Edge-transitive products[*]

Richard Hammack [a]   Wilfried Imrich [b]   Sandi Klavžar [c,d,e]

[a] Virginia Commonwealth University, Richmond, Virginia USA
rhammack@vcu.edu

[b] Montanuniversität Leoben, Leoben, Austria
wilfried.imrich@unileoben.ac.at

[c] Faculty of Mathematics and Physics, University of Ljubljana, Slovenia
sandi.klavzar@fmf.uni-lj.si

[d] Faculty of Natural Sciences and Mathematics, University of Maribor, Slovenia

[e] Institute of Mathematics, Physics and Mechanics, Ljubljana, Slovenia

Abstract

This paper concerns finite, edge-transitive direct and strong products, as well as infinite weak Cartesian products. We prove that the direct product of two connected, non-bipartite graphs is edge-transitive if and only if both factors are edge-transitive and at least one is arc-transitive, or one factor is edge-transitive and the other is a complete graph with loops at each vertex. Also, a strong product is edge-transitive if and only if all factors are complete graphs. In addition, a connected, infinite non-trivial Cartesian product graph G is edge-transitive if and only if it is vertex-transitive and if G is a finite weak Cartesian power of a connected, edge- and vertex-transitive graph H, or if G is the weak Cartesian power of a connected, bipartite, edge-transitive graph H that is not vertex-transitive.

1   Introduction

A graph G is *vertex-transitive* (resp. *edge-transitive*) if its automorphism group $\mathrm{Aut}(G)$ acts transitively on the vertex set $V(G)$ (resp. on the edge set $E(G)$). The vertex-transitivity of the four standard associative products of graphs (the Cartesian product, the direct product, the strong product, and the lexicographic product) is well-understood. If $*$ denotes any of the four standard products and G and H are arbitrary finite graphs, then $G * H$ has transitive automorphism group if and only if G and H

---

[*]Dedicated to Chris Godsil on the occasion of his 65th birthday

have transitive automorphism groups, see [6, Theorems 6.16, 7.19, 8.19, 10.14]. These classical results were proved about half a century ago, cf. [2, 8, 16, 17], hence it is surprising that not much can be found in the literature about the edge-transitivity of graph products, especially because edge-transitivity has been an active research topic since then, cf. the recent papers [4, 11, 13] and references therein.

The very recent article [10] characterizes finite edge-transitive Cartesian and lexicographic products. (That paper was in part motivated by a false, but related result from [12].) In this recent article it is shown that a Cartesian product is edge-transitive if and only if it is a Cartesian power of some vertex- and edge-transitive graph. For the lexicographic product it proves, if G is connected and not complete, then $G \circ H$ is edge-transitive if and only if G is edge-transitive and H is edgeless. Moreover, if G is complete, and $G \circ H$ is edge-transitive, then there is a complete graph K and an edgeless graph D for which $G \circ H = K \circ D$.

Here we continue these investigations. In Section 2 we consider and characterize edge-transitive, finite, non-bipartite direct products, while in Section 3 finite, edge-transitive strong products are characterized. In this way we round off the edge-transitivity of the standard graph products with the sole exception of bipartite direct products. For this remaining case we provide a sufficient condition for the direct product to be edge-transitive (Proposition 2.5) and conjecture that it is also necessary.

In the second part of the paper we turn to infinite graphs, more precisely to infinite (weak) Cartesian product graphs. As already mentioned, finite Cartesian products are vertex-transitive if and only if every factor is vertex-transitive. For the weak Cartesian product the situation is more complex. If all factors are vertex-transitive, then any of their weak Cartesian products is also vertex-transitive. However, a connected weak Cartesian product can be vertex-transitive, even when all factors are asymmetric. This was first observed in [9]. In Section 4 we also round off this story by characterizing vertex-transitive weak Cartesian products, and then characterize edge-transitive weak Cartesian products.

In the rest of this section we define the standard graph products and discuss several other concepts. The *Cartesian product* $G \square H$, the *direct product* $G \times H$, the *strong product* $G \boxtimes H$, and the *lexicographic product* $G \circ H$, each has as its vertex set the Cartesian product $V(G) \times V(H)$. Edges are as follows:

$$
\begin{aligned}
E(G \square H) &= \big\{(x,u)(y,v) \mid xy \in E(G) \text{ and } u = v, \text{ or, } x = y \text{ and } uv \in E(H)\big\}, \\
E(G \times H) &= \big\{(x,u)(y,v) \mid xy \in E(G) \text{ and } uv \in E(H)\big\}, \\
E(G \boxtimes H) &= E(G \square H) \cup E(G \times H), \\
E(G \circ H) &= \big\{(x,u)(y,v) \mid xy \in E(G), \text{ or, } x = y \text{ and } uv \in E(H)\big\}.
\end{aligned}
$$

The notions of *arc-transitive* and *half-transitive* graphs will arise several times. A graph is *arc-transitive* if for any two edges xy and uv there is an automorphism mapping $x \mapsto u$ and $y \mapsto v$, and another mapping $x \mapsto v$ and $y \mapsto u$. (In other words, the automorphism group acts transitively on the graph's arcs.) Clearly, arc-transitive graphs are both vertex-transitive and edge-transitive. A graph is *half-transitive* if it is vertex-transitive and edge-transitive, but not arc-transitive. For such a graph, the automorphisms mapping an edge xy to uv are either always $x \mapsto u$ and $y \mapsto v$, or

always $x \mapsto v$ and $y \mapsto u$. See [1] for further properties and examples of half-transitive graphs.

If a graph G is edge-transitive, then it need not be vertex-transitive. In that case it must be bipartite, as proved in [18, Proposition 2.2] under certain finiteness conditions. See also Lemma 3.2.1 in the classical book of Godsil and Royle [5] which is a fine source for general properties of vertex-, edge-, and arc-transitive graphs. In general we have the following.

Lemma 1.1 *A graph with no isolated vertices that is edge-transitive but not vertex-transitive is bipartite.*

Proof. Such a graph G must be non-trivial and every vertex must be incident to an edge. For a vertex x, let $\mathrm{Aut}(G)(x)$ denote the $\mathrm{Aut}(G)$-orbit of x. Take an edge xy and note that $\mathrm{Aut}(G)(x) \cup \mathrm{Aut}(G)(y) = V(G)$. If $\mathrm{Aut}(G)(x) \cap \mathrm{Aut}(G)(y)$ is non-empty, then G is vertex-transitive. Hence $\mathrm{Aut}(G)(x)$ and $\mathrm{Aut}(G)(y)$ form a partition of $V(G)$. Furthermore, by edge-transitivity, any edge of G joins $\mathrm{Aut}(G)(x)$ to $\mathrm{Aut}(G)(y)$. □

Note that a connected bipartite graph that is edge-transitive is also arc-transitive if and only if it admits an automorphism that reverses the bipartition of one of its components.

Consequently, a connected edge-transitive graph can be arc-transitive, or half-transitive, or neither. If it is neither, then it is bipartite and the partite sets are distinguishable in the sense that there is no automorphism mapping one to the other.

2 Finite direct products

This section's goal is to characterize edge-transitivity of direct products. Our main result is Theorem 2.3, below, stating that any connected non-bipartite direct product is edge-transitive if and only if both factors are edge-transitive and one is not half-transitive, or one factor is edge-transitive and the other is a complete graph with a loop at each vertex. We prepare for this by recalling some standard notions and results.

Denote by K_n^* the complete graph on n vertices with loops at each vertex and its complement by $\overline{K_n^*}$. Notice that $\overline{K_n^*}$ is completely disconnected, that is, it has n vertices and no edges.

Discussions of the direct product are often simplified with a certain equivalence relation R on the vertex set of a graph. Two vertices x and y of a graph are said to be in relation R if their open neighborhoods are identical, that is, if $N(x) = N(y)$. Clearly an R-equivalence class (an R-*class*) of vertices in a graph G induces either a $\overline{K_n^*}$ or a K_n^*. Also, for any two R-classes X and Y, either every vertex of X is adjacent to every vertex of Y, or no vertex of X is adjacent to any vertex of Y. For details, see Section 8.2 of [6], where it is also proved that any R-class of the direct product of two graphs is the Cartesian product of R-classes of the factors.

A connected edge-transitive graph G with more than one vertex cannot have any loops, because no automorphism can move a loop to a non-loop edge; therefore all

R-classes induce completely disconnected subgraphs. Consider an edge xy of an edge-transitive graph G, and say x (respectively y) belongs to the R-class X (respectively Y). Because any automorphism α of G sends R-classes to R-classes, the R-class containing $\alpha(x)$ has $|X|$ elements, and the R-class containing $\alpha(y)$ has $|Y|$ elements. By edge-transitivity, any R-class of G has size either $|X|$ or $|Y|$, and any edge joins an R-class of size $|X|$ to one of size $|Y|$. In particular, this means that if G has a odd cycle (or if it is arc-transitive) then all its R-classes have the same size. If G is bipartite, then any two R-classes in the same partite set have the same size.

Given a graph G, the quotient G/R is the graph whose vertices are the R-classes of G, with two classes being adjacent precisely if there is an edge in G between them. Any automorphism $\alpha : G \to G$ induces an automorphism $G/R \to G/R$ sending any R-class X to $\alpha(X)$. Conversely, if all R-classes of G have the same size, then we can lift any automorphism $\alpha : G/R \to G/R$ to an automorphism of G by simply declaring that each R-class X is mapped to $\alpha(X)$ via an arbitrary bijection. Moreover, if xy and $x'y'$ are two edges of G joining R-classes X and Y, then the transposition of x with x', and y with y' is an automorphism of G interchanging xy with $x'y'$. This implies a lemma.

Lemma 2.1 *If a graph G is edge-transitive, then G/R is edge-transitive. Conversely, if G/R is edge-transitive and non-trivial, and all R-classes of G have the same size, then G is edge-transitive.*

We call a graph R-thin if each R-class contains exactly one vertex. Clearly G/R is always R-thin. We remark also that $(A \times B)/R \cong A/R \times B/R$ (Section 8.2 of [6]).

Note that $G \times K_1^* \cong G$, so K_1^* is a unit for the direct product. A non-trivial graph is *prime* (over $\times$) if it cannot be expressed as a direct product of two graphs, neither of which is K_1^*. It is well-known that in the class of graphs where loops are allowed, connected non-bipartite graphs factor uniquely over the direct product into prime graphs. Moreover in the class of R-thin graphs, automorphisms have a particularly rigid structure.

Theorem 2.2 [6, Theorem 8.18] *Suppose φ is an automorphism of a connected nonbipartite R-thin graph G that has a prime factorization $G = G_1 \times G_2 \times \cdots \times G_k$. Then there exists a permutation π of $\{1, 2, \ldots, k\}$, together with isomorphisms $\varphi_i : G_{\pi(i)} \to G_i$, such that $\varphi(x_1, x_2, \ldots, x_k) = (\varphi_1(x_{\pi(1)}), \varphi_2(x_{\pi(2)}), \ldots, \varphi_k(x_{\pi(k)}))$.*

Now our main theorem's foundation is laid.

Theorem 2.3 *Suppose $A \times B$ is connected and non-bipartite. Then it is edge-transitive if and only if both factors are edge-transitive and at least one is arc-transitive, or one factor is edge-transitive (and non-trivial) and the other is a K_n^*.*

Proof. Note that A and B are connected and non-bipartite because their product is.

Say A and B are edge-transitive and A is arc-transitive. Take edges $(a_1, b_1)(a_1', b_1')$ and $(a_2, b_2)(a_2', b_2')$ of $A \times B$. We construct an automorphism of $A \times B$ carrying the first edge to the second. Begin with an automorphism β of B sending edge $b_1 b_1'$ to $b_2 b_2'$.

4

Case 1. Suppose $\beta(b_1) = b_2$ and $\beta(b_1') = b_2'$. Let α be an automorphism of A for which $\alpha(a_1) = a_2$ and $\alpha(a_1') = a_2'$. Then $(a, b) \mapsto (\alpha(a), \beta(b))$ is the desired automorphism.

Case 2. Suppose $\beta(b_1) = b_2'$ and $\beta(b_1') = b_2$. Let α be an automorphism of A for which $\alpha(a_1) = a_2'$ and $\alpha(a_1') = a_2$. Then $(a, b) \mapsto (\alpha(a), \beta(b))$ is the desired automorphism.

On the other hand, suppose that A is non-trivial and edge-transitive, and $B = K_n^*$. By this section's preliminary remarks (having in mind that A is non-bipartite), all R-classes of A have the same size and induce a $\overline{K_m^*}$. Note that the R-classes of $A \times B$ are $X \times V(K_n^*)$, where X is an R-class of A, and thus all R-classes of $A \times B$ have the same size mn. Also, $(A \times B)/R \cong A/R \times K_n^*/R \cong A/R \times K_1^* \cong A/R$, and A/R is edge-transitive by Lemma 2.1. So we have established that $(A \times B)/R$ is edge-transitive; another application of Lemma 2.1 reveals that $A \times B$ is edge-transitive.

Conversely, suppose $A \times B$ is edge-transitive. Assume first that $A \times B$ is R-thin. Because each R-class of $A \times B$ is the Cartesian product of R-classes in A and B, it follows that A and B are R-thin too. Let X be the product of the prime factors (over $\times$) of A that are also prime factors of B. Thus $A = A' \times X$ and $B = X \times B'$, where A' and B' have no prime factors in common. By Theorem 2.2, any automorphism φ of $A \times B = A' \times X \times X \times B'$ has form

$$\varphi(a, x, x', b) = \big(\, \alpha(a),\ \gamma_A(x, x'),\ \gamma_B(x, x'),\ \beta(b) \,\big) \tag{1}$$

for isomorphisms $\alpha : A' \to A'$, and $\beta : B' \to B'$ and homomorphisms $\gamma_A, \gamma_B : X \times X \to X$. Now we argue that A is edge-transitive. In $A = A' \times X$, fix two arbitrary edges $(a_1, x_1)(a_1', x_1')$ and $(a_2, x_2)(a_2', x_2')$. Then $A \times B$ has edges $(a_1, x_1, x_1, b_1)(a_1', x_1', x_1', b_1')$ and $(a_2, x_2, x_2, b_2)(a_2', x_2', x_2', b_2')$, and an automorphism (1) for which

$$\varphi(a_1, x_1, x_1, b_1) = (a_2, x_2, x_2, b_2) \text{ and } \varphi(a_1', x_1', x_1', b_1') = (a_2', x_2', x_2', b_2').$$

(Note that if it happened that $\varphi(a_1, x_1, x_1, b_1) = (a_2', x_2', x_2', b_2')$, then we could attain the above by relabeling (a_2, x_2) as (a_2', x_2') and vice versa, and b with b'.) From this,

$$(a, x) \mapsto \big(\, \alpha(a),\ \gamma_A(x, x) \,\big) \tag{2}$$

is an automorphism of A carrying (a_1, x_1) to (a_2, x_2), and (a_1', x_1') to (a_2', x_2'). This means A is edge-transitive. Similarly

$$(x, b) \mapsto \big(\, \gamma_B(x, x),\ \beta(b) \big)$$

is an automorphism of B carrying edge $(x_1, b_1)(x_1', b_1')$ to edge $(x_2, b_2)(x_2', b_2')$, mapping endpoints as $(x_1, b_1) \mapsto (x_2, b_2)$, and $(x_1', b_1') \mapsto (x_2', b_2')$. (In particular this implies B is edge-transitive.)

Now assume B is not arc-transitive (i.e., it is half-transitive); we show that A is arc-transitive. By the above paragraph no automorphism of B maps $(x_1, b_1) \mapsto (x_2', b_2')$ and $(x_1', b_1') \mapsto (x_2, b_2)$. Select an automorphism $\varphi' \in \text{Aut}(A \times B)$ sending the edge $(a_1, x_1, x_1, b_1)(a_1', x_1', x_1', b_1')$ to the edge $(a_2', x_2', x_2, b_2)(a_2, x_2, x_2', b_2')$, and having form

$$\varphi'(a, x, x', b) = \big(\, \alpha'(a),\ \gamma_A'(x, x'),\ \gamma_B'(x, x'),\ \beta'(b) \,\big).$$

5

Then it must be the case that

$$\varphi'(a_1, x_1, x_1, b_1) = (a_2', x_2', x_2, b_2) \text{ and } \varphi(a_1', x_1', x_1', b_1') = (a_2, x_2, x_2', b_2'),$$

because otherwise $(x, b) \mapsto \big(\gamma_B'(x, x), \beta'(b)\big)$ would be an automorphism of B with $(x_1, b_1) \mapsto (x_2', b_2')$ and $(x_1', b_1') \mapsto (x_2, b_2)$, which we have remarked is impossible. As a consequence, the automorphism $(a, x) \mapsto \big(\alpha'(a), \gamma_A'(x, x)\big)$ of A sends (a_1, x_1) to (a_2', x_2'), and (a_1', x_1') to (a_2, x_2). Comparing this to the automorphism (2), we see that A is arc-transitive.

Summary: If $A \times B$ is R-thin and edge-transitive, then A and B are edge-transitive, and at least one of them is arc-transitive.

Now consider the case in which $A \times B$ is not necessarily R-thin. Then $(A \times B)/R \cong A/R \times B/R$, so $A/R \times B/R$ is edge-transitive by Lemma 2.1. Also, A/R and B/R are R-thin; the previous paragraph implies both A/R and B/R are edge-transitive, and at least one of them is arc-transitive. The following cases finish the proof.

Case 1. Suppose both A/R and B/R have loops. Because they are connected edge-transitive graphs, the only possibility is that $A/R = K_1^* = B/R$. But then each of A and B is a K_n^*, and thus so is $A \times B$. But it is also edge-transitive, so $A \times B = K_1^*$. Thus $A = K_1^* = B$, so both factors are (trivially) edge-transitive and arc-transitive.

Case 2. Suppose one of A/R and B/R (say B/R) has a loop but the other does not. Then as in the previous case, $B/R = K_1^*$ and so $B = K_n^*$. Turning our attention to A, we note that A/R is nontrivial (for otherwise A has no edges and hence neither does $A \times B$). All R-classes of $A \times B$ have form $X \times V(K_n^*)$, where X is an R-class of A. But also we have noted that all R-classes of $A \times B$ are of the same size $|X| \cdot n$, whence all R-classes of A have the same size. Then A is edge-transitive by Lemma 2.1. We have now established that A is non-trivial and edge-transitive, and B is a K_n^*.

Case 3. Suppose neither A/R nor B/R has loops. Now, each is non-trivial because if one had no edges then neither would $A \times B$. Because all R-classes of $A \times B$ have the same size, and each is the Cartesian product of R-classes of A and B, it follows that all R-classes of A have the same size, and the same for B. Now Lemma 2.1 says that each of A and B is edge-transitive.

To finish Case 3, it remains to show that one of A or B is arc-transitive. As all R-classes of A have the same size, it is immediate that A is arc-transitive if and only if A/R is, and similarly for B. But we remarked above that one of A/R and B/R is arc-transitive. $\square$

Notice that in Lemma 2.1, every occurrence of the phrase "edge-transitive" can be replaced with either "half-transitive" or "arc-transitive," and the statement remains true. In Theorem 2.3 all graphs are non-bipartite, and in this context "edge-transitive" means either "half-transitive" or "arc-transitive." We can therefore focus the statement of the theorem by replacing all occurrences of "edge-transitive" with "half-transitive" (in the proof as well as the theorem). The proof goes through the same. Similarly, we can replace "edge-transitive" with "arc-transitive" with only the slightest modification of the proof. Following the consequences of this gives a corollary.

Corollary 2.4 *Every connected, non-trivial, edge-transitive non-bipartite graph G has form $G = K_n^* \times H$ (possibly with $n = 1$), where H is non-trivial, has no factor K_n^*, and at most one half-transitive (prime) factor, while all other (prime) factors, if any, are arc-transitive. Furthermore, G is half-transitive if H has a half-transitive factor. Otherwise G is arc-transitive.*

We now contemplate bipartite direct products. In what follows, suppose A and B are connected, A has an odd cycle, and B is bipartite. In such a case $A \times B$ is connected and bipartite. (If A were also bipartite, $A \times B$ would be disconnected; this result goes back to Weichsel [19], see also [6, Theorem 5.9].)

If A and B are edge-transitive and one is arc-transitive, then $A \times B$ is edge-transitive. This was established in the first part of the proof of Theorem 2.3, which did not use non-bipartiteness of the factors. However, this is a sufficient but not necessary condition for $A \times B$ to be edge-transitive. For example, consider the graphs in Figure 1. In each case $A \times K_2$ is the (edge-transitive) cube Q_3, but not all of the A are edge-transitive.


Figure 1: Three graphs A for which $A \times K_2 = Q_3$ is edge-transitive.

We might call a graph A (such as those in Figure 1) *quasi-edge-transitive* if $A \times K_2$ is edge-transitive, and regard this as a "weak" form of edge-transitivity. Observe that edge-transitive graphs are quasi-edge-transitive, but not conversely. For example, K_n^* is quasi-edge-transitive because $K_n^* \times K_2 = K_{n,n}$ is edge-transitive.

Proposition 2.5 *Suppose A has an odd cycle and B is bipartite. If both $A \times K_2$ and B are edge-transitive and one is arc-transitive, then $A \times B$ is edge-transitive.*

Proof. Notice that $(A \times K_2) \times B \cong K_2 \times (A \times B)$ is the disjoint union of two copies of $A \times B$. But also, as $A \times K_2$ and B are both edge-transitive, and one is arc-transitive, $(A \times K_2) \times B$ is edge-transitive. We have thus established that a disjoint union of two copies of $A \times B$ is edge-transitive. Certainly, then, one copy is edge-transitive too. $\square$

We conjecture that the converse holds.

3 Finite strong products

In treating the strong product we use an equivalence relation S on a graph's vertices. Recall that the *closed neighborhood* $N[x]$ of a vertex $x \in V(G)$ is the set containing

x and its neighbors in G. Then for two vertices x and y, we say xSy provided that $N[x] = N[y]$. See Chapter 7 of [6] for details. In particular recall that S-classes of $A \boxtimes B$ are characterized as Cartesian products $X \times Y$ of S-classes of the factors. A graph is called S-*thin* if each of its S-equivalence classes is a single vertex.

It is immediate that any S-equivalence class induces a complete graph. It follows that a connected graph with two S-classes, at least one of which is non-trivial, cannot be edge transitive, because an edge with endpoints in a non-trivial S-class cannot be moved to an edge joining two different S-classes.

The article [7] (also Chapter 7 of [6]) defines a so-called *Cartesian skeleton* $S[G]$ of a graph G having the following properties: If G is connected then $S[G]$ is a connected, spanning subgraph of G. Moreover if A and B are S-thin, then $S[A \boxtimes B] = S[A] \square S[B]$. And, regardless of S-thinness, any isomorphism $G \to H$ restricts to an isomorphism $S[G] \to S[H]$. (See [3] for a related construction of a Cartesian skeleton.)

Theorem 3.1 *The strong product $G = A \boxtimes B$ of two connected, non-trivial graphs is edge-transitive if and only if both factors are complete.*

Proof. If both factors are complete, then $A \boxtimes B$ is complete, so it is edge-transitive.

Conversely, suppose that $G = A \boxtimes B$ is edge-transitive. Our strategy is to show that G has only one S-class. Then G is complete and hence also are A and B.

Indeed, by the above remarks, if G had several S-classes, they would all be single vertices. As the S-classes of G are products of the S-classes of A and B, the S-classes of A and B are single vertices, and both A and B are S-thin. Then any automorphism of $A \boxtimes B$ restricts to an automorphism of $S[A \boxtimes B] = S[A] \square S[B]$. Because A and B are both nontrivial, it is immediate that $S[A] \square S[B] = S[A \boxtimes B]$ is a proper nontrivial subgraph of $A \boxtimes B$. But then no automorphism can move an edge of this subgraph to an edge not on it, contradicting our assumption that G is edge-transitive. Then G has only one S-equivalence class and is thus complete. $\square$

4   Weak Cartesian products

Having dealt with the edge-transitivity of finite product graphs, we now turn to infinite Cartesian products. In [10] it was shown that a finite, connected Cartesian product graph is edge-transitive if and only if it is the Cartesian power of a connected, edge- and vertex-transitive graph. Here we characterize connected, edge-transitive Cartesian product graphs of arbitrary cardinality. In addition, we do the same for vertex-transitivity. The key is that every connected graph is a weak Cartesian product of prime graphs, unique up to isomorphisms of the factors.

Let us first recall the definition of the Cartesian and the weak Cartesian product of graphs G_ι, for $\iota \in I$.

The Cartesian product $G = \square_{\iota \in I} G_\iota$ is defined on the vertex set $V(G)$ that consists of all functions $x : \iota \mapsto x_\iota$, with $x_\iota \in V(G_\iota)$. The x_ι are the *coordinates* of x, and two vertices x and y are adjacent in G if there exist a $\kappa \in I$ such that $x_\kappa y_\kappa \in E(G_\kappa)$ and $x_\iota = y_\iota$ for $\iota \in I \setminus \kappa$.

If I is finite, then G is connected if and only if all G_ι are connected. But if I is infinite and the G_ι are non-trivial, that is, if they have at least two vertices, then G cannot be connected, even if all factors G_ι are. The reason is that in this case G has at least two vertices x, y that differ in infinitely many coordinates, and they cannot be connected by a path of finite length, as the endpoints of any edge differ in only one coordinate.

This leads to the definition of the weak Cartesian product. A *weak Cartesian product* of graphs G_ι, $\iota \in I$, is a connected component of the Cartesian product of the G_ι. To identify the component it suffices to specify a vertex, say $a \in V(G)$, that it contains. We use the notation $\square_{\iota \in I}^a G_\iota$ or $\square_{\iota \in I}(G_\iota, a_\iota)$ for the connected component of $\square_{\iota \in I} G_\iota$ containing a.

Clearly $V(\square_{\iota \in I}^a G_\iota)$ consists of a and all vertices of $\square_{\iota \in I} G_\iota$ that differ from a in only finitely many coordinates.

For finite I the weak Cartesian product of connected graphs coincides with the Cartesian product. Notice that we did not order the factors in our definition of the product. Sometimes this may be useful though, and so we recall that the Cartesian product is commutative and associative with the trivial graph K_1 as a unit. Furthermore, if $J \subset I$, then

$$\square_{\iota \in I}(G_\iota, a_\iota) \cong \left(\square_{\iota \in J}(G_\iota, a_\iota) \right) \square \left(\square_{\iota \in I \setminus J}(G_\iota, a_\iota) \right).$$

A special case is $J = \{\kappa\}$. Then

$$\square_{\iota \in I}^a G_\iota \cong G_\kappa \square \square_{\iota \in I \setminus \kappa}^{a'} G_\iota,$$

where a' is the restriction of the function $a : I \to \bigcup_{\iota \in I} V(G_\iota)$ to $I \setminus \kappa$. We call a' the projection of a into $V(\square_{\iota \in I \setminus \kappa}(G_\iota, a_\iota))$.

Non-trivial graphs that cannot be represented as a product of two graphs with at least two vertices are called *prime*. It is well known (see [8, 14] or [6]) that any connected graph can be represented as a Cartesian or weak Cartesian product of prime graphs. The representation is unique up to isomorphisms of the factors, and if $\square_{\iota \in I}^a G_\iota = \square_{\iota \in I}^b G_\iota$, then a and b differ in at most finitely many coordinates.

4.1 Vertex-transitive Cartesian products

In order to characterize weak Cartesian products that are vertex-transitive, we need knowledge about the structure of the automorphism group of weak Cartesian products. By [6, p. 413] or [8, 9, 15] the description for finitely or infinitely many factors is the same, so we use [6, Theorem 6.10] although it was originally written for finitely many factors:

Proposition 4.1 *If $G = \square_{\iota \in I}^a G_\iota$, then every $\varphi \in \mathrm{Aut}(G)$ is of the form*

$$\varphi(x)_\iota = \varphi_\iota(x_{\pi(\iota)}),$$

where π is a permutation of the index set I and $\varphi_\iota \in \mathrm{Aut}(G_\iota)$ for all $\iota \in I$. Furthermore, only finitely many φ_ι move a coordinate of a in the sense that $\varphi(a)_\iota \neq a_\iota$.

The condition that only finitely many φ_ι move a coordinate of a ensures that φ preserves the connected component of a in $\square_{\iota \in I} G_\iota$.

Special cases occur when π consists of a transposition (κ, λ), and when all φ_ι are the identity automorphisms. We call such a mapping a *transposition of factors*. On the other hand, if π is the identity and all φ_ι are the identity mapping, with the sole exception of φ_κ, then we write $\varphi = (\varphi_\kappa)$ and say φ is *induced by the automorphism* φ_κ of G_κ. For finite I one can say that all automorphisms of G are generated by transpositions and automorphisms of the factors.

Notice that $x_\lambda = y_\lambda$ for two vertices x and y if and only if $\varphi(x)_{\pi^{-1}(\lambda)} = \varphi(y)_{\pi^{-1}(\lambda)}$. This means that if two vertices differ only in the λ-coordinate, then their images under φ differ only in the $\pi^{-1}(\lambda)$ coordinate.

Hence, the subgraph of G induced by all vertices that differ from a vertex a only in the λ-coordinate is mapped into the subgraph of G induced by all vertices that differ from $\varphi(a)$ only in the $\pi^{-1}(\lambda)$ coordinate. Such a subgraph of G that is induced by all vertices that differ from a vertex a only in the λ-coordinate is called the G_λ-*layer of* G *through* a and denoted by G_λ^a. Thus, φ maps the set of G_λ-layers into the set of $G_{\pi^{-1}(\lambda)}$-layers. Notice that every G_λ-layer is an isomorphic copy of G_λ. Furthermore, every edge ab joins two vertices that differ in exactly one coordinate. If this coordinate is κ, then ab is in the κ-layer through a, and this is the only layer that contains ab.

Before stating our main theorem on vertex-transitivity of weak Cartesian products, we give an example to put it in context. Consider the path P_3 of length 2. It is edge-transitive, but not vertex-transitive, and no finite Cartesian power of it is vertex-transitive. However, if we take infinitely many copies of $G_\iota = P_3$, where ι belongs to an index set I of any infinite cardinality, then $G = \square_{\iota \in I}^a G_\iota$ is vertex-transitive whenever infinitely many a_ι have degree 1 in G_ι and infinitely many degree 2. To see this, we just have to prove that for any vertex b of G with this property, and an adjacent vertex c, some automorphism of G carries b to c. Indeed, relabel a subset of I with the index set $\mathbb{Z}$, so that b and c differ only in the coordinate 0, and b_i has degree 1 if $i > 0$ and degree 2 if $i < 0$. Then there is a series of isomorphisms

$$\cdots \to G_{-3} \to G_{-2} \to G_{-1} \to G_0 \to G_1 \to G_2 \to G_3 \to \cdots$$

such that either this series or its reverse maps b to c, when we restrict them to the indices $\mathbb{Z} \subseteq I$. Now extend this to an automorphism φ of G by declaring that φ is the identity on $I \setminus \mathbb{Z}$. Then $\varphi(b) = c$, proving that G is vertex-transitive.

We will later show that this also ensures that G is edge-transitive, and that G belongs to the interesting class of *half-transitive graphs*.

We continue with a result about the number of vertex-orbits in the powers of a connected graph H with two orbits. In [10] it was shown that every finite power H^i, $2 \leq i < \aleph_0$, of such a graph H has at least three vertex-orbits. Here we extend it to infinite powers.

Lemma 4.2 *Let H be a connected graph with two vertex-orbits and G be the weak Cartesian product $\square_{\iota \in I}^{a} G_\iota$, where $G_\iota \cong H$ and $\aleph_0 \leq |I|$. Then G has either only one or infinitely many vertex-orbits.*

Proof. Let V_1 and V_2 be the vertex-orbits of H. Recall that P_3 from the above example also has two orbits, one consisting of the vertices of degree 1, and the other of the vertex of degree 2. We assumed that $G_\iota = P_3$, where ι belongs to an index set I of any infinite cardinality, and showed that $G = \square_{\iota \in I}^{a} G_\iota$ is vertex-transitive whenever infinitely many a_λ have degree 1 in G_λ and infinitely many degree 2. By practically the same argument one sees that G is vertex-transitive if infinitely many components a_ι of a are in the vertex-orbit of H_ι that corresponds to V_1 and infinitely many in the one that corresponds to V_2.

If only finitely many components a_ι of a are in the vertex-orbit of H_ι that corresponds to, say V_1, then we can replace the root a in the definition of G by b, where all components of b_ι of b are in the vertex-orbit of H_ι that corresponds to V_1.

We now choose a natural number n and a vertex v of G that has exactly n components in the vertex-orbit of H_ι that corresponds to V_2. Then this also holds for all vertices $\mathrm{Aut}(G)(v)$, that is, for all vertices in the orbit of v under the action of $\mathrm{Aut}(G)$. Since there are infinitely many natural numbers, we can thus construct infinitely many orbits. $\qquad\square$

We are ready for our theorem on vertex-transitivity of weak Cartesian products.

Theorem 4.3 *The weak Cartesian product $G = \square_{\iota \in I}^{a} G_\iota$ of connected, prime graphs G_ι is vertex-transitive under the following necessary and sufficient condition:*

If v is a vertex of a G_κ that is not vertex-transitive, then there is an infinite set $K \subseteq I$, and isomorphisms $\varphi_{\lambda,\kappa} : G_\lambda \to G_\kappa$ for every $\lambda \in K$, with $\varphi_{\lambda,\kappa}(a_\lambda) = v$.

Proof. It is easy to see that the G is vertex-transitive if it satisfies the condition of the theorem. We merely adapt the argument used in the above discussion involving the product of copies of P_3.

Conversely, assume that the condition is not met. Then there must be a factor G_κ that has at least two vertex-orbits under the action of $\mathrm{Aut}(G_\kappa)$, say X and Y. Let $K \subseteq I$ be the set of indices λ such that G_λ is isomorphic to G_κ, say via an isomorphism $\varphi_{\lambda,\kappa}$. For any vertex $v \in V(G)$ let $v(X)$ denote the number of coordinates v_λ of v that are in $(\varphi_{\lambda,\kappa})^{-1}(X)$.

Clearly $v(X)$ remains constant under automorphisms of G. However, if there is a v for which $v(X)$ is finite, there is a vertex w for which $w(X) \neq v(X)$ that differs from v only in the κ-coordinate. For, if $v_\kappa \in X$, then we choose $w_\kappa \in Y$. Otherwise, if $v_\kappa \notin X$, then we choose $w_\kappa \in X$.

The observation that all $v(X)$ are finite if the set $K \subseteq I$ of indices to which there are isomorphisms $\varphi_{\lambda,\kappa} : G_\lambda \to G_\kappa$ for every $\lambda \in K$, such that $\varphi_{\lambda,\kappa}(a_\lambda) \in \mathrm{Aut}(G_\kappa)(v)$, is finite completes the proof.

Notice that this is possible, even if there are infinitely many factors that are isomorphic to G_κ. $\qquad\square$

Corollary 4.4 *If all G_ι, $\iota \in I$, are vertex-transitive, then $G = \square^a_{\iota \in I} G_\iota$ is also vertex transitive.*

Corollary 4.5 *If I is finite, then G is vertex-transitive if and only if all G_ι, $\iota \in I$, are vertex-transitive.*

4.2 Edge-transitive Cartesian products

The article [10, Theorem 3.1 and Corollary 3.2] shows that a finite, connected Cartesian product G is edge-transitive if and only if it is the power of a connected, edge- and vertex-transitive prime graph H, and that G is half-transitive if and only if H is half-transitive. Now we extend this result to infinite graphs. We begin with two lemmas.

Lemma 4.6 *Let H be a connected, edge- and vertex-transitive graph and $G = \square^a_{\iota \in I} H_\iota$, where $H_\iota \cong H$ and $2 \leq |I|$. Then G is also edge- and vertex-transitive. Furthermore, G is half-transitive if and only if H is half-transitive.*

Proof. By Lemma 4.4 G is vertex transitive. Now, let uv, xy be arbitrary edges of G, where uv is in the κ-layer G^v_κ and xy in the λ-layer G^x_λ. By Proposition 4.1 there is an automorphism, say α, that maps G^v_κ into G^x_λ. Then $\alpha(u)\alpha(v) \in E(G^x_\lambda)$, and since H is edge-transitive there is an automorphism β_λ of $G_\lambda \cong H$ that maps $\alpha(u)_\lambda \alpha(v)_\lambda$ into $x_\lambda y_\lambda$. Let (β_λ) denote the automorphism of G that is induced by β_λ, then $(\beta_\lambda)\alpha$ maps uv into xy, and therefore G is edge-transitive.

Furthermore, if G is half-transitive, the the restriction of $\mathrm{Aut}(G)$ to any layer, say to G^x_λ, is also half-transitive, and if G is edge-transitive, but not half-transitive, the the restriction is also not half-transitive. The observation that $G^x_\lambda \cong G_\lambda \cong H$ completes the proof. $\square$

Lemma 4.7 *Let H be a connected, edge-transitive but not vertex-transitive graph, bi-partitioned by its two vertex orbits V_1 and V_2. Let $G = \square^a_{\iota \in I} H_\iota$, where $H_\iota \cong H$, $2 \leq |I|$, and where infinitely many of the a_ι are in the vertex-orbit of G_ι corresponding to V_1, and infinitely many in the vertex-orbit of G_ι corresponding to V_2. Then G is edge-transitive (but only half-transitive) and vertex-transitive.*

Proof. The edge-transitivity of G is shown as in the proof of Lemma 4.6. This means that G can have only one or two vertex-orbits. By Lemma 4.2 it can only have one or infinitely many vertex-orbits, hence it has only one and is vertex-transitive. If G were not half-transitive, that is, if it were arc-transitive, then it would have to be arc-transitive on every layer, and hence H would have to be arc-transitive, contrary to assumption. Hence G is half-transitive. $\square$

Theorem 4.8 *Let G be a connected, edge-transitive graph that is not prime with respect to the Cartesian product. Then G is the Cartesian or weak Cartesian power of a connected, edge-transitive graph H.*

For vertex-transitive H the structure of G is described by Lemma 4.6, otherwise by Lemma 4.7. In both cases G is vertex-transitive.

Proof. Let G be a connected, edge-transitive graph that is not prime with respect to the Cartesian product, and $\square_{\iota \in I}^a G_\iota$ be its prime factorization. We show first that all factors G_ι must be isomorphic.

Take any two indices $\kappa, \lambda \in I$, and two arbitrary edges uv, xy of G, where uv is in the κ-layer G_κ^v and xy in the λ-layer G_λ^x. Any automorphism that maps uv into xy has to map G_κ^v into G_λ^x. Hence these layers, and thus G_κ and G_λ are isomorphic. Thus G is the Cartesian power or weak Cartesian power of a connected graph H.

Also, H must be edge-transitive. To see this, consider any layer G_ι^v and two, not necessarily distinct, edges uv and xy of that layer. The automorphism of G that maps uv into xy has to preserve that layer and thus induces an automorphism of G_ι^v. Since uv and xy were arbitrarily chosen in G_ι^v, this layer, and hence H, must be edge-transitive.

Thus $G = \square_{\iota \in I}^a H_\iota$, where each H_ι is isomorphic to an edge-transitive graph H. Clearly H must be connected. We have to consider two cases.

Case 1. H is vertex-transitive. Clearly this implies that G is vertex-transitive, independently of the choice of a, and whether I is finite or not. Half-transitivity is dealt with as in Lemma 4.6.

Case 2. H is not vertex-transitive. As it is edge-transitive, it must be bipartite by Lemma 1.1, and the bipartition is given by the vertex-orbits.

If I is finite, then this implies that G has a least three vertex-orbits by [10]. But, since G is edge-transitive, it can have at most two vertex-orbits. So I must be infinite.

If I is infinite, then G has either one or infinitely many vertex-orbits by Lemma 4.2. Hence G must be vertex-transitive and, again by Lemma 4.2, the root a must have the structure described in Lemma 4.7. The arguments there also show that G is half-transitive. $\square$

Acknowledgment

References

[1] J. Chen, C .H. Li, Á. Seress, A family of half-transitive graphs, Electron. J. Combin. 20(1) (2013) Paper #56, 10 pp.

[2] W. Dörfler, Primfaktorzerlegung und Automorphismen des Kardinalproduktes von Graphen, Glasnik Mat. 9 (1974) 15–27.

[3] J. Feigenbaum, A. A. Schäffer, Finding prime factors of strong direct product graphs in polynomial time, Discrete Math. 109 (1992) 77–102.

[4] M. Giudici, P. Potočnik, G. Verret, Semiregular automorphisms of edge-transitive graphs, J. Algebraic Combin. 40 (2014) 961–972.

[5] C. Godsil, G. Royle, Algebraic Graph Theory, Springer-Verlag, New York, 2001.

[6] R. Hammack, W. Imrich, S. Klavžar, Handbook of Product Graphs, Second Edition, CRC Press, Boca Raton, FL, 2011.

[7] R. Hammack, W. Imrich, On Cartesian skeletons of graphs, Ars Math. Contemp. 2 (2009) 191–205.

[8] W. Imrich, Automorphismen und das kartesische Produkt von Graphen, Österreich. Akad. Wiss. Math.-Natur. Kl. S.-B. II 177 (1969) 203–214.

[9] W. Imrich, Embedding graphs into Cartesian products, Ann. N.Y. Acad. Sci. 576 (1989) 266–274.

[10] W. Imrich, A. Iranmanesh, S. Klavžar, A. Soltani, Edge-transitive lexicographic and Cartesian products, submitted.

[11] C. H. Li, Z. H. Heng, Finite vertex-biprimitive edge-transitive tetravalent graphs, Discrete Math. 317 (2014) 33–43.

[12] F. Li, W. Wang, Z. Xu, H. Zhao, Some results on the lexicographic product of vertex-transitive graphs, Appl. Math. Lett. 24 (2011) 1924–1926.

[13] G. X. Liu, Z. P. Lu, On edge-transitive cubic graphs of square-free order, European J. Combin. 45 (2015) 41–46.

[14] D. J. Miller, Weak cartesian product of graphs, Colloq. Math. 21 (1970) 55–74.

[15] D. J. Miller, The automorphism group of a product of graphs, Proc. Amer. Math. Soc. 25 (1970) 24–28.

[16] G. Sabidussi, The composition of graphs, Duke Math. J. 26 (1959) 693–696.

[17] G. Sabidussi, Vertex-transitive graphs, Monatsh. Math. 68 (1964) 426–438.

[18] M.E. Watkins, Edge-transitive strips, Discrete Math. 95 (1991) 359–372.

[19] P. M. Weichsel, The Kronecker product of graphs, Proc. Amer. Math. Soc. 13 (1962) 47–52.