Expanders, sorting in rounds and superconcentrators of limited depth

Noga Alon*

Department of Mathematics

Massachusetts Institute of Technology

Cambridge, MA 02139

Abstract

Expanding graphs and superconcentrators are relevant to theoretical computer science in several ways. Here we use finite geometries to construct explicitly highly expanding graphs with essentially the smallest possible number of edges.

Our graphs enable us to improve significantly previous results on a parallel sorting problem, by describing an explicit algorithm to sort n elements in k time units using $O(n^{\alpha_k})$ processors, where, e.g., $\alpha_2 = 7/4$.

Using our graphs we can also construct efficient n-superconcentrators of limited depth. For example, we construct an n superconcentrator of depth 3 with $O(n^{4/3})$ edges; better than the previous known results.

1. Introduction

A graph G is called (n, α, β)-expanding, where $0 < \alpha \leq \beta \leq n$, if it is a bipartite graph on the sets of vertices I (inputs) and O (outputs), where $|I| = |O| = n$, and every set of at least α inputs is joined by edges to at least β different outputs.

Expanding graphs with a small number of edges, which are the subject of an extensive literature, are relevant to theoretical computer science in several ways. Here we merely point out two examples. A family of linear expanders of density k and expansion d is a set $\{G_n\}_{n=1}^{\infty}$ of graphs, where G_n has $\leq (k+o(1))n$ edges and is $(n, \alpha, \alpha(1 + d(1 - \alpha/n)))$-expanding for all $\alpha \leq n/2$, where $d > 0$ and k are fixed. Such a family is the basic building block used in the constructions of graphs with special connectivity properties and small number of edges (see, e.g., Chung [12]). An example of a graph of this type is an n-superconcentrator, which is a directed acyclic graph with n inputs and n outputs such that for every $1 \leq r \leq n$ and every two sets A of r inputs and B of r outputs there are r vertex disjoing paths from the vertices of A to the vertices of B. Superconcentrators have been used in the construction of graphs that are hard to pebble (see Lengauer and Tarjan [22], Pippenger [27] and Paul, Tarjan and Celoni [29]), in the study of lower bounds (see Valiant [34]), and in the establishment of time space tradeoffs for computing various functions (Abelson [1], Ja'Ja' [20] and Tompa [32]).

A family of linear expanders is also essential in the recent parallel sorting network of Ajtai, Komlós and Szemerédi [2].

It is not too difficult to prove the existence of a family of linear expanders using probabilistic arguments (see, e.g., Chung [12], Pinkser [25] and Pippenger [26]). However, for applications an explicit construction is desirable. Such a construction is far more difficult and was first given in Margulis [23] and modified in Gabber and Galil [14]. (See also Alon and Milman [4], [5] for a more general construction.)

The expanding graphs used in [14] to construct superconcentrators and those used in the sorting network of [2] are (n, α, β)-expanding for some fixed (independent of n) ratio of β/α, i.e., they are rather weakly expanding. For some applications, however, a higher amount of expansion is necessry and $(n, \alpha(n), \beta(n))$-expanding graphs are needed, where $\beta(n)/\alpha(n) \mapsto \infty$ as $n \mapsto \infty$. A possible (and essentially the only known) method to obtain (explicitly) highly expanding graphs with a small number of edges is an "iteration" of the known expander of [14] (see Pippenger [28]). Unfortunately, this method is a poor substitute for the probabilistic construction since it supplies graphs with too many edges. This makes some of the applications impossible.

*Research supported in party by the Weizmann Fellowship for Scientific Research

Here we use finite geometries to explicitly construct highly expanding graphs with essentially the smallest possible number of edges. Specifically, we show that the points versus hyperplanes incidence graph of a finite geometry of dimension d is an $(n, x, n - n^{1+1/d}/x)$-expanding graph, for all $0 < x < n$. As pointed out by Pippenger [28], the results of Guy and Znam [15] imply that any such graph must have at least $\Omega(n^{2-1/d})$ edges. Our graphs have $(1+o(1))n^{2-1/d}$ edges; only a constant times the theoretical lower bound. The previous methods were not sufficient to construct graphs with this amount of expansion having $o(n^2)$ edges.

By a theorem of Singer [16], the edges of our graphs can be defined by a set of $\approx n^{1-1/d}$ translations modulo n, in contrast to the result of Klawe [21] that asserts that no family of linear expanders can have this form. this reveals a difference between weakly expanding and highly expanding graphs.

Our new expanding graphs enable us to obtain an explicit algorithm for sorting n elements in two time units using $O(n^{7/4})$ parallel processors. this improves results of Bollobás and Rosenfeld [10], Häggvist and Hell [18] and Pippenger [28] who gave explicit algorithms to this problem using $2/5n^2 + O(n^{3/2})$, $13/30(n^2 - n)$ and $O(n^{1.943\cdots}(\log n)^{0.943\cdots})$ processors, respectively. It also enables us to improve the best known algorithms for sorting n elements in k time units, for all (fixed) $k \geq 4$. Very recently Pippenger has found a slightly better way of using our expanding graphs to get an explicit algorithm for sorting n elements in two time units using only $O(n^{26/15})$ parallel processors.

Using our graphs we also construct explicitly n-superconcentrators of depth 3 with $O(n^{4/3})$ edges - better than those having $O(n^{3/2})$ edges obtained from the results of Slepian and Duguid (cf [7, pp. 86-88]) and Meshulam [24]. This also enables one to construct better explicit superconcentrators of depth $2r+1$ for all fixed $r \geq 1$.

Our paper is organized as follows: in Section 2 we construct our geometric expanders. In Section 3 we describe how they can be applied to the problem of sorting in rounds and in Section 4 we discuss briefly superconcentrators of limited depth.

Our expanders also enable us to construct several graphs relevant to Ramsey Theory and obtain a strengthened version of the well known deBruijn-Erdös Theorem [8]. This will appear in another paper [3].

2. The geometric expanders

Let $d, q \geq 2$ be integers. Let I and O be, respectively, the sets of points and hyperplanes of a finite geometry of dimension d and order q. As is well known, such a geometry always exists if q is a prime power [16]. Let $G = G(q,d)$ denote the bipartite graph with classes of vertices I and O in which $p \in I$ is joined to $h \in O$ iff p is incident with h. When q is a prime power G has the following easy explicit construction. Let V be the set of all non-zero vectors of length $d+1$ over the finite field $GF(q)$. Two vectors $\bar{x} = (x_1, \ldots, x_{d+1})$ and $\bar{y} = (y_1, \ldots, y_{d+1})$ are equivalent iff $x_i = cy_i$ for some $c \in GF(q)$ and $1 \leq i \leq d+1$. Let $\tilde{V}$ denote the set of all equivalence classes of V under this relation. Put $I = O = \tilde{V}$. Two vertices $[(x_1, x_2, \ldots, x_{d+1})] \in I$ and $[(y_1, y_2, \ldots, y_{d+1})] \in O$ are joined iff $\Sigma_{i=1}^{d+1} x_i y_i = 0$. The next theorem shows that $G(q,d)$ is a highly expanding graph.

Theorem 2.1

Put $n = (q^{d+1}-1)/(q-1)$, $k = (q^d - 1)/(q-1)$, $\lambda = (q^{d-1} - 1)/(q-1)$.

1. $G = G(q,d)$ is k-regular and $|I| = |O| = n$; thus G has $(1+o(1))n^{2-1/d}$ edges. (As $q \mapsto \infty$ for fixed d.) Every two distinct vertices of O have precisely λ common neighbors in I.

2. If $X \subseteq I$, $|X| = x$ then $|N(X)| \geq n - n^{1+1/d}/x$. Thus G is $(n, x, n - n^{1+1/d}/x)$-expanding for all $0 < x < n$.

3. If $Z \subseteq O$ then

$$|\{i \in I : |N(i) \bigcap Z| \leq \frac{1}{2}|Z|/n^{1/d}\}| \leq 4n^{1+1/d}/|Z|.$$

Proof

Part 1 of the above theorem is well known (and easy - see, e.g., [16]). Parts 2 and 3 (and in fact even slightly stronger assertions) can be proved using a relation, similar to the one proved in [31] or in [4], between the eigenvalues of the adjacency matrix of a graph and its expansion properties. Here we present an easier proof that uses linear algebra and a certain "second moment" method. Put $G = G(q,d) = (I, O; E)$. For $i \in I$, $o \in O$ put $c_{io} = q-1$ if $io \in E$ and $c_{io} = -1$ if $io \notin E$. Suppose $X \subseteq I, Z \subseteq O$ then

$$\Sigma_{i \in X}(\Sigma_{o \in Z} c_{io})^2 \leq \Sigma_{i \in I}(\Sigma_{o \in Z} c_{io})^2 = \Sigma_{o \in Z}\Sigma_{i \in I} c_{io}^2 +$$

$$+2\Sigma_{\{o,o'\} \subset Z}\Sigma_{i \in I} c_{io} \cdot c_{io'} = |Z|(k \cdot (q-1)^2 + n - k) +$$

$$+|Z|(|Z|-1)(\lambda(q-1)^2 - 2(k-\lambda)(q-1) + n - 2k + \lambda) =$$

$$= |Z|(q^{d+1} - q + 1) - |Z|(|Z| - 1)(q - 1) \leq |Z| \cdot n^{1+1/d}.$$

To prove 2, apply the last inequality to X and $Z = O - N(X)$. (Note that for $i \in X$ $\Sigma_{o \in Z} c_{io} = -|Z|$). To prove 3, apply the last inequality to $X = \{i \in I : |N(i) \bigcap Z| < \frac{1}{2}|Z|/n^{1/d}\}$ and Z. (Note that here for $i \in X$ $\Sigma_{o \in Z} c_{io} \leq -|Z|/2$). ∎

Remarks

1. The known results about the distribution of primes (see, e.g., [6, p. xx]) clearly imply that for every fixed $d \geq 2$ and every integer n there exists a prime p such that $n \leq (p^{d+1} - 1)/(p - 1) \leq n + O(n^{1-1/(3d)})$. Any x inputs in any induced subgraph of $G(p,d)$ with n inputs and n outputs have $\geq n - (1 + o(1))\frac{n^{1+1/d}}{x}$ neighbors. Thus we have for every $d \geq 2$, an explicit construction of a family of graphs $\{H(n,d)\}_{n=1}^{\infty}$ where $H(n,d)$ has $(1 + o(1))n^{2-1/d}$ edges and is $(n, x, n - (1 + o(1))n^{1+1/d}/x)$-expanding for all $0 < x < n$.

2. by using the results of [15] on the problem of Zarenkiewicz one can easily show that any graph that has the expansion properties of $G(q,d)$ must have at least $(1 + o(1)) \cdot \ell n2 \cdot n^{2-1/d}$ edges. Note that the number of edges of $G(q,d)$ (or of $H(n,d)$) is $(1 + o(1)) \cdot n^{2-1/d}$ and thus these graphs have (up to a constant of $1/\ell n2$) the smallest possible number of edges.

3. Let $PG(d,q)$ be the finite geometry of dimension d over the field $GF(q)$ and let $G(q,d)$ be the corresponding expander. Let n, k be as in Theorem 2.1. By Singer's Theorem ([16]) there exist $0 \leq a_1 < a_2 < ... < a_k < n$ such that $G(q,d)$ is isomorphic to the bipartite graph with classes of vertices $A = B = \{0, 1, 2, ..., n-1\}$ in which $a \in A$ is joined to $b \in B$ iff $b = (a + a_i) (\mathrm{mod}\ n)$ for some $1 \leq i \leq k$. This contrasts with the result of [21] that implies that no family of linear expanders can have this form and thus shows a difference between highly expanding and weakly expanding graphs.

3. *Sorting in rounds.*

Suppose we are given n elements with a linear order unknown to us. In the first round we ask m_1 simultaneous questions, each a binary comparison. Having the answers we deduce all implications and ask, in the next round, another m_2 questions, deduce their implications, and so on. A choice of our questions that guarantees that after r rounds we will know the complete order of the elements is an algorithm for sorting in r rounds. The need for such algorithms with fixed r arises in structural modeling (see Häggvist and Hell [19]). Since all comparisons within a round are evaluated simultaneously, such algoorithms have obvious connection to parallel sorting, as defined by Valiant [33], and seem to be practical in situations like testing consumer preferences (see Scheele [30]), where the communication between our sorting computer and the consumers is being performed by correspondence. Many results about sorting in rounds can be found in the survey article [9].

Let $f_r(n)$ denote the minimum possible number of comparisons sufficient to sort n elements in r rounds. Clearly, $f_1(n) = \binom{n}{2}$. In Häggvist and Hell [17,18] and Bollobás and Thomason [11], probabilistic arguments are used to obtain estimates of $f_r(n)$ for $r \geq 2$. In particular it is known that $f_2(n) = O(n^{3/2} \log n)$ and $f_2(n) = \Omega(n^{3/2})$ (see [11]). For practical applications, however, a probabilistic argument is not enough and an explicit sorting algorithm is desirable. Häggvist and Hell observed this fact and in [19] they gave explicit algorithms for sorting in k rounds with $O(n^{s_k})$ comparisons, where $s_k \mapsto 1$ as $k \mapsto \infty$ and, e.g., $s_3 = 8/5$, $s_4 = 20/13$, and $s_5 = 28/19$. It seems more difficult to find an efficient sorting algorithm in two rounds. In [18] such an algorithm with $13/30(n^2 - n)$ comparisons is given. A somewhat better algorithm is given in Bollobás and Rosenfeld [10] - with $2/5n^2 + O(n^{3/2})$ comparisons. The only construction with $o(n^2)$ comparisons is due to Pippenger [28] - $O(n^{1.943\cdots}(\log n)^{0.943\cdots})$.

In some situations it may be undesirable to allow deducing all implications, since conclusions derived from relations themselves derived by transitivity may be unreliable. Thus one may be willing to allow only direct implications (i.e., if we find in the first round that $x < y$, $y < z$ and $z < t$ we conclude that $x < z$ and $y < t$ but not necessarily that $x < t$). In [11] a lower bound of $\Omega(n^{5/3})$ is proved for such an algorithm in 2 rounds. Using our geometric expanders we obtain:

Theorem 3.1

By an explicit construction that uses only direct implications

$$f_2(n) = O(n^{7/4}).$$

Note that by the lower bound mentioned above our construction is not fare from being best possible.

Proof.

Let A be the set of n objects we have to sort. Clearly we may assume that n is of the form $(q^5 - 1)/(q - 1)$ for some prime power q (otherwise, add $o(n)$ dummy objects to obtain an n of this form). Let $G = G(q, 4)$ be a geometric expander corresponding to a finite geometry of dimension 4 and order q. Let $I = \{v_1, v_2, ..., v_n\}$ and $O = \{u_1, u_2, ..., u_n\}$ be the sets of inputs and outputs of G, respectively. In the first round we compare the i-th element of A to the j-th element if $v_i u_j$ is an edge of G. There are $O(n^{7/4})$ such comparisons.

We proceed to show that even by deducing only direct implications, we will have to compare in the second round only $O(n^{7/4})$ pairs. For $X \subseteq A$ put $N(X) = \{y \in A : y$ is compared in the first round to some $x \in X\}$. The following two facts follow directly from Theorem 2.1.

Fact 1.

If $Z \subseteq A$, $|Z| = (4 + O(1))n^{3/4}$ and

$$X = \{x \in A : |N(x) \bigcap Z| \leq n^{1/2}\}$$

then $|X| \leq (1 + o(1))n^{1/2}$.

Fact 2.

If $Y \subseteq A$, $|Y| > n^{1/2}$ then $|N(Y)| \geq n - n^{3/4}$.

Define a partition of A into $\ell = \lceil n^{1/4}/4 \rceil$ blocks of $A_1, \ldots, A_\ell$, each of size $(4 + o(1))n^{3/4}$, such that each A_i consists of consecutive objects (in the linear order we have to find) and the maximal element of A_i is smaller than the minimal element of A_{i+1}. Call an element $a \in A_{i+1}$ *bad* if $|N(a) \bigcap A_i| \leq n^{1/2}$, otherwise call it *good*. By Fact 1 the number of bad elements in A_{i+1} is $\leq (1 + o(1))n^{1/2}$. Let $a \in A_{i+1}$ be good and suppose $b \in \bigcup_{j=1}^{i-1} A_j$. If

$$N(b) \bigcap N(a) \bigcap A_i \neq \phi \tag{3.1}$$

then, by direct implication from the first round, $b < a$. However, $|N(a) \bigcap A_i| > n^{1/2}$, and thus, by Fact 2 the number of b-s that violate (3.1) is $\leq n^{3/4}$. It follows that the total number of comparisons of an element $a \in A_{i+1}$ to elements in $\bigcup_{j=1}^{i-1} A_j$ left for the second round is bounded by n (of course) if a is bad and by $|A_i| + |A_{i+1}| + n^{3/4} = (9 + o(1))n^{3/4}$ if a is good. The total number of these comparisons is thus bounded by

$$\ell \cdot (1 + o(1))n^{1/2} \cdot n + n(9 + o(1))n^{3/4} = O(n^{7/4}).$$

Since the first round also requires $O(n^{7/4})$ comparisons, the total number of comparisons is $O(n^{7/4})$. ∎

Very recently Pippenger has shown that by using indirect implications of arbitrary length, the number of comparisons can be reduced to $O(n^{26/15})$. The first round of his algorithm uses our expanders arising from finite geometries of dimension 3.

Our new results, together with the recursive construction of Häggvist and Hell [19, theorem 3], enable us to improve the best known explicit algorithm for sorting in k rounds for all (fixed) $k \geq 4$.

4. *Superconcentrators of limited depth.*

Recall the definition of an n-superconcentrator (= s.c.) given in Section 1. The *depth* of an s.c. is the number of edges in the longest directed path from an input to an output. The *size* of an s.c. is the number of its edges. It is well known that s.c.'s of linear size exist (see [26], [34]), and in [14] an explicit construction of an n-s.c. with size $\approx 271.8n$ is given. This was improved in [5] to $\approx 158n$. However, the minimal possible size of an n-s.c. of depth r is not linear with n, for all fixed $r \geq 1$. This was shown by Dolev, Dwork, Pippenger and Wigderson [13] and, independently, by Ajtai. Meshulam [24] constructed explicitly an n-s.c. of depth 2 and size $O(n^{3/2})$. The results of Slepian, Duguid and LeCorre (cf [7, pp. 86-88]) supply an explicit n-s.c. of depth 3 and size $O(n^{3/2})$. (This is also obtained, of course, by [24].) Our geometric expanders enable us to prove;

Theorem 4.1.

By an explicit construction there is an n-s.c. of depth 3 and size $O(n^{4/3})$.

The method described in [7, pp. 136-144] enables one to use Theorem 4.1 for explicit constructions of n-s.c.'s of depth $2r + 1$ and size $O(n^{(r+3)/(r+2)})$ for all fixed $r \geq 1$; better than the previous known results. We omit the detailed constructions.

References

[1] H. Abelson, A note on time space tradeoffs for computing continuous functions, Infor. Proc. Letters 8 (1979), 215-217.

[2] M. Ajtai, J. Komlós and E. Szemerédi, Sorting in $C \log n$ parallel steps, Combinatorica 3 (1983), 1-19.

[3] N. Alon, Eigenvalues, geometric expanders, sorting in rounds and Ramsey theory, preprint.

[4] N. Alon and V.D. Milman, λ_1, isoperimetric inequalities for graphs and superconcentrators, J. Combinatorial theory, Ser. B, to appear.

[5] N. Alon and V.D. Milman, eigenvalues, expanders and superconcentrators, Proc. 25th Annual Symp.on Foundations of Comp. Sci., Florida (1984), 320-322.

[6] B. Bollobás, "Extremal Graph Theory," Academic Press, London and New York (1978).

[7] V.E. Beneš, "Mathematical Theory of Connecting Networks and Telephone Traffic," Academic Press, London and New York (1965).

[8] N.G. deBruijn and P. Erdös, On a combinatorial problem, Indagationes Math. 20 (1948), 421-423.

[9] B. Bollobás and P. Hell, Sorting and Graphs, preprint.

[10] B. Bollobás and M. Rosenfeld, sorting in one round, Israel J. Math 38 (1981), 154-160.

[11] B. Bollobás and A. Thomason, Parellel sorting, Discrete Appl. Math 6 (1983), 1-11.

[12] F.R.K. Chung, On Concentrators, superconcentrators, generalizers, and nonblocking networks, Bell Sys. Tech. J. 58 (1978), 1765-1777.

[13] D. Dolev, C. Dwork, N. Pippenger and A. Wigderson, Superconcentrators, Generalizers and generalized connectors with limited depth, preprint.

[14] O. Gabber and Z. Galil, Explicit construction of linear sized superconcentrators, J. Comp. and Sys. Sci. 22 (1981), 407-420.

[15] R.K. Guy and S. Znam, A problem of Zarenkiewicz, in "Recent Progress in combinatorics" (W.T. Tutte, ed.), Academic Press (1969), 237-243.

[16] M. Hall, Jr., "Combinatorial Theory," Wiley and Sons, New York and London (1967), 128-131.

[17] R. Häggvist and P. Hell, Graphs and parallel comparison algorithms, Congr. Num. 29 (1980), 497-509.

[18] R. Häggvist and P. Hell, Parallel sorting with constant time for comparisons, SIAM J. Comp. 10 (1981), 465-472.

[19] R. Häggvist and P. Hell, Sorting and merging in rounds, SIAM J. Alg. and Disc. Meth. 3 (1982), 465-473.

[20] J. Ja'Ja', time space tradeoffs for some algebraic problems, Proc. 12th Ann. ACM Symp. on Theory of Computing, 1980, 339-350.

[21] M. Klawe, Non-existence of one-dimensional expanding graphs, Proc. 22nd Ann. Symp. Found. Comp. Sci. Nashville (1981), 109-113.

[22] T. Lengauer and R.E. Tarjan, Asymptotically tight bounds on time space tradeoffs in a pebble game, J. ACM 29 (1982), 1087-1130.

[23] G.A. Margulis, Explicit constructions of concentrators, Prob. Per. Infor. 9 (4), (1973), 71-80 (english translation in Problems of Infor. Trans. (1975), 325-332.).

[24] R. Meshulam, A geometric construction of a superconcentrator of depth 2, preprint.

[25] M. Pinkser, On the complexity of a concentrator, 7th International Teletraffic Conference, Stockholm, June 1973, 318/1-318/4.

[26] N. Pippenger, Superconcentrators, SIAM J. Computing 6 (1977), 298-304.

[27] N. Pippenger, Advances in pebbling, Internat. Colloq. on Autom. Lang. and Prog. 9 (1982), 407-417.

[28] N. Pippenger, Explicit construction of highly expanding graphs, preprint.

[29] W.J. Paul, R.E. Tarjan and J.R. Celoni, Space bounds for a game on graphs, Math. Sys. Theory 20 (1977), 239-251.

[30] S. Scheele, Final report to office of environmental education, Dept. of Health, Education and Welfare, Social Engineering Technology, Los Angeles, CA 1977.

[31] R.M. Tanner, Explicit construction of concentrators from generalized N-gons, SIAM J. Alg. discr. Meth. 5 (1984), 287-293.

[32] M. Tompa, time space tradeoffs for computing functions, using connectivity properties of their circuits, J. Comp. and Sys. Sci. 20 (1980), 118-132.

[33] L.G. Valiant, Parallelism in comparison networks. SIAM J. Comp. 4 (1975), 348-355.

[34] L.G. Valiant, Graph theoretic properties in computational complexity, J. Comp. and Sys. Sci. 13 (1976), 278-285.