Correspondence

Small Diameter Symmetric Networks from Linear Groups

Lowell Campbell, Gunnar E. Carlsson, Michael J. Dinneen, Vance Faber, Michael R. Fellows, Michael A. Langston, James W. Moore, Andrew P. Mullhaupt, and Harlan B. Sexton

Abstract— **In this note is reported a collection of constructions of symmetric networks that provide the largest known values for the number of nodes that can be placed in a network of a given degree and diameter. Some of the constructions are in the range of current potential engineering significance. The constructions are Cayley graphs of linear groups obtained by experimental computation.**

Index Terms—**Cayley graphs, interconnection networks, linear groups.**

I. Introduction

The problem of constructing large graphs of a given degree and diameter has received much attention, and is significant for parallel processing because it models two important constraints in the design of massively parallel processing systems: 1) there are limits on the number of processors to which any processor in the network can be directly connected, and 2) the distance between any two processors in the network should not be too great. Other applications of such networks include shared-key cryptographic protocols and the design of local area networks. See [3] and [9] for recent surveys.

In this paper we give evidence that the table of largest known constructions for small values of the two parameters can be improved for many parameter values by methods based on finite linear groups. In many cases the networks we describe here are dramatically larger than those previously known.

Many of our improvements are in the range of the numbers of processors currently being considered for large parallel processing systems, suggesting that some of these constructions may merit further investigation for such applications. This is the focus of continuing research by some of our party. In this note we present only our accumulated results on the now classic problem of network construction. In particular, we do not address the many interesting problems concerning routing and data exchange that would be crucial for most parallel processing applications.

Manuscript received January 12, 1989; revised December 15, 1990. This work was supported in part by the National Science Foundation under Grant MIP-8693879, by the Office of Naval Research under Contracts N00014-88-K-0343 and N00014-88-K-0546, and by the National Aeronautics and Space Administration under Grant NAGW-1406. A preliminary version of this paper was presented at the Second Symposium on the Frontiers of Massively Parallel Computation, Fairfax, VA. Inquiries should be directed to M. Fellows.

L. Campbell is with the Department of Electrical Engineering, University of Idaho, Moscow, ID 83843.

G. E. Carlsson is with the Department of Mathematics, Princeton University, Princeton, NJ 08540.

M. J. Dinneen and M. R. Fellows are with the Department of Computer Science, University of Victoria, Victoria, BC, Canada V8W 3P6.

V. Faber and J. W. Moore are with Los Alamos National Laboratories, Los Alamos, NM 87545.

M. A. Langston is with the Department of Computer Science, University of Tennessee, Knoxville, TN 37996.

A. P. Mullhaupt is with the Department of Mathematics, University of New Mexico, Albuquerque, NM 87131.

H. B. Sexton is with Lucid, Inc., Menlo Park, CA 94025.

IEEE Log Number 9101691.

For an overview of our results see Table I, which represents an updated version, obtained from Bermond [5] of the table published in [3]. Interested readers are advised that a "current" table incorporating the results of many workers on this problem is maintained by and available from that helpful source. The entries in the table that are due to our efforts and reported on in this note are marked in bold. Other entries that have been obtained by Cayley graph techniques are marked with an asterisk. In particular, two other groups of researchers have recently and independently obtained record-breaking constructions based on linear groups [4], [8].

II. Algebraic Symmetry as an Organizing Principle for Parallel Processing

There are important considerations apart from degree and diameter that must figure in any choice of network topology for parallel computation. A network is (*vertex-*) *symmetric* if for any two nodes u, v there is an automorphism of the network mapping u to v. Our approach yields symmetric constructions, and we believe that in this may lie their greater value. Symmetry is one of the most powerful and natural tools to apply to the central problem of massively parallel computation: how to organize and coordinate computational resources.

The symmetries of the networks we describe are represented by simple algebraic operations (such as 2×2 matrix multiplications and modulo arithmetic). The main advantage of algebraically constructed networks is that the developed mathematical resources of algebra are available to structure the problems of

1) design and description
2) testing
3) data exchange and routing
4) scheduling and computation mapping.

The appeal of hypercubes, cube-connected cycles, butterfly networks, and others rests in large part on the availability of easily computed (and comprehended) symmetries. These popular network designs and those that we describe all belong to a class of algebraic networks based on vector spaces and their symmetry groups. For recent algebraic approaches to routing algorithms, deadlock avoidance, emulation, and scheduling for algebraically described networks of this kind see [2], [1], [8], [11], and [12].

Our main result in this brief paper is a demonstration that algebraic symmetry provides a powerful approach to problem 1), design and description. Our approach centers on the following definition.

Definition: If A is a group and $S \subseteq A$ is a generating set that is closed under inverses, i.e., $S = S \cup S^{-1}$, then the (undirected) *Cayley graph* (A, S) is the graph with vertex set A and with an edge between elements a and b of A if and only if $as = b$ for some $s \in S$.

Every Cayley network is symmetric (symmetries are given by group multiplication). The degree of a Cayley graph (A, S) is $\Delta = |S|$ and the diameter of (A, S) is

$$D = \max_{a \in A} \left\{ \min_{t} : a = s_1 \cdots s_t, s_i \in S, i = 1, \cdots, t \right\}.$$

It is remarkable (but, indeed, natural) that most networks that have been considered for large parallel processing systems (including hypercubes, torus grids, cube-connected-cycles and butterfly networks)

TABLE I
ALGEBRAIC SYMMETRY AS AN ORGANIZING PRINCIPLE FOR PARALLEL PROCESSING

Δ	D								
	2	3	4	5	6	7	8	9	10
3	10	20	38	70	128	184	320	540	938
4	15	40	95	364	734	1081*	2943*	7439*	15657*
5	24	70	182	532	2742	**4368**	11200	33600	**123120**
6	32	105	**355**	**1081**	7832	**13310**	**50616**	**202464**	**682080**
7	50	128	**506**	**2162**	10554	**39732**	140000	**911088**	2002000
8	57	**203**	842	**3081**	39258	**89373**	**455544**	**1822175**	3984120
9	74	585	1248	**6072**	74954	215688	910000	3019632	15686400
10	91	650	1820	**12144**	132932	486837	2002000	7714494	47059200

are Cayley graphs. A standard reference on Cayley graphs is [7]. For a Cayley graph description of the cube-connected-cycles see [10].

Symmetry immediately provides the following advantage for the design problem considered here: to compute the diameter of a Cayley graph it is only necessary to compute the distances from a single node to all others. Furthermore, the compactness of an algebraic description allows for an efficient computational search strategy.

Our results were obtained by experimental computing with relatively simple programs on small machines (an IBM PC and a VAX 11/780). The programs followed closely the above expression for the diameter of a Cayley graph. Having focused (by setting the appropriate program parameters) on a particular kind of matrix group, and on a choice of cardinality for the generating set (hence the degree of the resulting graph), the diameter was computed for repeated random choices of the generating set until (in the favorable case) a new record was obtained. Consonant with the above expression for the diameter, this is done by starting with the identity of the group as the *live* set, multiplying the elements of the live set with the elements of the generator set, recording any new elements obtained (the new live set) in a large array representing all elements in the target group, and repeating this until no new elements are obtained. The number of repetitions until this occurs is the diameter.

The reader may reasonably wonder about several things, beginning with the large number of authors of this note and including perhaps the question of whether some voodoo was employed in choosing the target groups and in exploring the search space of generator sets. The explanation of the first is simply that exploration of this approach to this design problem has continued among us at a low level for a number of years beginning with the seminal work of the author subset: Carlsson and Sexton. Although we have tried several "sophisticated" heuristics for choosing groups and generators, we must honestly report that none of these has proven better than simple and straightforward random search, with the exception of the nearly obvious guidance that one should choose a nonabelian group! Several of our record-breaking constructions employ upper-triangular matrix groups, but we are unable to explain why these worked better than other possibilities.

Thus, in some sense these results are less interesting than one might at first suppose, although the above information may underscore our main point: the power of an algebraic approach (even a simple one). A sophisticated understanding of what is possible by the method of Cayley graphs would be highly desirable, but it seems to present a difficult mathematical problem.

The next section describes some examples of our constructions and the associated costs of our computational explorations.

III. EXAMPLE CONSTRUCTIONS

Given that a "winning" set of generators exists for a group it would be interesting to know the expected time for random search to discover a winning set. We have no real information on this (it would seem to be a difficult mathematical problem to give any bounds), but we do indicate in the example descriptions that follow the time that was required for the particular search that uncovered the construction as a rough indication of the amount of computational effort involved. About half of the record-breaking constructions that we report here (the ones of smaller order!) were obtained on a PC, by a search program running in some cases for only a few minutes and in some cases for a few days. For the approach that we have taken memory is a more important computational bottleneck than speed.

In what follows $\mathrm{GL}[n,q]$ denotes the (*general linear*) group of $n \times n$ matrices with entries in the finite field with q elements (since below q is always a prime, this is just the integers mod q), and $\mathrm{SL}[n,q]$ is the *special linear* subgroup of $\mathrm{GL}[n,q]$ consisting of those matrices with determinant 1.

Example 1: Degree 5, diameter 7: 4368 vertices.

This is a Cayley graph on the subgroup of GL[2,13] consisting of the matrices with determinant in the set $\{1,-1\}$. The generators are the following elements together with their inverses.

$$\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \text{order } 2 \quad \begin{bmatrix} 11 & 2 \\ 8 & 12 \end{bmatrix} \text{order } 52 \quad \begin{bmatrix} 11 & 4 \\ 7 & 5 \end{bmatrix} \text{order } 14.$$

The discovery time for this construction was approximately 10 hours on an IBM PC for a small Pascal program.

Example 2: Degree 8, diameter 7: 89 373 vertices

This is a Cayley graph on a subgroup of GL[3,31]. The generators are the following elements together with their inverses.

$$\begin{bmatrix} 1 & 12 & 10 \\ 0 & 1 & 15 \\ 0 & 0 & 25 \end{bmatrix} \text{order } 93 \quad \begin{bmatrix} 1 & 25 & 15 \\ 0 & 1 & 4 \\ 0 & 0 & 5 \end{bmatrix} \text{order } 93$$
$$\begin{bmatrix} 1 & 29 & 29 \\ 0 & 1 & 16 \\ 0 & 0 & 5 \end{bmatrix} \text{order } 93 \quad \begin{bmatrix} 1 & 27 & 5 \\ 0 & 1 & 8 \\ 0 & 0 & 1 \end{bmatrix} \text{order } 31$$

The discovery time for this construction was approximately 3 hours of CPU time on a VAX 11/780.

Example 3: Degree 10, diameter 5: 12 144 vertices.

This is a Cayley graph on the group SL[2,23]. The generators are the following elements together with their inverses.

$$\begin{bmatrix} 9 & 0 \\ 18 & 18 \end{bmatrix} \text{order } 11 \quad \begin{bmatrix} 13 & 10 \\ 18 & 21 \end{bmatrix} \text{order } 11 \quad \begin{bmatrix} 9 & 10 \\ 0 & 17 \end{bmatrix} \text{order } 22$$
$$\begin{bmatrix} 14 & 7 \\ 19 & 3 \end{bmatrix} \text{order } 22 \quad \begin{bmatrix} 18 & 13 \\ 17 & 20 \end{bmatrix} \text{order } 24.$$

The discovery time for this construction was approximately 2 hours on an IBM PC.

TABLE II

Parameters	Order	Group	Generators: order $S = S \cup S^{-1}$
degree 5 diameter 7	4368	subgroup of GL[2,13]	[0,1,1,0]:2 [11,4,7,5]:14 [11,2,8,12]:52
degree 5 diameter 8	8788	subgroup of GL[3,13]	[1,0,4,0,1,0,0,0,12]:2 [1,5,6,0,1,9,0,0,5]:52 [1,2,12,0,1,8,0,0,5]:52
degree 5 diameter 9	25308	PSL[2,37]	[0,36,1,0]:2 [34,26,34,1]:37 {2,16,11,33]:37
degree 5 diameter 10	123120	GL[2,19]	[0,1,1,0]:2 [11,16,0,15]:18 [16,11,2,0]:45
degree 6 diameter 4	355	subgroup of GL[2,71]	[54,66,0,1]:5 [5,43,0,1]:5 [57,38,0,1]:5
degree 6 diameter 5	1081	subgroup of GL[2,47]	[7,20,0,1]:23 [6,33,0,1]:23 [9,42,0,1]:23
degree 6 diameter 7	13310	subgroup of GL[3,11]	[1,2,7,0,1,0,0,0,10]:22 [1,5,2,0,1,2,0,0,4]:55 [1,6,10,0,1,3,0,0,5]:55
degree 6 diameter 8	50616	SL[2,37]	[32,24,35,2]:19 [23,16,28,34]:36 [12,24,15,27]:37
degree 6 diameter 9	202464	subgroup of GL[2,37]	[25,1,31,1]:36 [12,35,23,30]:76 [12,4,28,16]:152
degree 6 diameter 10	682080	GL[2,29]	[28,10,8,8]:28 [17,13,16,27]:28 [3,4,27,14]:840
degree 7 diameter 4	506	subgroup of GL[2,23]	[22,1,0,1]:2 [13,16,0,1}:11 [3,16,0,1]:11 [19,12,0,1]:22
degree 7 diameter 5	2162	subgroup of GL[2,47]	[46,1,0,1]:2 [4,20,0,1]:23 [20,27,0,1]:46 [29,14,0,1]:46
degree 7 diameter 7	39732	PSL[2,43]	[0,42,1,0]:2 [18,16,38,41]:22 [34,2,37,6]:22 [8,28,14,33]:43
degree 7 diameter 8	101232	subgroup of GL[2,37]	[0,1,1,0]:2 [21,34,17,17]:6 [21,1,4,2]:9 [27,26,4,8]:74
degree 7 diameter 9	911088	subgroup of GL[2,37]	[0,1,1,0]:2 [23,17,14,26]:18 [25,16,13,6]:36 [27,33,19,22]:684
degree 7 diameter 10	1822176	GL[2,37]	[0,1,1,0]:2 [1,19,14,16]:17 [36,1,12,0]:18 [35,28,34,12]:456
degree 8 diameter 3	203	subgroup of GL[2,29]	[16,9,0,1]:7 [16,21,0,1]:7 [25,15,0,1]:7 [25,9,0,1]:7
degree 8 diameter 4	812	subgroup of GL[2,29]	[12,1,0,1]:4 [20,24,0,1]:7 [6,27,0,1]:14 [15,18,0,1]:28
degree 8 diameter 5	3081	subgroup of GL[2,79]	[46,43,0,1]:13 [49,72,0,1]:39 [19,26,0,1]:39 [13,13,0,1]:39
degree 8 diameter 7	89373	subgroup of GL[2,31]	[1,4,25,0,1,23,0,0,1]:31 [1,29,29,0,1,16,0,0,5]:93 [1,12,10,0,1,15,0,0,25]:93 [1,6,17,0,1,24,0,0,5]:93
degree 8 diameter 8	455544	subgroup of GL[2,37]	[21,9,17,5]:57 [0,26,3,1]:171 [28,32,33,33]:171 [9,34,25,16]:342
degree 8 diameter 9	1822176	GL[2,37]	[12,13,34,33]:18 [36,6,20,10]:36 [35,3,19,35]:684 [26,10,36,31]:1368
degree 9 diameter 5	6072	PSL[2,23]	[0,22,1,0]:2 [2,18,4,2]:11 [10,1,21,16]:24 [6,19,4,9]:24 [22,0,1,22]:46
degree 9 diameter 8	682080	GL[2,29]	[0,1,1,0]:2 [5,22,18,26]:14 [17,15,21,4]:840 [2,5,10,21]:840 [23,12,11,21]:840
degree 10 diameter 5	12144	SL[2,23]	[9,0,18,18]:11 [13,10,18,21]:11 [9,10,0,17]:22 [14,7,19,3]:22 [18,13,17,20]:24
degree 10 diameter 8	1822176	subgroup of GL[2,37]	[21,12,22,5]:57 [9,12,6,26]:456 [35,10,17,32]:684 [5,31,35,14]:684 [11,3,33,7]:1368

IV. The Constructions

During the publication process for this note we have become aware of a new approach to this design problem, not based on Cayley graphs, that shares with our approach the aspects of 1) a significant exploitation of symmetry, and 2) computational exploration [5]. This has had the effect on this note of removing from "bold" seven entries of the original version of Table I. We have retained the descriptions of the Cayley graphs that gave those entries in the table that follows, as they may still be of interest by virtue of their vertex symmetry or other properties.

V. Conclusions

Our main contribution in this brief presentation is the demonstration of the power of an algebraic approach to the problem of constructing large networks of a given degree and diameter. The success of the relatively limited search we have so far conducted seems to indicate that further exploration based on Cayley graphs may be productive. Major problems relevant to applicants in parallel processing and not addressed here concern message routing and data exchange. Solutions are likely to be much more complicated in such networks as we have described than in the familiar (Cayley graph) networks of hypercubes and cube-connected cycles, and this remains an area for further research.

References

[1] F. Annexstein, M. Baumslag, and A. L. Rosenberg, "Group action graphs and parallel architectures," *COINS Tech. Rep. 87-133,* Univ. of Massachusetts, Amherst, MA, 1987.

[2] S. B. Akers and B. Krishnamurthy, "On group graphs and their fault-tolerance," *IEEE Trans. Comput.,* vol. C-36, pp. 885–888, 1987.

[3] J. C. Bermond, C. Delorme, and J. J. Quisquater, "Strategies for interconnection networks: Some methods from graph theory," *J. Parallel and Distributed Comput.,* vol. 3, pp. 433–449, 1986.

[4] J. Bond, C. Delorme, and W. F. de La Vega, "Large Cayley graphs with small degree and diameter," Rapport de Recherche no. 392, LRI, Orsay, France, 1987.

[5] J. C. Bermond, private communication, June 1989.

[6] R. Bar-Yehuda and T. Etzion, "Connections between two cycles—A new design of dense processor interconnection networks," Tech. Rep. 528, Dep. Comput. Sci., Technion, Haifa, Israel, 1988.

[7] N. Biggs, *Algebraic Graph Theory.* London, England: Cambridge University Press, 1974.

[8] D. V. Chudnovsky, G. V. Chudnovsky, and M. M. Denneau, "Regular graphs with small diameter as models for interconnection networks," in *Proc. Third Int. Conf. Supercomput.,* Boston, MA, May 1988, pp. 232–239.

[9] F. R. K. Chung, "Diameters of graphs: Old problems and new results," *Congressus Numerantium 60,* pp. 295–317, 1987.

[10] G. E. Carlsson, J. E. Cruthirds, H. B. Sexton, and C. G. Wright, "Interconnection networks based on a generalization of cube-connected cycles," *IEEE Trans. Comput.,* vol. C-34, pp. 769–777, 1985.

[11] V. Faber, "Global communication algorithms for hypercubes and other Cayley coset graphs," Tech. Rep., Los Alamos National Lab., Los Alamos, NM, 1988.

[12] M. Fellows, "A category of graphs for parallel processing," Tech. Rep. Dep. Comput. Sci., Univ. of Idaho, Moscow, ID, 1988.