Information Sciences 181 (2011) 239–256

Contents lists available at ScienceDirect

Information Sciences

journal homepage: www.elsevier.com/locate/ins

Augmented k-ary n-cubes

Yonghong Xiang, Iain A. Stewart *

Department of Computer Science, University of Durham, Science Labs, South Road, Durham DH1 3LE, UK

ARTICLE INFO

Article history:
Received 7 January 2009
Received in revised form 13 April 2010
Accepted 1 September 2010

Keywords:
Interconnection networks
Parallel computing
k-ary n-cubes
Augmented cubes

ABSTRACT

We define an interconnection network $AQ_{n,k}$ which we call the augmented k-ary n-cube by extending a k-ary n-cube in a manner analogous to the existing extension of an n-dimensional hypercube to an n-dimensional augmented cube. We prove that the augmented k-ary n-cube $AQ_{n,k}$ has a number of attractive properties (in the context of parallel computing). For example, we show that the augmented k-ary n-cube $AQ_{n,k}$: is a Cayley graph, and so is vertex-symmetric, but not edge-symmetric unless $n = 2$; has connectivity $4n - 2$ and wide-diameter at most $max\{(n-1)k-(n-2), k+7\}$; has diameter $\lfloor \frac{k}{3} \rfloor + \lceil \frac{k-1}{3} \rceil$, when $n = 2$; and has diameter at most $\frac{k}{4}(n+1)$, for $n \geqslant 3$ and k even, and at most $\frac{k}{4}(n+1)+\frac{n}{4}$, for $n \geqslant 3$ and k odd.

1. Introduction

Hypercubes are perhaps the most well known of all interconnection networks for parallel computing, given their basic simplicity, their generally desirable topological and algorithmic properties, and the extensive investigation they have undergone (not just in the context of parallel computing but also in discrete mathematics in general; see, for example, [27] for some essential properties of hypercubes). However, a multitude of different interconnection networks have been devised and developed in a continuing search for improved performance, with many of these networks having hypercubes at their roots. Amongst these generalisations of hypercubes are k-ary n-cubes [14], augmented cubes [12], cube-connected cycles [26], twisted cubes [19], twisted n-cubes [18], crossed cubes [16], folded hypercubes [17], Mcubes [30], Möbius cubes [13], generalised twisted cubes [11], shuffle cubes [24], k-skip enhanced cubes [31], twisted hypercubes [22], supercubes [29], and Fibonacci cubes [20].

Perhaps the most popular of these generalisations are the k-ary n-cubes [14]. A k-ary n-cube Q_n^k is essentially a 'k-bit version' of a '2-bit' hypercube in that vertices are represented by n-tuples of integers from $\{0, 1, \ldots, k-1\}$ so that two vertices are joined by an edge if, and only if, their representations are identical save in one bit position, where in that position the bits differ by 1 modulo k (thus, a k-ary 2-cube, for example, is just a $k \times k$ torus). It turns out that k-ary n-cubes have similar properties to hypercubes yet provide more flexibility with regard to incorporating more processors; for the two parameters available, k and n, allow us to regulate the degree of the nodes yet still incorporate large numbers of processors, although usually at a cost to some other property such as the diameter or the connectivity. Some properties of the k-ary n-cube are that it: has k^n vertices and nk^n edges; has diameter $n\lfloor \frac{k}{2} \rfloor$; has wide-diameter $n\lfloor \frac{k}{2} \rfloor + 1$, when $n \geqslant 3$ or when $n = 2$ and $k \geqslant 6$ [21]; has connectivity $2n$ [10]; is a Cayley graph, and so is vertex-symmetric [7], and also edge-symmetric [4]; and has an $O(nk)$ time optimal routing algorithm [3,15]. A number of distributed memory multiprocessors have been built with a k-ary n-cube forming the underlying topology, such as the Mosaic [28], the iWARP [9], the J-machine [25], the Cray T3D [23], the Cray T3E [2], and the IBM Blue Gene [1].

* Corresponding author.
E-mail addresses: yonghong.xiang@durham.ac.uk (Y. Xiang), i.a.stewart@durham.ac.uk (I.A. Stewart).

0020-0255/$ - see front matter
doi:10.1016/j.ins.2010.09.005

Another generalisation of hypercubes are augmented cubes, recently proposed by Choudum and Sunitha [12] as improvements over hypercubes. Hypercubes and augmented cubes (of the same dimensions) have the same sets of vertices. However, whereas the recursive construction of an n-dimensional hypercube is to take two copies of an $(n-1)$-dimensional hypercube and join corresponding pairs of vertices, the recursive construction of an n-dimensional augmented cube AQ_n is to take two copies of an $(n-1)$-dimensional augmented cube and as well as joining corresponding pairs of vertices, pairs of vertices of Hamming distance $n-1$ are also joined (that is, vertices that are different in every component). Choudum and Sunitha show that an n-dimensional augmented cube AQ_n: has 2^n vertices and $n2^n$ edges; has diameter $\lceil \frac{n}{2} \rceil$; has connectivity $2n-1$; is a Cayley graph and so is vertex-symmetric; and has an $O(n)$ time optimal routing algorithm.

In this paper, and inspired by [12], we extend a k-ary n-cube in a manner analogous to the extension of an n-dimensional hypercube to an n-dimensional augmented cube. Our definition of an *augmented k-ary n-cube* $AQ_{n,k}$, in comparison with that in [12], is not a straightforward generalisation; however, we believe that it does reflect the essence of the extension in [12], and our structural results bear this out. We give two different definitions of an augmented k-ary n-cube in Section 2 and show that they yield the same interconnection network. In Section 3, we show that an augmented k-ary n-cube $AQ_{n,k}$ is vertex-symmetric and, furthermore, a Cayley graph, though not edge-symmetric unless $n = 2$. In Section 4, we show that an augmented k-ary n-cube $AQ_{n,k}$ has connectivity $4n-2$, and that we can build a set of $4n-2$ mutually disjoint paths joining any two distinct vertices so that the path of maximal length has length at most $max\{(n-1)k-(n-2), k+7\}$; that is, $AQ_{n,k}$ has wide-diameter at most $max\{(n-1)k-(n-2), k+7\}$. In Section 5, we examine the diameter of the augmented k-ary n-cube $AQ_{n,k}$ and show that the diameter of the augmented k-ary 2-cube $AQ_{2,k}$ is $\lfloor \frac{k}{3} \rfloor + \lceil \frac{k-1}{3} \rceil$. We also show that the diameter of the augmented k-ary n-cube $AQ_{n,k}$ is at most $\frac{k}{4}(n+1)$, when $n \geqslant 3$ and k is even, and at most $\frac{k}{4}(n+1)+\frac{n}{4}$, when $n \geqslant 3$ and k is odd. Our conclusions are presented in Section 6.

2. Basic definitions

We assume throughout that arithmetic on tuple elements is modulo k, and we denote tuples of elements by bold type. Recall the definition of the *k-ary n-cube* Q_n^k: the vertex set $V(Q_n^k)$ is $\{(a_n, a_{n-1}, \ldots, a_1): 0 \leqslant a_i \leqslant k-1\}$; and the edge set $E(Q_n^k)$ is $\{(\mathbf{u}, \mathbf{v}) : \mathbf{u} = (u_n, u_{n-1}, \ldots, u_1), \mathbf{v} = (v_n, v_{n-1}, \ldots, v_1)$, either $u_i = v_i - 1$ or $u_i = v_i + 1$, for some i, and $u_j = v_j$, for all $i \neq j\}$. Whilst we regard all graphs defined in this paper as undirected, our definitions define all edges from the perspective of a given vertex. Thus, in our definition of Q_n^k we define the (undirected) edge $(\mathbf{u}, \mathbf{v})$ twice: once from the perspective of $\mathbf{u}$, as the edge $(\mathbf{u}, \mathbf{v})$; and once from the perspective of $\mathbf{v}$, as the edge $(\mathbf{v}, \mathbf{u})$. The reason we do this is that later we shall define paths in our graphs and an undirected edge will be regarded differently depending upon the direction it is being traversed in the path. The following definition adheres to this convention.

Definition 1. Let $n \geqslant 1$ and $k \geqslant 3$ be integers. The *augmented k-ary n-cube* $AQ_{n,k}$ has k^n vertices, each labelled by an n-bit string $(a_n, a_{n-1}, \ldots, a_1)$, with $0 \leqslant a_i \leqslant k-1$, for $1 \leqslant i \leqslant n$. There is an edge joining vertex $\mathbf{u} = (u_n, u_{n-1}, \ldots, u_1)$ to vertex $\mathbf{v} = (v_n, v_{n-1}, \ldots, v_1)$ if, and only if:

- $v_i = u_i - 1$ (resp. $v_i = u_i + 1$), for some $1 \leqslant i \leqslant n$, and $v_j = u_j$, for all $1 \leqslant j \leqslant n, j \neq i$; call the edge $(\mathbf{u}, \mathbf{v})$ an $(i,-1)$-*edge* (resp. an (i,+1)-*edge*); or
- for some $2 \leqslant i \leqslant n$, $v_i = u_i - 1$, $v_{i-1} = u_{i-1} - 1, \ldots, v_1 = u_1 - 1$ (resp. $v_i = u_i + 1$, $v_{i-1} = u_{i-1} + 1, \ldots, v_1 = u_1 + 1$), $v_j = u_j$, for all $j > i$; call the edge $(\mathbf{u}, \mathbf{v})$ a $(\leqslant i,-1)$-*edge* (resp. a $(\leqslant i,+1)$-*edge*).

We emphasise that the graph $AQ_{n,k}$ is undirected but that edges are *labelled* differently, as an $(i,+1)$-edge or as an $(i,-1)$-edge, for example, according to the perceived orientation.

The augmented k-ary n-cube $AQ_{n,k}$ can also be recursively defined as follows (the proof of this fact is a simple induction).

Definition 2. Fix $k \geqslant 3$. The augmented k-ary 1-cube $AQ_{1,k}$ has vertex set $\{0, 1, \ldots, k-1\}$ and there is an edge joining vertex u to vertex v if, and only if, $v = u + 1$ or $v = u - 1$. Fix $n \geqslant 2$. Take k copies of an augmented k-ary $(n-1)$-cube $AQ_{n-1,k}$ and for the ith copy, add an extra number i as the nth bit of each vertex (all vertices have the same nth bit if they are in the same augmented k-ary $(n-1)$-cube). Four more edges are added for each vertex, namely the $(n,-1)$-edge, the $(n,+1)$-edge, the $(\leqslant n,-1)$-edge and the $(\leqslant n,+1)$-edge (as defined in Definition 1).

With respect to the above definition, we refer to the subgraph of $AQ_{n,k}$ induced by the vertices whose first component is i, for some fixed $i \in \{0, 1, \ldots, k-1\}$, as $AQ_{n-1,k}^i$ (this subgraph is clearly a copy of $AQ_{n-1,k}$).

Clearly, when $n \geqslant 2$, $AQ_{n,k}$ has k^n vertices, $(2n-1)k^n$ edges, and every vertex has degree $4n-2$.

We adopt the following notation with regard to identifying specific vertices relevant to a given vertex in $AQ_{n,k}$. Let $\mathbf{v} = (v_n, v_{n-1}, \ldots, v_1)$ be some vertex of $AQ_{n,k}$. For each $i \in \{0, 1, \ldots, k-1\}$ and each $j \in \{1, 2 \ldots, n\}$, we denote the vertex $(v_n, v_{n-1}, \ldots, v_{j+1}, i, v_{j-1}, \ldots, v_1)$ by $\mathbf{v}|_j^i$. For $j \in \{1, 2, \ldots, n\}$, we refer to the neighbour $(v_n, \ldots, v_{j+1}, v_j + 1, v_{j-1}, \ldots, v_1)$ (resp. $(v_n, \ldots, v_{j+1}, v_j - 1, v_{j-1}, \ldots, v_1)$, $(v_n, \ldots, v_{j+1}, v_j + 1, v_{j-1} + 1, \ldots, v_1 + 1)$, $(v_n, \ldots, v_{j+1}, v_j - 1, v_{j-1} - 1, \ldots, v_1 - 1)$) as $\mathbf{v}_{(j,+1)}$ (resp. $\mathbf{v}_{(j,-1)}, \mathbf{v}_{(\leqslant j,+1)}, \mathbf{v}_{(\leqslant j,-1)}$). We can combine our notation as the following example shows: $\mathbf{v}_{(j,+1)}|_n^3$ denotes the vertex obtained by taking the vertex $\mathbf{v}_{(j,+1)}$ and fixing its nth component at 3 whilst leaving all other components as they were.