Tight Bounds on the Diameter of Gaussian Cubes

DING-MING KWAI AND BEHROOZ PARHAMI

Department of Electrical and Computer Engineering, University of California, Santa Barbara, CA 93106–9560, USA
Email: parhami@ece.ucsb.edu

Gaussian cubes are derived by removing links from a hypercube in a periodic fashion. By varying the partition parameter, one can obtain networks with different characteristics, while maintaining a basic framework for computation and communication. Unfortunately, such networks are in general not regular, making it difficult to derive their topological properties explicitly. In this paper, we study the diameter of Gaussian cubes and show the trade-off between cost and performance.

Received February 2, 1997; revised November 20, 1997

1. INTRODUCTION

Gaussian cubes [1] have been proposed as interconnection networks for scalable distributed systems in view of their low node degree, small diameter, and large bisection width. Like 2^n-node cube connected cycles (CCC) [2], a Gaussian cube can be viewed as a subgraph of an n-dimensional hypercube or binary n-cube, derived by removing certain links from the nodes [3, 4]. The removal of links is performed in a periodic fashion so that nodes lacking different types of dimensional links can be clustered into groups; each group has the complete set of dimensional links, making them locally accessible to group members.

Using the group size as a parameter, one can obtain architectures with different characteristics while maintaining the same basic framework for computation and communication. To construct a Gaussian cube, one begins with a hypercube of the same size and proceeds in steps. By successively removing some of the existing links from a Gaussian cube, the group size is increased in powers of 2. Figures 1 and 2 show examples of a binary 4-cube, cube connected cycles, and Gaussian cubes with three different group sizes, where nodes are ordered by a Gray code.

The construction of Gaussian cubes by removing links from a hypercube, however, does not preserve the regularity of node degree. Because of this, the determination of diameter is quite difficult and a loose upper bound on diameter is all that is known [1]. In this paper, we derive a tighter bound. Our presentation is organized as follows. Section 2 defines the Gaussian cube. Section 3 shows special cases in which the exact diameter can be easily obtained. Section 4 derives upper and lower bounds on the diameter. Section 5 contains our conclusions.

2. DEFINITION

A Gaussian n-cube has 2^n nodes; each node is identified by an n-bit number $A = (a_{n-1} \ldots a_1 a_0)_2$. Here and throughout, a node, its label, and the binary representation of the label will be used interchangeably. Let $Q_{(n,m)}$ denote the Gaussian n-cube with group size 2^m, where $0 \le m \le n-1$. We define $Q_{(n,m)}$ as a *pruned* n-dimensional hypercube in the following way. The dimension-i link, $0 \le i \le n-1$, exists between nodes $A = (a_{n-1} \ldots a_i \ldots a_0)_2$ and $A^{(i)} = (a_{n-1} \ldots \bar{a}_i \ldots a_0)_2$ whose addresses differ exactly in bit position i, if and only if $\sum_{h=0}^{i-1} a_h 2^h \equiv i (\mathrm{mod} 2^m)$. This definition is different from, but equivalent to, the one in [1], given that the group size must be a power of 2.

Note that the dimension-0 links remain at all nodes and the dimension-1 links at odd-numbered nodes, independent of the group size. A complete n-cube, Q_n, is simply a Gaussian cube $Q_{(n,0)}$.

In the above definition, we have restricted the maximum group size to be 2^{n-1}, instead of 2^n. This is because $Q_{(n,n)}$ is isomorphic to $Q_{(n,n-1)}$ in view of the fact that $n-1 < 2^{n-1}$ for $n \ge 1$, thus making $i \mathrm{mod} 2^n$ equal to $i \mathrm{mod} 2^{n-1}$ for $0 \le i \le n-1$. Intuitively, $Q_{(n,n-1)}$ is composed of two 2^{n-1}-node groups connected by a dimension-$(n-1)$ link between node $n-1$ and node $2^{n-1}+n-1$. Each 2^{n-1}-node group can be further decomposed into two 2^{n-2}-node subgroups connected by a dimension-$(n-2)$ link. Continuing the recursive decomposition leads to the conclusion that $Q_{(n,n-1)}$ can be treated as a tree rooted at node $n-1$. This will be further clarified in Section 4.

More generally, for $n < 2^m$, each node will be assigned a link from at most one of the remaining $n-2$ dimensions (excluding 0 and 1 as stated above). The maximum node degree is 3 for $n \ge 4$, 2 for $2 \le n \le 3$, and 1 for $n = 1$. In deriving the diameter, we shall exclude the trivial cases of $n \le 3$ whose pruning results in a linear array.

In the Gaussian cube $Q_{(n,m)}$ where $n \ge 2^m$, node $2^m - 1$ has $m+1$ neighbours along dimensions $0, 1, 3, \ldots, 2^m - 1$, respectively. Any other neighbour, if it exists, must be along dimension $x2^m - 1$ with x being a positive integer. Since $2^m - 1 < x2^m - 1 \le n-1$, we have $1 < x \le \lfloor n/2^m \rfloor$,

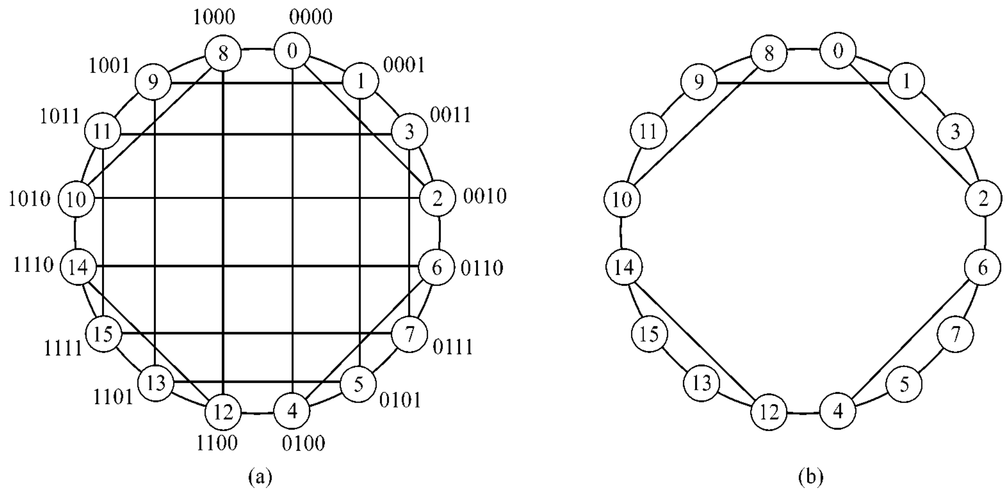

FIGURE 1. (a) Binary 4-cube Q_4 or $Q_{(4,0)}$; (b) cube connected cycles.

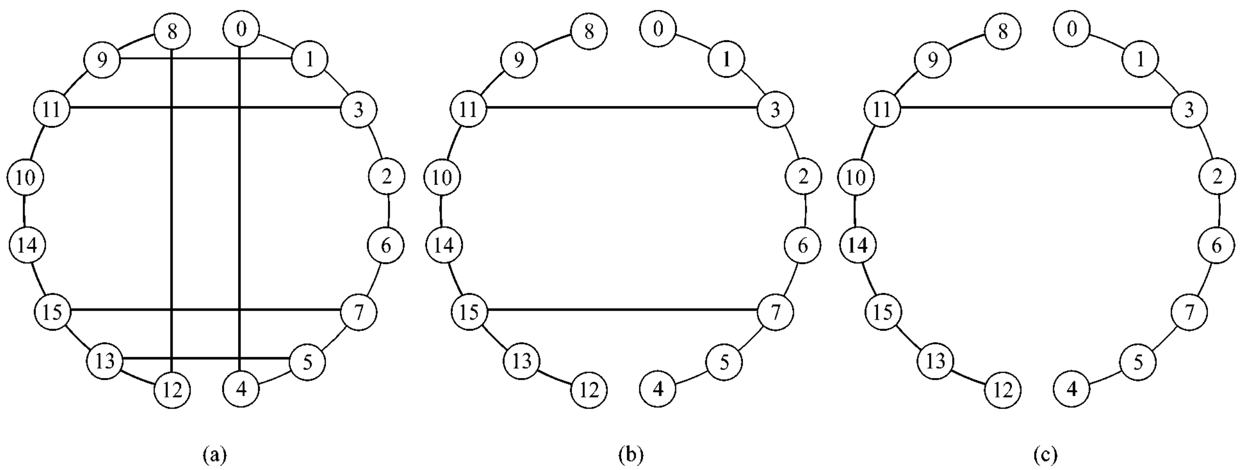

FIGURE 2. Gaussian 4-cubes: (a) $Q_{(4,1)}$; (b) $Q_{(4,2)}$; (c) $Q_{(4,3)}$.

leading to the maximum node degree $\lfloor n/2^m \rfloor + m$. In a similar manner, the minimum node degree $\lfloor (n-1)/2^m \rfloor + 1$ can be derived by considering the connectivity of node 0.

3. SPECIAL CASES

Given the maximum and minimum node degrees of $Q_{(n,m)}$, there are two special cases where the selection of group size parameter m results in a regular Gaussian n-cube. For these special cases we can derive the exact diameter. Clearly $Q_{(n,0)}$ is the binary n-cube from which no link has been removed; both the node degree and diameter are equal to n. If $m = 1$ and n is odd, then $Q_{(n,1)}$ is regular of degree $(n+1)/2$. In such a case, the even-numbered nodes are assigned even-dimension links and the odd-numbered nodes assigned odd-dimension links, in addition to the dimension-0 links for all nodes. Since the even and odd dimensions can be permuted without changing the connectivity, the network is node symmetric as well.

The diameter of $Q_{(n,1)}$, where n is odd, can be seen to be $n+1$ by counting the number of routing steps required to send a message from a source node $(a_{n-1} \ldots a_1 0)_2$ to a destination node $(\bar{a}_{n-1} \ldots \bar{a}_1 0)_2$ whose addresses are even numbers that differ in the $n-1$ most significant bits. It takes $(n-1)/2$ steps to traverse all even dimensions, leading to the node $(\bar{a}_{n-1} a_{n-2} \ldots \bar{a}_2 a_1 0)_2$, one step to route to the odd-numbered node $(\bar{a}_{n-1} a_{n-2} \ldots \bar{a}_2 a_1 1)_2$, then $(n-1)/2$ steps to traverse all odd dimensions, leading to $(\bar{a}_{n-1} \bar{a}_{n-2} \ldots \bar{a}_2 \bar{a}_1 1)_2$, and finally, one step to route to the destination node $(\bar{a}_{n-1} \bar{a}_{n-2} \ldots \bar{a}_2 \bar{a}_1 0)_2$ via dimension 0.

One can easily verify that the diameter of $Q_{(n,1)}$, where n is even, is also $n+1$. Such Gaussian cubes, which have node degree of $n/2 - 1$ or $n/2 + 1$ for even- or odd-numbered nodes, respectively, may be compared with a two-level hierarchical cubic network [5] which has regular node degree $n/2 + 1$ and a smaller diameter $\lfloor 3n/4 \rfloor + 1$. The regularity of node degree and reduction in diameter can be attributed to the links added between nodes whose addresses are complements of each other. Generalization that allows more than two hierarchical levels is certainly viable and has been shown in [6, 7].

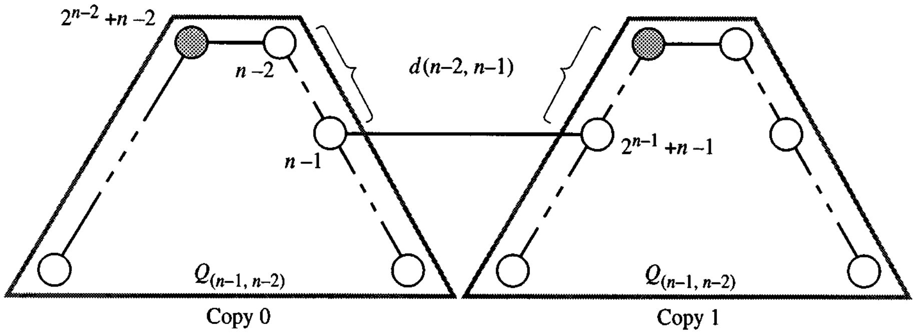

FIGURE 3. The construction of Gaussian cube $Q_{(n,n-1)}$ from $Q_{(n-1,n-2)}$.

4. BOUNDS ON DIAMETER

Before deriving the bounds on the diameter of Gaussian cube $Q_{(n,m)}$, it is helpful to discuss $Q_{(n,n-1)}$ whose diameter is the largest of all $Q_{(n,m)}$, $0 \le m \le n-1$, given the dimension n. This is obvious in view of the fact that $Q_{(n,m)}$ contains additional links, compared with $Q_{(n,n-1)}$.

$Q_{(n,n-1)}$ can be constructed by taking two copies of $Q_{(n-1,n-2)}$, appending one with the most significant bit 0 and the other with 1, and connecting the two copies with a new dimension $n-1$ link between node $n-1$ and node $2^{n-1}+n-1$ (see Figure 3). For $n \ge 4$, $Q_{(n,n-1)}$ has the maximum node degree of 3.

Let $D_{(n,m)}$ denote the diameter of $Q_{(n,m)}$ and $d(u,v)$ denote the distance from node u to node v along a shortest path in $Q_{(n,m)}$. Since for $n \ge 4$ the diameter of $Q_{(n,n-1)}$ is larger than the diameter of $Q_{(n-1,n-2)}$ by the distance between the two shaded nodes in Figure 3, we have the inductive formula

$$D_{(n,n-1)} = D_{(n-1,n-2)} + 2[d(n-2, n-1) + 1]. \quad (1)$$

Expanding the above expression and substituting the initial condition $D_{(3,2)} = 7$, we get

$$D_{(n,n-1)} = 2\sum_{h=3}^{n-1} d(h-1, h) + 2n + 1. \quad (2)$$

Clearly $d(h-1, h)$ cannot be less than the Hamming distance between $h-1$ and h. The sum of Hamming distances for consecutive numbers from 0 to $n-1$ can be calculated exactly as

$$n-1+\lfloor (n-1)/2\rfloor + \lfloor (n-1)/4\rfloor + \ldots + 1 = \sum_{h=0}^{\lfloor \log_2(n-1)\rfloor} \lfloor (n-1)/2^i\rfloor \ge 2n-3-\lfloor \log_2(n-1)\rfloor .$$

Hence the sum of Hamming distances for consecutive numbers from 2 to $n-1$ is lower bounded by $2n-6-\lfloor \log_2(n-1)\rfloor$. Furthermore, if h is a multiple of 4 (i.e. a leaf node of $Q_{(3,2)}$, the routing from $h-1$ to h takes at least one more step to gain access to the required dimensional link. This leads to

$$\sum_{h=3}^{n-1} d(h, h-1) \ge 2n-6+\lfloor (n-1)/4\rfloor - \lfloor \log_2(n-1)\rfloor . \quad (3)$$

By noting that the sum of Hamming distances for consecutive numbers from 0 to $n-1$ is strictly less than $2(n-1)$, the sum of Hamming distances for consecutive numbers from 2 to $n-1$ can be upper bounded by $2n-6$. Again, consider the case where h is a multiple of 4. The number of extra steps needed to gain access to the required dimensional link is at most seven. Excluding at least two steps taken locally along dimensions 0 and 1, we have

$$\sum_{h=3}^{n-1} d(h, h-1) \le 2n-6+5\lfloor (n-1)/4\rfloor . \quad (4)$$

Combining eqns (2)–(4) yields

$$6n + 2\lfloor (n-1)/4\rfloor - 2\lfloor \log_2(n-1)\rfloor - 11 \le D_{(n,n-1)} \le 6n + 10\lfloor (n-1)/4\rfloor - 11, \quad (5)$$

implying that $Q_{(n,m)}$ also has logarithmic diameter.

$Q_{(n,m)}$ can be viewed as an $(n-m)$-dimensional hypercube in which each node has been replaced with a $Q_{(m,m-1)}$. Any longest path of $Q_{(n,m)}$ then includes a longest path of $Q_{(m,m-1)}$ and traverses $n-m$ dimensions:

$$D_{(n,m)} \ge n-m+D_{(m,m-1)} \ge n+5m+2\lfloor (m-1)/4\rfloor - 2\lfloor \log_2(m-1)\rfloor - 11. \quad (6)$$

For $m \le 3$, the 2^m nodes of $Q_{(m,m-1)}$ are connected in a line. Hence to visit the $n-m$ dimensions requires the traversal of all the 2^m nodes. An upper bound can be immediately obtained as

$$D_{(n,m)} \le n-m+2(2^m-1) \quad (7)$$

which is tight for $0 \le m \le 3$.

However, for $m \ge \log_2 n$, the above upper bound becomes too large to be useful. Recall that each $Q_{(m,m-1)}$ group of $Q_{(n,m)}$ has links connected to $n-m-1$ other groups. Thus $D_{(n,m)}$ is at least $n-m-1$ less than $D_{(n,n-1)}$:

$$D_{(n,m)} \le D_{(n,n-1)} - n + m + 1. \quad (8)$$

Combining eqns (7) and (8) yields

$$D_{(n,m)} \le \min\{n-m+2^{m+1}-2,\, 5n+10\lfloor (n-1)/4\rfloor + m - 10\}. \quad (9)$$

TABLE 1. Diameter $D_{(n,m)}$ of Gaussian cube $Q_{(n,m)}$

m \ n	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
1		3	4	5	6	7	8	9	10	11	12	13	14	15	16	17
2			7	8	9	10	11	12	13	14	15	16	17	18	19	20
3				11	16	17	18	19	20	21	22	23	24	25	26	27
4					23	24	25	26	31	32	33	34	39	40	41	42
5						27	32	33	44	45	48	49	58	59	60	61
6							33	34	45	46	49	50	59	60	61	62
7								37	48	49	52	53	62	63	64	65
8									51	52	55	56	65	66	67	68
9										55	60	61	70	71	72	73
10											61	62	71	72	73	74
11												65	74	75	76	77
12													77	78	79	80
13														81	86	87
14															87	88
15																91

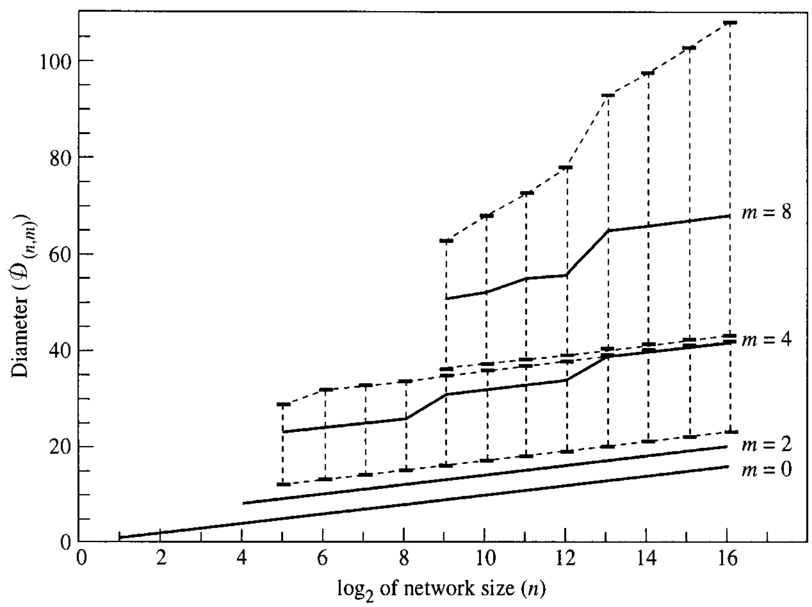

FIGURE 4. Diameter $D_{(n,m)}$ of Gaussian cube $Q_{(n,m)}$.

Table 1 lists the exact diameter of $Q_{(n,m)}$ with n ranging from 1 to 16 and m from 0 to $n-1$. Figure 4 shows selected data from Table 1 along with the corresponding lower and upper bounds, drawn as error bars.

5. CONCLUSION

We have derived tighter bounds on the diameter of the Gaussian cube $Q_{(n,m)}$. It is intuitively clear that the diameter is a monotone increasing function of dimension n and group size parameter m. However, it seems difficult, if not impossible, to express the exact diameter in closed form. Because of the irregularity and asymmetry, routing on Gaussian cubes is bound to be more difficult than on a hypercube.

To show the possibility of trading cost with performance in Gaussian cubes, we use a composite measure, defined as the product of maximum node degree and diameter. Figure 5 plots the degree-diameter product of Gaussian 16-cubes with the group size parameter m varying from 0 to 15. There actually exists an optimal group size that minimizes this figure of merit. We note that this optimal group size cannot be larger than n, since beyond this point the maximum node degree becomes a constant but the diameter continues to grow.

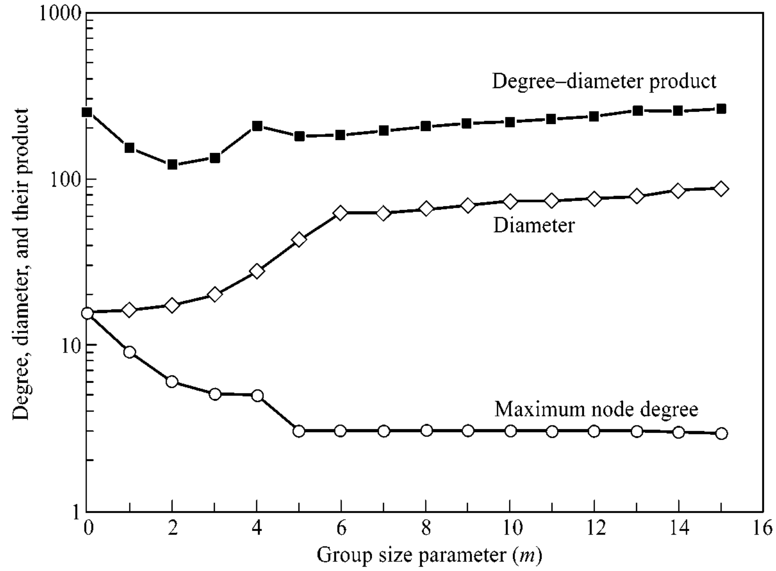

FIGURE 5. Cost–performance trade-offs in Gaussian 16-cubes.

After deriving the degree–diameter products for all $Q_{(n,m)}$, $1 \leq n \leq 16$ and $0 \leq m \leq n-1$, we found that the optimal group size is equal to one for $1 \leq n \leq 3$, two for $4 \leq n \leq 9$, and four for $10 \leq n \leq 16$. Though not evident from the data in Table 1, we have verified that the optimal group size remains at four for $17 \leq n \leq 19$.

A similar figure may be drawn for maximum node degree versus bisection width. In this case, the degraded performance can hardly be justified by the reduced cost, since each doubling of the group size reduces the bisection bandwidth to half, while there is no corresponding reduction in the maximum node degree.

REFERENCES

[1] Hsu, W.-J., Chung, M. J., and Hu, Z. (1996) Gaussian networks for scalable distributed systems. *Comput. J.*, **39**, 417–426.

[2] Preparata, F. P. and Vuillemin, J. E. (1981) The cube-connected cycles: a versatile network for parallel computation. *Commun. ACM*, **24**, 300–309.

[3] Kwai, D.-M. and Parhami, B. (1996) A generalization of hypercubic networks based on their chordal ring structures. *Parallel Process. Lett.*, **6**, 469–477.

[4] Kwai, D.-M. and Parhami, B. (1997) A class of fixed-degree Cayley-graph interconnection networks derived by pruning k-ary n-cubes. *Proc. Int. Conf. on Parallel Processing*, pp. 92–95.

[5] Ghose, K. and Desai, R. (1995) Hierarchical cubic networks. *IEEE Trans. Parallel Distrib. Sys.*, **6**, 427–435.

[6] Ziavras, S. G. (1994) RH: a versatile family of reduced hypercube interconnection networks. *IEEE Trans. Parallel Distrib. Sys.*, **5**, 1210–1220.

[7] Yeh, C.-H. and Parhami, B. (1996) Recursive hierarchical swapped networks: versatile interconnection architectures for highly parallel systems. *Proc. IEEE Symp. on Parallel and Distributed Processing*, New Orleans, LA, pp. 453–460.