Hypercubes: Properties, Applications, Problems

Mark Ramras and Beth Donovan

May 31, 2007

1. Preliminaries (Beth)

2. Structure of $\mathrm{Aut}(Q_n)$

For $z \in Q_n$ define $\sigma_z : Q_n \longrightarrow Q_n$ by $\sigma_z(x) = z \bigoplus x$. $\Sigma_n = \{\sigma_z : z \in Q_n\} \simeq Z_2^n$. Σ_n = the subgroup of "complementations". For $\theta \in \mathcal{S}_n = \mathrm{Perm}\,\{1, 2, \ldots, n\}$, define $\rho_\theta(x_1, x_2, \ldots, x_n) = (x_{\theta(1)}, x_{\theta(2)}, \ldots, x_{\theta(n)})$. Abusing notation, call this subgroup of $\mathrm{Aut}(Q_n)$ $\mathcal{S}_n$.

Fact: Every $\phi \in \mathrm{Aut}(Q_n)$ can be uniquely expressed as $\sigma_z \circ \rho_\theta$. $\therefore |\,\mathrm{Aut}\,(Q_n)\,| = 2^n \cdot n!$.

In the subsets of $\{1, 2, \ldots, n\}$ description of Q_n, $\sigma_A(B) = A \Delta B$ and $\rho_\theta(B) = \theta(B)$.
Easy to see: $\boxed{\rho_\theta \circ \sigma_A \circ \rho_{\theta^{-1}} = \sigma_{\theta(A)}}$. $\therefore \Sigma_n$ is a *normal* subgroup of $\mathrm{Aut}(Q_n)$.
Note: If e is an edge with direction i, then direction$(\sigma_A(e)) = i$, and direction$(\rho_\theta(e)) = \theta(i)$.
Remark: Q_n is "vertex-transitive", i.e. for any $A, B \in V(Q_n), \exists \phi \in \mathrm{Aut}\,(Q_n)$ s.t. $\phi(A) = B$. Namely, take $\phi = \sigma_{B\Delta A}$.

3. Walks, paths, cycles and edge directions
Walk: a seq $x_1, x_2, \ldots, x_k + 1$ of adjacent vertices (not nec. distinct). *Length* of the walk = # edges = k .
Closed walk: $x_{k+1} = x_1$.
Path: walk with *no* repetition of vertices.

k- cycle: Closed walk $x_1, x_2, \ldots, x_k, x_1$ s.t. $x_1, \ldots, x_k$ is a path.

On Q_n any walk uniquely determined by initial vertex A and a seq. of edge directions $i_1, i_2, \ldots, i_k$.

Example: On Q_3, for any A the edge dir. seq. 1,2, 1,3,2,3 gives a 6-cycle.

Note: A walk in Q_n is closed iff every edge dir. in the seq occurs an *even* no. of times.
A walk $x_1, x_2, \ldots, x_k$ is a path iff in each subwalk $x_j, x_{j+1}, \ldots, x_{j+q}$ $\exists$ at least one edge dir. which occurs an *odd* no. of times.

4. Hamiltonian paths and cycles
Ham. path(cycle): path (cycle) containing all the vertices of the graph.

Proposition 1 *For $n \geq 2$, if x and y are any two vertices of opposite parity, $\exists$ a Ham. path from x to y.*

Proof. By induction on n, using $Q_{n+1} \simeq Q_n \times Q_1$.

Corollary 1 *For $n \geq 2, Q_n$ has Ham. cycles.*

Binary reflected Gray codes (B.R.G.C.)

For $n \geq 2$, BRGC(n) is the edge dir. seq. of a Ham. cycle for Q_n, defined inductively:
BRGC(2)= $1, 2, 1, 2$
BRGC(3)= $1, 2, 1, \mathbf{3}, 1, 2, 1, \mathbf{3}$.
Assume BRGC(n) has n as its last term. Let P(n) be the seq with the final n deleted. (P(n) is a palindrome.) Then
BRGC($n+1$) $= P(n)$,$n+1$,P(n),$n+1$=(P(n), $n+1)^2$.

Application to "Towers of Hanoi"

Rule 1: Odd numbered disks move 1 peg CW, even numbered disks move 1 peg CCW.

Rule 2: On i^{th} move, move the disk whose number is the i^{th} term of $P(n)$.

5. Hamming Codes

R.W. Hamming (AT&T Bell Labs, 1949): "Perfect single-error-correcting codes". Now known as Hamming codes. For $n = 2^k - 1$, constructed subgroup H_n of Q_n s.t. the closed neighborhoods of the elts of H_n (the "codewords") form a vertex partition of $V(Q_n)$. So every n-bit string is either a codeword or distance 1 from a unique codeword. $\therefore$ if a single error is made in transmitting a codeword, it can be corrected. Closed nbhds are pairwise disjoint iff minimum distance bet. codewords ≥ 3. Then their union is $V(Q_n)$ iff $|H_n|(n+1) = |V(Q_n)| = 2^n$. $\therefore n+1|2^n$, so $n = 2^k - 1$. $\therefore |H_n| = 2^{n-k}$.

Simple, elegant construction: $A = k \times n$ 0-1 matrix. Columns are the $n = 2^k - 1$ distinct non-0 binary k-tuples. Define $H_n = \text{Ker}(A) = \{\vec{x} \in Z_2^n : A \cdot \vec{x} = \vec{(0)}\}$. Rank$(A) = k$ since among the cols of A are the standard basis vectors $\vec{e_1}, \ldots, \vec{e_k}$. So dim(Ker $(A)) = n - k$, and $\therefore$ $|H_n| = 2^{n-k}$. If $\vec{x} \in H_n$ then $x_1 A_1 + \cdots + x_n A_n = \vec{0}$, where $A_j = j^{\text{th}}$ column of A. If wt$(\vec{x}) = 1$ then $\vec{x} = \vec{e_j} \Longrightarrow A_j = \vec{(0)}$ for some j. Contradiction.
If wt$(\vec{x}) = 2$, then $A_i + A_j = \vec{0}$ for some $i \neq j$. So $A_i = A_j$. Contradiction. $\therefore$ wt $(\vec{x}) \geq 3$.

6. Edge decompositions and Fundamental Sets of edges

In any graph, $\sum_{v \in V} \deg(v) = 2|E(G)|$. In Q_n, $\deg(v) = n, \forall v$. $\therefore 2(|E(Q_n)| = n \cdot 2^n$.
Def. H, G graphs. H *edge decomposes* G if $E(G) =$ disjoint union of $E(H_i), 1 \leq i \leq k$, where each $H_i \simeq H$.

Stronger notion: Let $E' \subset E(G)$. E' is a *fundamental set* for $E(G)$, with respect to group Γ, if Γ is a subgroup of Aut(G) and the subgraph $(V(E', E')$ edge decomposes G.

Proposition 2 [RAM1, Graphs and Combinatorics, 1991] *Let T be any tree with n edges. Let $\Gamma = \{\sigma_x \in \Sigma_n : \text{wt}(x) \text{ is even}\}$. Then there is a distance-preserving embedding ϕ of T into Q_n s.t. $E(\phi(T))$ is a fundamental subset of $E(Q_n)$, w.r.t. Γ.*

Idea of proof. Embed T in Q_n s.t. no 2 edges of $\phi(T)$ have same direction.

First: assign the labels $1, 2, \ldots, n$ in an arbitrary $1-1$ fashion to the n edges of T. Pick any $v_0 \in V(T)$ and define $\phi(v_0) = \vec{0}$. For any other $w \in V(T), \exists$ a unique $v_0 - w$ path in T. This gives a sequence of labels $i_1, \ldots, i_n$. Viewing this as a sequence of distinct edge directions in Q_n gives a path from $\vec{0}$ to a unique vertex $\phi(w)$.

$|\Sigma_n| = 1/2(2^n) = 2^{n-1}$. Must check $\sigma_x(E(T)) \cap \sigma_y(E(T)) = \emptyset$, if $x, y \in \Gamma, x \neq y$. Since Γ a group, suff. to check that $\sigma_x(E(T) \cap \sigma_{\vec{0}}(E(T) = \emptyset$ if $x \neq \vec{0}$. If $e \in \sigma_x(E) \cap E$, then $\text{dir}(\sigma_x(e) = \text{dir}\,(e)$. But no 2 different edges of $E(T)$ have the same direction. $\therefore$ $\sigma_x(e) = e$. Since $\text{wt}(x)$ is even, σ_x cannot transpose the endpoints of e. $\therefore$ it must leave each endpoint fixed. But $\sigma_x(v) = x \bigoplus v$, so $x \bigoplus v = v \Longrightarrow x = \vec{0}$. Contradiction.

7. Edge decompositions by paths. Double stars, cycles as fundamental sets. (Beth)

8. Vertex partitions: generalization of Hamming Codes.

Restatement of Hamming's thm:
For $n = 2^k - 1, k \geq 2, Q_n$ has a vertex partition into n-stars. (Centers of the stars are the codewords.) The centers form a subgroup Γ of $\Sigma_n \subset \text{Aut}\,(Q_n)$, and if T_0 is the n-star with center $\vec{0}$ then $\{\gamma(T_0) :\ \gamma \in \Gamma\}$ is a vertex partition of Q_n.

Generalization:

Theorem 1 [RAM2, J.Comb.Th. Ser. B 1992] *For $n = 2^k - 1, k \geq 2$, if T is* **any** *tree on $n+1$ vertices (and therefore n edges), $\exists$ a distance-preserving embedding φ of T in Q_n and a subgroup Γ of Aut(Q_n) s.t. $\{\gamma(\varphi(T)) :\ \gamma \in \Gamma\}$ is a vertex partition of Q_n.*

Idea of proof. The embedding φ is the same as for the earlier proposition. However, the subgroup Γ varies with T and is not constructed explicitly. The idea is to show that for any two trees T_1 and T_2, if Γ_i is the subgroup of $\text{Aut}(Q_n)$ of maximum cardinality s.t. the family $\{\gamma_i(T_i) :\ \gamma_i \in \Gamma_i\}$ is pairwise disjoint, $i = 1, 2$ then $|\Gamma_1| = |\Gamma_2|$. Letting T_1 be the n-star and $T_2 = T$, the result then follows from Hamming's Theorem.

9. Q_n as popular architecture for massively parallel computers.

Since the 1980's, the biggest spur to research on hypercubes is the fact that they are a popular architecture for parallel processing computers. For example, the Connection Machine, invented by Tom Leighton and his grad students from MIT's Comp. Sci. dept, who formed the company Thinking Machines, has 2^{16} (approx. 65000) processors linked together as the vertices of Q_{16}.

Advantages of Q_n.

1. Diameter $= n$; small compared to $|V(Q_n)| = 2^n$.
2. Scalable (For $m > n, Q_m$ is built from Q_n in a very natural way.)
3. Good "fault tolerance". System will remain connected if there are $\leq n-1$ faulty processors or $\leq n-1$ faulty links. (Vertex connectivity = edge connectivity $= n$.)

Def. Generalized Cutset: A set of vertices and edges whose deletion leaves a disconnected graph.

Proposition 3 [Ram3, Disc. Math. 2002] *The minimum size of a generalized cutset of Q_n is n. Any such cutset consists of vertices adjacent to, and edges incident on, a fixed vertex v_0.*

10. Problems

(a) Hamiltonian cycles in subgraphs of Q_n: "Middle level question". If $n = 2k+1$, does the subgraph of Q_n consisting of the vertices of weights k and $k+1$ have a Ham. cycle? (Known true for $1 \leq k \leq 15$.)

(b) Minimum size of a dominating set.
A set of vertices $S \subset V(G)$ *dominates* graph G if every $x \in V(G) \setminus S$ is adjacent to at least one $s \in S$. dom(G) is the minimum size of a dominating set for G.
Question: What is dom(Q_n)? Hamming's Thm $\Longrightarrow$: for $n = 2^k - 1$, dom$(Q_n) = 2^{n-k}$. So, for example, dom$(Q_3) = 2$ and dom$(Q_7) = 2^4 = 16$. For small values of n exact values are known, though for most others there are only upper and lower bounds. For example, it's not hard to see that dom$(Q_4) = 4$, and it can be proved that dom$(Q_5) = 7$. An example of Dan Ramras (my son) shows that dom$(Q_6) \leq 12$:

Example. The example given below shows that dom$(Q_6) \leq 12$. It also has an interesting algebraic structure. For readability, we will denote the set

$\{i_1, i_2, \ldots, i_k\}$ by $i_1 i_2 \ldots i_k$, e.g. $\{\{1\}, \{3\}, \{1, 2, 3\}\}$ will be written $1, 3, 123$.
$S = 1, 3, 13, 24, 25, 26, 456, 1456, 3456, 12345, 12356, 123456$.
Let $S_1 = 1, 3, 13, 24, 26, 12346$ and $S_2 = 25, 456, 3456, 1456, 12345, 123456$. Then $S_1 \cup S_2 = S, S_1 \cap S_2 = \emptyset$ and $S_1 \Delta \{1, 3, 4, 5, 6\} = S_2$. Note that half the vertices in S are even and half are odd.

(c) Maximal "balanced independent subsets"
Def. $S \subset V(G)$ is *independent* if no 2 vertices of S are adjacent.
Remark. A maximal indep. set in G must be a dominating set. But a minimal dominating set need not be indep. In fact the example showing $\text{dom}(Q_6) \leq 12$ is *not* indep.
Def. If $G = (X, Y)$ is bipartite, call a subset S of $V(G)$ *balanced* if $|S \cap X| = |S \cap Y|$.

So $S \subset V(Q_n)$ is balanced iff exactly half its elements have even weight.

Example. Hamming codes are balanced maximal indep. sets. In general, if $\mathcal{C}$ is a subgroup of Q_n, with $\min\{wt(x) : x \in \mathcal{C}\} \geq 2$, and $\text{wt}(z)$ is odd for some $z \in \mathcal{C}$, then $\mathcal{C}$ is a balanced indep. set.

Some balanced maximal independent sets in Q_n are the end result of a 2-player game between Even and Odd. Starting with Even, they alternate choosing vertices (even wt for E, odd wt for O) so that with each move, set of chosen vertices remains indep. The last player able to make a legal move wins.

O has an easy winning strategy: he chooses a fixed odd vertex z of wt≥ 3. When E pays even vertex x, O plays $x + z$. $\text{wt}(x + z) \equiv wt(x) + wt(z) \equiv 1 \pmod 2$.

For let $x_k =$ E's k^{th} move. If O's k^{th} move $x_k + z$ were illegal, $x_k + z$ would be adj. to some earlier chosen vertex, nec. $x_i, i < k$. Adding z to both, $x_k = (x_k + z) + z$ is adj. to $x_i + z$. But then E's move x_k was illegal. So O must win, and the resulting set of chosen vertices is a balanced maximal indep. set.

Even restricting to sets produced by this strategy, different odd z's can produce bal. max. indep. sets of different sizes.

Example. Let $n = 5$.
(1) $z_1 = \{1, 2, 3\} = 123$. $S_1 = \{\emptyset, 123, 14, 234, 15, 235, 1245, 345\}$

(2) $z_2 = 12345$. $S_2 = \{\emptyset, 12345, 12, 345, 13, 245, 14, 235, 15, 234\}$

Problem: What is the max size of a balanced independent subset? (Beth)

11. (If time permits) Communicating on the network Q_n: Broadcasting and Permutation Routing

(a) Broadcasting: Assume edges are bi-directional, and a message at one node can be sent simultaneously to any subset of its neighbors. Also assume that for each directed edge $\langle x, y\rangle$ at most one message can be sent from x to y at a given moment. One messge is sitting at one node (processor) v_0. Want to transmit it to all other nodes. Each node only needs to receive the message from 1 neighbor. So want a spanning tree T for Q_n (i.e. T contains all nodes of Q_n) of diameter n, so as to minimize total broadcast time. If there are several messages to be broadcast, look for as many edge-disjoint spanning trees as possible, for with k such trees, can send k messages simultaneously, provided each node can hold k messages at a single moment.

(b) Permutation Routing: Suppose each node x has a message that must be sent to node $\pi(x)$, where π is a permutation of $V(Q_n)$. Assume a global clock with which the 2^n processors are synchronized. As before, edges are bi-directional, and in one tick of the clock, 2 adj. nodes MAY exchange messages. ASSUME: at any given time, a node contains *at most one* (and therefore *exactly one*) message. Also, with each tick of the clock, a node may transmit its message to ONE of its neighbors, or else keep that message. So the movement of messages in 1 tick of the clock amounts to a permutation φ of the nodes, with the property that: $\forall x, d(x, \varphi(x)) \leq 1$.

Def. $k(\pi) = \text{Max}\{d(x, \pi(x)) : x \in Q_n\}$.
$\Delta = \{\text{perms } \varphi : k(\varphi) = 1\}$. This means that when φ is factored (algebraically) into a product of cycles, each of these cycles of length ≥ 2 corresponds to a directed cycle in Q_n in the graph sense. (Transpositions correspond to adj. nodes swapping messages along the 2 directed edges be-

tween them.)

A *congestion-free routing* of π amounts to a factorization: $\pi = \varphi_t \varphi_{t-1} \cdots \varphi_1$ where each $\varphi \in \Delta$, i.e. $\pi \in \Delta^t$.

Fact: Every π has such a factorization.

Def. $t_\Delta(\pi) = \min\{t : \ \pi \in \Delta^t\}$.

Clearly: $t_\Delta(\pi) \geq k(\pi)$.

For one class of permutations, routings arise in a nice way.

Linear Permutations: Let M be an invertible $n \times n$ matrix over Z_2. Define $\pi_M : Q_n \Longrightarrow Q_n$ by $\pi_M(\vec{x}) = M \cdot \vec{x}$.

Proposition 4 [RAM, SIAM J. Disc. Math., 1998] *If M has an LU decomposition, then from this we can obtain an n-step congestion-free routing of π_M.*