Graham and E. Szemeredi, "On sparse graphs and dense long paths", Comp. Math. Appl. 1(1975), pp. 365-369.
[FP] J. Friedman and N. Pippenger, "Expanding graphs contain all small trees", Combinatorica 7 (1987), pp. 71-76.
[GG] O. Gabber and Z. Galil, "Explicit construction of linear-sized superconcentrators", JCSS 22 (1981), pp. 407-420.
[L] F. T. Leighton, New Lower Bound Techniques for VLSI, Math Systems Theory 17(1984), pp. 47-70.
[LPS] A. Lubotsky, R. Philips, and P. Sarnak, "Explicit expanders and the Ramanujan conjectures", Combinatorica 8 (1988), pp. 261-278.
[P] M. Pinsker, "On the complexity of a concentrator", Proc. 7th Int. Teletraffic Conf. Stockholm, 1973, pp. 318/1-318/4.
[Pi] N. Pippenger, "Communication networks", Handbook of Theor. Comp. Sci., ed. J. van Leeuwen, MIT Press, 1990, pp. 807-833.
[PPST] W. J. Paul, N.J. Pippenger, E. Szemeredi, and W. T. Trotter, "On determinism versus non-determinism and related problems, Proc. IEEE Found. Comp. Sci., 1983, pp. 429-438.
[S] G. Schnitger, "On depth-reduction and grates, Proc. IEEE Found. Comp. Sci., 1983, pp. 323-328.
[UW] E. Upfal and A. Wigderson, "How to share memory in a distributed system", JACM 34, 1, pp 116-127 (1987).
[V] L. G. Valiant, "Graph-theoretic arguments in low-level complexity", Proc. 6th Symp. Math. Found. Comp. Sci., Tratranska Lomnica, Czechoslovakia (78).