Finite and Infinite Hypercubes as Direct Products

Wilfried Imrich
Montanuniversität Leoben, A-8700 Leoben, Austria
imrich@unileoben.ac.at

Douglas Rall
Furman University, Greenville, SC 29613
drall@herky.furman.edu

Abstract

We characterize the factorizations of finite or infinite hypercubes with respect to the direct product in the class of unoriented simple graphs with loops. The paper extends the corresponding result for finite graphs [1]. It is based on a new approach that yields a simple, unified proof for both the finite and infinite case of the main lemma in [1].

Key words: direct product, hypercube, automorphism, involution

AMS subject classification (2000): 05C70, 05C25

1 Introduction

This paper is concerned with the decomposition of finite and infinite hypercubes with respect to the direct product. It generalizes the results of [1] to infinite graphs. It is shown that every decomposition of the k-dimensional hypercube Q_k with respect to the direct product, where k is a finite or infinite cardinal, is of the form $K_2 \times G$ where K_2 is the complete graph on two vertices and G a non-bipartite graph that contains a spanning subgraph S that is isomorphic to Q_{k-1}. Furthermore, G is characterized by the condition that the endpoints of the edges in $E(G) \setminus E(S)$ have even distance in S and induce an involution of S.

The main result in [1] is the existence of a spanning hypercube in every graph G with the property that $K_2 \times G$ is a hypercube. It is effected by an application of Graham's density lemma for (the number of edges of) subgraphs of the hypercube, a method that does not extend to the infinite case. Here Graham's density lemma is replaced by application of metric properties of the sets of parallel edges in hypercubes and symmetries induced by these sets.

The fact that every factorization has exactly two factors, where one of them is always a K_2, depends on properties of the so-called Cartesian skeleton of hypercubes that was originally introduced for the investigation of finite non-bipartite graphs. It also exists for infinite hypercubes. Just as in the finite case it can be used to show that there are no factorizations other than those just described.

The paper is organized as follows. The next section contains the main definitions and basic facts about the direct product. Then comes the proof of the main result, that is Lemma 2. Formally this lemma is identical to the Lemma 2 in [1], except that the new proof is also valid for the infinite case as will be displayed in the the section on infinite hypercubes. For the reader not interested in infinite graphs the main difference to the proof in [1] is the fact that it is shorter and that it does not refer to Graham's density lemma.

The proofs of the consequences of Lemma 2 for finite graphs, that is the characterization of all factorizations of the hypercube with respect to the direct product, are identical to those in [1] and therefore omitted.

Then follows the definition of the Cartesian product of infinitely many graphs, of the weak Cartesian product and the infinite hypercube. It will then be clear that Lemma 2 also holds in the infinite case.

For the final, complete characterization of all factorizations of the finite or infinite hypercube the Cartesian skeleton is invoked again. As it was introduced for finite graphs the last section is devoted to show that the infinite hypercube also has a Cartesian skeleton with the properties needed for our purposes.

2 Preliminaries

All graphs considered here are undirected graphs that may contain loops but not multiple edges.

The *direct product* $G \times H$ of two graphs G and H is defined on the Cartesian product $V(G) \times V(H)$ of the vertex sets of the factors. Its edge set is the set of all pairs of vertices $(a, x), (b, y) \in V(G) \times V(H)$ where $ab \in E(G)$ and $xy \in E(H)$. It is commutative, associative and the one-vertex graph with a loop is a unit.

Figure 1 depicts the direct products $K_2 \times P_3^*$, where P_3^* is a path of length 2 with loops added to its endpoints, and $K_2 \times K_3$. In both cases the product is a cycle of length 6.

The *Cartesian product* $G \Box H$ has the same vertex set as the direct product. Its edge set consists of all pairs (a, x), (b, y) with $ab \in E(G)$ and $x = y$, or $a = b$ and $xy \in E(H)$. It is also commutative and associative with K_1 as a unit.

The subgraph of $G \Box H$ induced by the vertices $(a, x), x \in V(H)$, is called an

Figure 1: Two decompositions of C_6

H*-layer* of $G \Box H$ and denoted by $H^{(a,x)}$. Note that any H-layer is isomorphic to H. Analogously one defines G-layers. The d-dimensional *hypercube* or d*-cube* Q_d is the Cartesian product of d copies of the complete graph K_2 on two vertices. So $Q_1 = K_1$ and we also set $Q_0 = K_1$. The vertices of Q_d can be considered to be all binary vectors of length d. Two such vertices $x = (x_1, x_2, \dots, x_d)$ and $y = (y_1, y_2, \dots, y_d)$ are adjacent in Q_d if and only if there exists an index i such that $x_i = 1 - y_i$ and $x_j = y_j$ for $j \neq i$.

Let $Q_{d-1} \Box K_2$ be an arbitrary factorization of Q_d. The edges between the two Q_{d-1}-layers are said to be of the same *color* or are *parallel* in Q_d. The set of all edges between two Q_{d-1}- layers will be referred to as a *color class* or a *parallel class* of the edge set of the factorization. Such classes are also equivalence classes with respect to the transitive closure of the Djoković-Winkler relation Θ as defined in [2, p. 48]. We will therefore denote such classes containing the edge e by Θ_e. The main property (cf. [2, Lemma 2.3]) of this relation that we shall use is the fact that in a bipartite graph two edges $e = [u, v]$, $f = [x, y]$ are in the relation Θ if and only if the notation can be chosen such that

$$d(u, x) = d(v, y) = d(u, y) - 1 = d(v, x) - 1.$$

Layers of direct products are defined analogously to those of the Cartesian product. In the case of the direct product the layer $H^{(a,x)}$ is isomorphic to H only if a carries a loop (in G), otherwise the edge-set of $H^{(a,x)}$ is empty, cf. Figure 1.

For a given graph G it follows directly from the definition of the direct product that interchanging vertices between the G-layers is an automorphism of $K_2 \times G$. In fact, even more can be said.

Lemma 1 *Suppose that H is a bipartite graph. There exists a graph G such that H is isomorphic to $K_2 \times G$ if and only if H has an involution, that is an automorphism of order two, that interchanges the color classes of $V(H)$.*

Proof. If H can be factored as $K_2 \times G$ for some graph G, then the involution φ defined by $\varphi(i, x) = (1 - i, x)$ for $i = 0, 1$ is such an automorphism of H. Conversely, suppose H has color classes $C_0 = \{u_1, u_2, \ldots, u_n\}$ and $C_1 = \{v_1, v_2, \ldots, v_n\}$ together with an automorphism φ such that $\varphi(u_i) = v_i$ and $\varphi(v_i) = u_i$ for each i. We construct a graph G on the vertex set $\{w_1, w_2, \ldots, w_n\}$ by including the edge w_iw_j in G if and only if $u_iv_j \in E(H)$. It follows immediately from the definition of the direct product and the assumptions about φ that H is isomorphic to $K_2 \times G$. $\square$

We note that if the involution φ in the above lemma has the property that $\varphi(x)$ is not adjacent to x for any vertex x of H, then the resulting factor G has no loops. However, each pair of vertices x and $\varphi(x)$ that are adjacent in H give rise to a loop in G. The former case is illustrated in the second factorization of C_6 and the latter by the first factorization in Figure 1.

3 The main result

For a bipartite graph G with bipartition $V(G) = X + Y$ we call an involution α *bipartite* if $\alpha(X) = X$. For such an involution we let G^α denote the graph obtained from G by addition of the edges $\{uv \mid u = \alpha(v), v \in V(G)\}$.

It is not hard to see that $K_2 \times H^\alpha_{k-1}$ is isomorphic to H_k for any bipartite involution. We wish to show that every graph G with $H_k \cong K_2 \times G$ is of that form. The main step in the proof is the following lemma.

Lemma 2 *If G is a connected, nonbipartite graph such that $K_2 \times G$ is a k-dimensional hypercube, then G has a spanning subgraph isomorphic to a $(k-1)$-dimensional hypercube.*

Proof. Assume that G is connected and nonbipartite such that $H = K_2 \times G$ is isomorphic to Q_k. We denote the vertex set of K_2 by $\{0, 1\}$ and the projection map from H onto G by p_G. The idea of the proof is to find a parallel class Θ of edges in the hypercube Q_k such that for each $e \in \Theta$ either $p_G(e)$ is a loop or there is another edge $f \in \Theta$ with $p_G(e) = p_G(f)$.

Choose a odd cycle $C = v_1, v_2, \ldots, v_{2\ell+1}, v_1$ of shortest length in G. If C is a loop, then there is an edge e in H projecting onto it, so assume that $\ell \geq 1$. Consider the subgraph $K_2 \times C \simeq C_{4\ell+2}$ of H. The edges $e = [a, d]$ and $f = [b, c]$ belong to $K_2 \times C$ where $a = (0, v_1), b = (0, v_{2\ell+1}), c = (1, v_1)$ and $d = (1, v_{2\ell+1})$. Denote by Θ_e the parallel class of H containing the edge e.

Suppose $d_H(a, b) < 2\ell$. Since a and b belong to the same color class of H, $d_H(a, b) = 2r < 2\ell$. If P is a shortest path in H from a to b, then P projects under the homomorphism p_G to a walk Q of length $2r$ joining v_1 and $v_{2\ell+1}$. If P does not contain a pair $(0, x)$ and $(1, x)$ for some vertex x of G, then adding the edge $v_1v_{2\ell+1}$ to Q yields an odd cycle of length $2r + 1$. Hence, there exist vertices $(0, v)$ and $(1, v)$ in P. From among all such pairs choose $(0, w)$ and $(1, w)$ that are closest together along P. The segment of P from $(0,w)$ to $(1,w)$ projects onto an odd cycle in G whose length is less than $2r$. This contradicts the choice of C, and so $d_H(a,b) = 2\ell = d_H(c,d)$. It follows that $f \in \Theta_e$.

Let R be a shortest path in H from a to c. Since a and c are from different color classes of H, the length of R is odd. The image $p_G(R)$ is a closed walk beginning and ending at v_1 that must contain an odd cycle. For if not, then its edges induce a bipartite subgraph of G that has a closed walk of odd length. Therefore, the length of $p_G(R)$, and hence also of R, is at least $2\ell+1$. We conclude that $d_H(c,d) < d_H(c,a)$, and so c and d belong to the same component of $H \setminus \Theta_e$.

Let $g = [x,y]$ be an arbitrary edge of H where $x = (0,u)$ and $y = (1,v)$, and suppose that $g \in \Theta_e$. The involution φ that interchanges vertices $(0,t)$ and $(1,t)$ for all $t \in V(G)$ leaves the parallel classes of H invariant and $\varphi(e) = f$. Thus, $\varphi(g) = [(1,u),(0,v)]$ is also in Θ_e. That is, for all $u, v \in V(G)$, $[(0,u),(1,v)] \in \Theta_e$ implies that $[(1,u),(0,v)] \in \Theta_e$.

The graph $H \setminus \Theta_e$ consists of two components, say S_1 and S_2 each of which is a hypercube of dimension $k-1$. To complete the proof we will now show that p_G is injective when restricted to S_1 or to S_2. We assume without loss of generality that $a, b \in S_1$ and that $c, d \in S_2$. Consider the following partition $V(H) = A \cup B \cup C \cup D$ depicted in Figure 2 where

$$A = \bigcup_{(0,x)\in S_1,\,(1,x)\in S_1} \{(0,x),(1,x)\} \qquad B = \bigcup_{(0,x)\in S_2\;(1,x)\in S_1} \{(0,x),(1,x)\}$$

$$C = \bigcup_{(0,x)\in S_1\;(1,x)\in S_2} \{(0,x),(1,x)\} \qquad D = \bigcup_{(0,x)\in S_2\;(1,x)\in S_2} \{(0,x),(1,x)\}.$$

Since $a \in S_1$ and $c \in S_2$ and $p_G(a) = p_G(c)$, it follows that $a, c \in C$ and hence $C \neq \emptyset$. Suppose $A \neq \emptyset$. Let $u = (0,x) \in A$. By definition of A, $v = (1,x) \in A$ and both of u and v belong to S_1. Each vertex of S_1 and each vertex of S_2 is incident with exactly one edge that belongs to Θ_e. Let $w, z \in S_2$ such that $[u,w] \in \Theta_e$ and $[v,z] \in \Theta_e$. Since H is bipartite it follows that $w \in \{1\} \times G$ and that $z \in \{0\} \times G$, and therefore $w, z \in D$. The hypercube S_1 is connected so there must be an edge h in S_1 from A to $B \cap S_1$ or $C \cap S_1$. Without loss of generality we can assume that h connects A with $C \cap S_1$ and that it is of the form $h = [(1,t),(0,s)]$, where $(1,t) \in A \subset S_1$ and $(0,s) \in S_1 \cap C$.

This implies that s and t are adjacent in G, and hence $g = [(1,s),(0,t)]$ is an edge of H. But $(0,t) \in S_1$ and $(1,s) \in S_2$. Hence g is the unique edge of H incident with $(0,t)$ that belongs to the parallel class Θ_e. By the above argument it follows that h must also belong to Θ_e, but this is a contradiction since h is incident with two vertices of S_1.

Figure 2: The partition $V(H) = A \cup B \cup C \cup D$

Therefore, $A = \emptyset$, and we conclude that p_G is injective on S_1. That is, G has a spanning hypercube. $\square$

Corollary 3 *The hypercube Q_k is representable as a product of the form $G \times K_2$ if and only if G is isomorphic to Q_{k-1}^{α} for some bipartite involution α of Q_{k-1}.*

Proof. Suppose $G \times K_2 = H_k$. Then the projections of endpoints of the edges in Θ_e induce α. In other words, let $[u, v] \in \Theta$. Then $\alpha(p_G u) = v$.

The converse is just as easy. $\square$

In order to show that H_k has no other decompositions with respect to the direct product the Cartesian skeleton as defined in [2] is invoked in [1]. For a more detailed, self-contained discussion that also holds for infinite hypercubes see Section 5.

4 Infinite hypercubes

We have introduced the hypercube as the Cartesian product of finitely many complete graphs on two vertices. So we begin with the generalization of the Cartesian product of finitely many factors to that of infinitely ones. Clearly the vertex set of $G_1 \Box \cdots \Box G_n$ consists of the coordinate vectors $v = (v_1, v_2, \ldots, v_n)$ where $v_i \in V(G_i)$, and the edge set of $G_1 \Box \cdots \Box G_n$ is the set of all unordered pairs $[u, v] = [(u_1, u_2, \ldots, u_n), ((v_1, v_2, \ldots, v_n)]$ for which there exists a $k \in \{1, 2, \ldots n\}$ such that $[u_k, v_k] \in E(G_k)$ and $u_i = v_i$ for $i \neq k, i \in \{1, 2, \ldots n\}$.

For infinitely many factors, we replace the coordinate vector by a function from an index set into the sets of vertices of the factors. Thus, let I be an index set and G_ι, $\iota \in I$, be a family of graphs. Then the Cartesian product

$$G = \prod_{\iota \in I} G_\iota$$

is defined on the set x of all functions $x : \iota \mapsto x_\iota$, $x_\iota \in V(G_\iota)$, where two vertices x, y are adjacent if there exist a $\kappa \in I$ such that $[x_\kappa, y_\kappa] \in E(G_\kappa)$ and $x_\iota = y_\iota$ for $\iota \in I \setminus \kappa$. We call the x_ι the coordinates of x.

For finite I this clearly coincides with the usual definition, and in this case the product is connected if and only if the factors are. However, if we have infinitely many nontrivial factors, there are vertices that differ in infinitely many coordinates. They cannot be connected by paths of finite length, since the endpoints of every edge differ in just one coordinate, so the product is disconnected.

The connected components of G are called *weak Cartesian products*. To identify a component, it suffices to specify an arbitrary vertex of it. Thus the weak Cartesian product $G = \prod_{\iota \in I}^{a} G_\iota$ is the connected component of $\prod_{\iota \in I} G_\iota$ containing a.

It is easy to see that the components of the Cartesian product of infinitely many factors K_2 are pairwise isomorphic. We can thus define the $\mathfrak{n}$-dimensional hypercube $Q_{\mathfrak{n}}$ as the weak Cartesian product $\prod_{\iota \in I}^{a} G_\iota$, where all G_ι are isomorphic to K_2, $a_\iota = 0$ for all $\iota \in I$, and $|I| = \mathfrak{n}$.

One can show that every connected graph is uniquely representable as a weak Cartesian product of connected prime graphs. This result is due to Imrich and Miller, for a proof cf. [2]. For us this implies that there are no other representations of $Q_{\mathfrak{n}}$ as a weak Cartesian product.

More importantly, we note that $K_2 \Box Q_{\mathfrak{n}} = Q_{\mathfrak{n}}$ and that we can define parallel edges as in the finite case. Removal of a set of parallel edges from $Q_{\mathfrak{n}}$ yields two isomorphic copies of $Q_{\mathfrak{n}}$ and the set of parallel edges induces an isomorphism between these copies in a natural way, just as in the case of finite graphs. Moreover, the characterization of parallel edges by the distances between their endpoint is the same as in the finite case.

Therefore, both Lemma 2 and Corollary 3 also hold for the infinite hypercube.

5 The Cartesian skeleton for finite and infinite hypercubes

Since the Cartesian product is well understood and well behaved it is tempting to apply results about the Cartesian product to other products, in our case to the direct one. The first main obstacle is that the subgraphs induced by the layers may be empty, that is, may contain no edges. So the question arises how to characterize - in given direct product - pairs of vertices that have the same projection into exactly one of the factors.

We begin with a product $G_1 \times G_2$ and look for pairs of vertices $\{u, v\}$ that have the same projection into one of the factors, say G. What characterizes them? We observe that the neighborhood $N(u)$ of $u = (u_1, u_2) \in G_1 \times G_2$ is the Cartesian product of the neighborhoods $N_{G_1}(u_1) \times N_{G_2}(u_2)$ of the neighborhood $N_{G_1}(u_1)$ of u_1 in G_1 by $N_{G_2}(u_2)$ of u_2 in G_2.

Clearly $N(v) = N_{G_1}(v_1) \times N_{G_2}(v_2)$ and

$$N(u) \cap N(v) = (N_{G_1}(u_1) \cap N_{G_1}(v_1)) \cap (N_{G_2}(u_2) \cap N_{G_2}(v_2)).$$

(Of course this generalizes to any number - finite or infinite - of factors.)

In a hypercube Q - be it finite or infinite - the neighborhoods of two different vertices are either disjoint or have exactly two vertices in common. Thus, in our case, $N(u) \cap N(v)$ is empty or has exactly two elements if u and v are distinct. Also, $N(u) \cap N(v) = N(u) \cap N(w) \neq \emptyset$ and $u \neq v$ implies that $v = w$, otherwise Q would contain a $K_{2,3}$.

We use this to show that the projections of u and v into (exactly) one factor are different and into all the others the same. Since the number of elements of $N(u) \cap N(v)$ is independent of the factorization this must hold for any factorization.

This needs more detail!!!!!!

In any decomposition of the hypercube with respect to the direct product two vertices of distance two have all projection, with the exception of exactly one, in common.

We say any two vertices u, v in $Q_{\mathfrak{n}}$ that have distance two are a *Cartesian pair* and call the graph H on $V(Q_{\mathfrak{n}})$ whose edges are the Cartesian pairs of $Q_{\mathfrak{n}}$ the *Cartesian skeleton*. The name comes from the fact that any decomposition of $Q_{\mathfrak{n}}$ into a direct product $G_1 \times G_2$ gives rise to a factorization $H_1 \Box H_2$ of H where the H_i-layers of H are identical with the G_i-layers; $i \in \{1, 2\}$. To see this consider the vertices $u = (u_1, u_2)$, $v = (u_1, v_2)$, $y = (y_1, u_2)$,and $z = (y_1, v_2)$, where the coordinates are taken with respect to the decomposition $G_1 \times G_2$. It suffices to show that $\{y, z\}$ is a Cartesian pair if $\{u, v\}$ is one. In other words, we have to show that y and z have distance two if this is the case of u and v. Since y and z have the same projections into G_1 this is only possible if their projections into G_2 have intersecting neighborhoods. Since these projections are u_2 and v_2, that is the same as those of u and v, $d_{G_1 \times G_2}(u, v) = 2$ exactly if $d_{G_1 \times G_2}(y, z) = 2$.

This decomposition of H is inherited by the decomposition of $Q_{\mathfrak{n}}$. If H is disconnected, then its connected components are the Cartesian products of the connected components of the factors. Since $Q_{\mathfrak{n}}$ is bipartite and since adjacent vertices in H have distance two in $Q_{\mathfrak{n}}$, it is clear that H gas exactly two connected components. Hence one of the Cartesian factors of $H = H_1 \Box H_2$ has two components and the other one is connected. Suppose H_1 has two connected components, say $H_1 = A + B$. Then $H = A \Box H_2 + B \Box H_2$.

The connected components of H are isomorphic, because the hypercube is vertex transitive. One calls them *halved cubes*. Halved cubes are prime with respect to the Cartesian product. To see this, we first observe that incident edges in a Cartesian product are always contained in exactly on square without diagonals; this follows directly from the definition. Thus, edges in a triangle are always in one and the same layer. Now, it is not hard to see that any two incident edges $[a, b]$, $[a, c]$ of a halved cube are either contained in a triangle or there is a third edge $[a, d]$ such that abd and adc are triangles.

But then both A and B consist of just one vertex each, say v_0 and v_1. This is the vertex set of G_1. Clearly, $[v_0, v_1] \in E(G_1)$, otherwise $G_1 \times G_2$ would have to be disconnected. Also, neither v_0 nor v_1 can carry a loop. For, suppose otherwise, assume that v_0 carries a loop, consider an edge $[x, y]$ of G and its projection into G_2. If the projection is a loop, then $G_1 \times G_2 = Q_{\mathfrak{n}}$ would have a loop, which is not possible. So the projection is an edge $[x_2, y_2]$ with $x_2 \neq y_2$. But then $[(v_0, x_2), (v_0, y_2)]$ is an edge with both endpoints in the halved cube H_2, which is not possible. Thus $G_1 = K_2$. Furthermore, any factorization of G_2 with respect to the direct product would entail a factorization of the halved cube H_2 with respect to the Cartesian product, which is not possible, since it is prime.

We have thus shown

Theorem 4 *Every decomposition of a nontrivial hypercube Q into a direct product has exactly two factors. On factor is always K_2 and the other one any of the graphs Q^{α}_{k-1} for a bipartite involution α of Q_{k-1} if Q has finite dimension k, or $Q^{\alpha}_{\mathfrak{n}}$ for a bipartite involution α of $Q_{\mathfrak{n}}$ if Q has infinite dimension $\mathfrak{n}$.*

References

[1] B. Brešar,W. Imrich, S. Klavžar and B. Zmazek, Hypercubes as direct products, SIAM J. Discrete Math., to appear.

[2] W. Imrich and S. Klavžar, *Product Graphs*, Wiley-Interscience, New York, NY, 2000.