2017

Edge-Transitive Bipartite Direct Products

Cameron M. Crenshaw
Virginia Commonwealth University

Edge-Transitive Bipartite Direct Products

A thesis submitted in partial fulfillment of the requirements for the degree of Master of Science at Virginia Commonwealth University.

by

Cameron Crenshaw
Master of Science

Director: Richard Hammack, Professor
Department of Mathematics and Applied Mathematics

Virginia Commonwealth University
Richmond, Virginia
May 2017

Acknowledgements

I want to thank my mother for her help with all things non-mathematical, and my father for teaching me to love asking questions.

I also want to thank all of the faculty in the math department at VCU, particularly Dr. Hammack, whose patience, knowledge, and creativity inspire me.

Table of Contents

List of Figures

Abstract

In their recent paper "Edge-transitive products," Hammack, Imrich, and Klavžar showed that the direct product of connected, non-bipartite graphs is edge-transitive if and only if both factors are edge-transitive, and at least one is arc-transitive. However, little is known when the product is bipartite. This thesis extends this result (in part) for the case of bipartite graphs using a new technique called ″stacking.″ For R-thin, connected, bipartite graphs A and B, we show that $A \times B$ is arc-transitive if and only if A and B are both arc-transitive. Further, we show $A \times B$ is edge-transitive only if at least one of A, B is also edge-transitive, and give evidence that strongly suggests that in fact both factors must be edge-transitive.

Chapter 1

Introduction

1.1 Essentials from Graph Theory

A **graph** is a set of **vertices**, or "points," and a set of **edges** that connect them. For a graph G, we denote its set of vertices as $V(G)$, and edges as $E(G)$. Specifically, we consider edges as unordered pairs, so an edge between vertices a and b would be written $\{a, b\}$, or typically ab for short. In this case we say the vertices a and b are **adjacent** since there is an edge that connects them.

The **neighborhood** of a vertex v in a graph G, denoted $N_G(v)$, is the set of all vertices in G adjacent to v, and its members are called **neighbors** of v. This set is often abbreviated as $N(v)$ when the graph is clear from context. If it happens that there is an edge between every distinct vertex of a graph then we call the graph **complete**, and denote the complete graph on n vertices by K_n. We say H is a **subgraph** of G if $V(H) \subseteq V(G)$, and $E(H) \subseteq E(G)$.

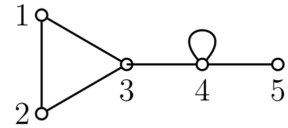

Figure 1.1: An example of a graph

An edge that joins a vertex to itself is called a **loop**. Frequently it is necessary to restrict graphs from having loops, so the class of graphs without loops allowed is denoted by Γ, and the class of graphs with loops allowed is Γ_0, thus $\Gamma \subset \Gamma_0$.

Walks are finite sequences of vertices $(v_1, v_2, \ldots, v_n)$, where at each step, $v_i v_{i+1}$ is an edge in the graph. A walk that begins and ends at the same vertex is called a **closed walk**. Walks are free to repeat vertices, but if it happens that no vertices are repeated then we have a **path**. Closed paths are called **cycles**. Note that the number of vertices on a cycle always equals the number of edges; this quantity is called the **length** of the cycle, where cycles with even length are called **even cycles** and those with odd length are called **odd cycles**. The cycle of length ℓ is denoted C_ℓ.

A graph is called **connected** if and only if for all pairs of distinct vertices $u, v \in V(G)$, there exists a path connecting them. A **component** of a graph is a maximally connected subgraph.

We are particularly interested in the class of bipartite graphs. A graph B is **bipartite** if and only if there exists a partition $\{X, Y\}$ of $V(B)$ such that every edge has exactly one endpoint on each side of the partition. We call X and Y **partite sets**. In any cycle C of a bipartite graph, the vertices must alternate between the partite sets as we travel around the cycle. Since cycles are closed, C must finish in the same partite set in which it began, so it must have even length. This necessary condition is also sufficient, giving an essential characterization of bipartite graphs. This is summarized in our first theorem:

Theorem 1. *A graph is bipartite if and only if it contains no odd cycles.*

[1] is an excellent reference for definitions, notation, or as a general introduction to graph theory.

Chapter 2

Background

2.1 Automorphisms and Transitivity

We begin this section with the notion of a graph isomorphism. For graphs G and H, an **isomorphism** φ from G to H is a function $\varphi : V(G) \to V(H)$ which is bijective and preserves adjacency. That is, φ is one-to-one, onto, and gg' is an edge of G if and only if $\varphi(g)\varphi(g')$ is an edge of H. When such a function exists, we say G and H are **isomorphic** and write $G \cong H$. Further, if G has finitely many vertices, then a bijective function $\varphi : V(G) \to V(G)$ for which $gg' \in E(G)$ implies $\varphi(g)\varphi(g') \in E(G)$ is also an isomorphism.

If $\varphi : V(G) \to V(G)$ is an isomorphism, then we call φ an **automorphism**. Note that the identity map is an automorphism, composition of automorphisms is again an automorphism, and the inverse of an automorphism is an automorphism. Thus, the set of all automorphisms of a graph G form a group under function composition, called the **automorphism group** of G, $\mathrm{Aut}(G)$.

A graph G is called **vertex-transitive** if for every pair of vertices u and v of G, there is an automorphism $\varphi \in \mathrm{Aut}(G)$ such that $\varphi(u) = v$. It is **edge-transitive** if for every pair of edges $g_1g_2, g_1'g_2' \in E(G)$, there is an automorphism $\varphi \in \mathrm{Aut}(G)$ such that either

$\varphi(g_1) = g_1'$ and $\varphi(g_2) = g_2'$, or $\varphi(g_1) = g_2'$ and $\varphi(g_2) = g_1'$. Finally, G is **arc-transitive** if for every pair of edges $g_1g_2, g_1'g_2' \in E(G)$, there is an automorphism $\varphi \in \mathrm{Aut}(G)$ such that $\varphi(g_1) = g_1'$ and $\varphi(g_2) = g_2'$, and *another* automorphism $\varphi' \in \mathrm{Aut}(G)$ with $\varphi'(g_1) = g_2'$ and $\varphi'(g_2) = g_1'$. Note that arc-transitive implies edge-transitive, and implies vertex-transitive provided the graph has no isolated vertices.

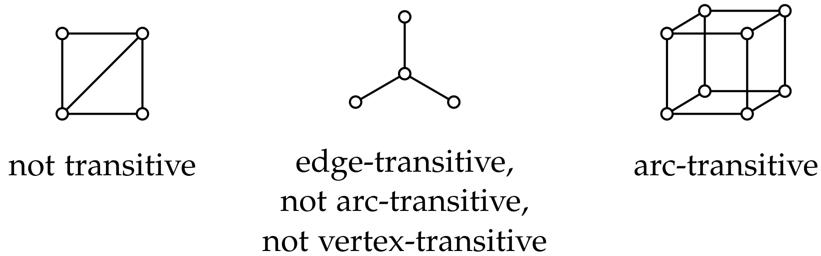

Figure 2.1: Some examples of transitivity

Definition 1. *Let gg' be an edge of a graph G. If there exists an automorphism φ of G such that $\varphi(g) = g'$ and $\varphi(g') = g$, then we call gg' a **reversible edge** of G.*

With this definition in mind, we will use the following Proposition as a shortcut for showing arc-transitivity:

Proposition 1. *If G is edge-transitive with a reversible edge, then G is arc-transitive.*

Proof. Let the edge uv of G be reversible by the automorphism φ, and aa', bb' be arbitrary edges of G. Since G is edge-transitive there is some automorphism φ_1 sending aa' to bb', say by $\varphi_1(a) = b$ and $\varphi_1(a') = b'$. We show there exists another automorphism φ_2 with $\varphi_2(a) = b'$ and $\varphi_2(a') = b$, proving that G is arc-transitive.

By the edge-transitivity of G, we have automorphisms α and β sending aa' to uv, and uv to bb', respectively. If these automorphism behave according to either of the following cases,

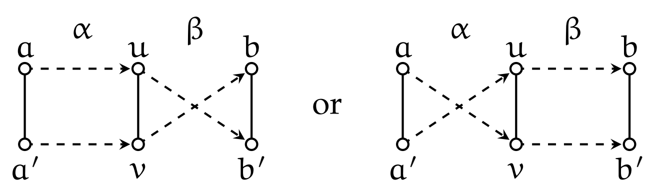

then the composition $\beta \circ \alpha$ produces the automorphism φ_2 desired. Therefore we may assume one of the following:

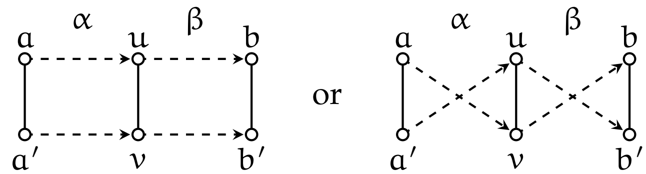

In both of these cases, the composition $\beta \circ \varphi \circ \alpha$ gives the automorphism φ_2 desired. Thus, G is arc-transitive. □

2.2 The Direct Product

Direct products of graphs are our main objects of interest, so we define this product and discuss some useful results. When working with the direct product, we consider graphs from the larger class Γ_0.

Given a pair of graphs G and H, the **direct product** of G and H in Γ_0, $G \times H$, is the graph with vertex set $V(G) \times V(H)$ (the set Cartesian product), and edge set

$$E(G \times H) = \{(g,h)(g',h') : gg' \in E(G) \text{ and } hh' \in E(H)\}.$$

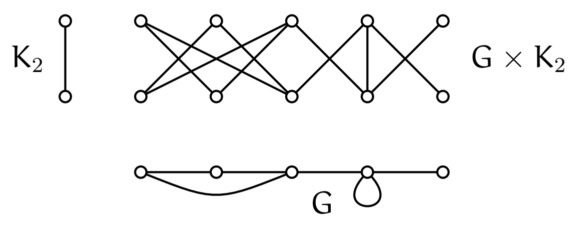

Figure 2.2: The direct product of graphs

The direct product is commutative, associative, and is distributive over "+", the disjoint union of graphs, i.e.,

$$A \times (B + C) \cong A \times B + A \times C.[4]$$

Bipartite graphs have a particular significance for the direct product. First, the direct product of a bipartite graph with an arbitrary graph is always bipartite. Next, bipartiteness characterizes the connectivity of a direct product:

Theorem 2 (Weichsel, Theorem 5.9 of [4])**.** *Let* G *and* H *be connected graphs. Then* $G \times H$ *is connected if and only if at most one of its factors is bipartite. If* G *and* H *are both bipartite, then* $G \times H$ *has exactly two components.*

We denote the graph of a single vertex with a loop as K_1^*. From the definition of the direct product, K_1^* serves as an identity in that $K_1^* \times G \cong G$ for all $G \in \Gamma_0$. We say a graph G is **prime** with respect to the direct product if it has at least two vertices and $G \cong H \times K$ implies that either H or K is K_1^*. While it is known that connected non-bipartite graphs have unique prime factorization under the direct product [4], the same is not true of bipartite graphs. For example, below are two prime factorizations of the connected bipartite graph C_6:

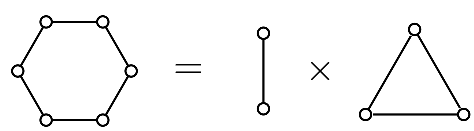

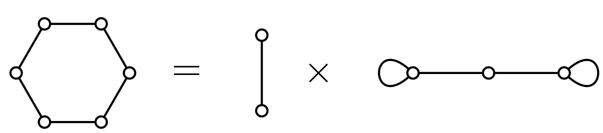

It is mainly for this reason that we turn to the Cartesian product of graphs.

2.3 The Cartesian Product

We use the Cartesian product to bypass some of the obstacles presented with the direct product. Given a pair of graphs G and H in Γ, the **Cartesian product** of G and H, $G \,\square\, H$, is the graph with vertex set $V(G) \times V(H)$, and edges

$$E(G \,\square\, H) = \{(g,h)(g',h') : gg' \in E(G) \text{ and } h = h', \text{ or } hh' \in E(H) \text{ and } g = g'\}.$$

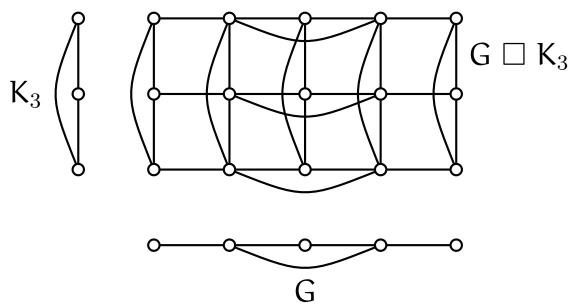

Figure 2.3: The Cartesian product of graphs

As with the direct product, the Cartesian product is also associative, commutative, and distributes over disjoint unions:

$$A \,\square\, (B + C) \cong A \,\square\, B + A \,\square\, C.$$

Connectivity is simpler for the Cartesian product. We have the following as a Corollary in [4]:

Corollary 1. *A Cartesian product is connected if and only if both of its factors are connected.*

We will also need the following cancellation law at one point:

Theorem 3 (Theorem 6.21 of [4])**.** *Let* G*,* H*, and* K *be graphs, with* K *nonempty. If* $G \,\square\, K \cong H \,\square\, K$*, then* $G \cong H$.

Primality with respect to the Cartesian product is the same as for the direct product, only with K_1 instead of K_1^* as the identity. As noted in the section on direct products, one critical advantage of the Cartesian product is uniqueness of prime factorizations for connected graphs. The next theorem is Theorem 6.6 of [4]:

Theorem 4 (Sabidussi-Vizing, [4])**.** *Let G be any connected graph. Then G has a unique representation as the [Cartesian] product of prime graphs, up to isomorphism and ordering of its factors.*

We will work extensively with automorphisms, particularly when applied to coordinatizations of the vertices of a product. This makes Theorem 6.10 from [4] indispensable:

Theorem 5. *Let φ be an automorphism of a connected graph G with prime factorization $G = G_1 \square G_2 \square \cdots \square G_n$. Then there is a permutation π of $\{1, 2, \ldots, n\}$ and isomorphisms $\varphi_i : G_{\pi(i)} \to G_i$ such that*

$$\varphi(g_1, g_2, \ldots, g_n) = \big(\varphi_1(g_{\pi(1)}), \varphi_2(g_{\pi(2)}), \ldots, \varphi_n(g_{\pi(n)})\big).$$

To make the transition from a direct product to a Cartesian product, we employ what is called the **Cartesian skeleton**.

2.4 Cartesian Skeletons and R-thin Graphs

The Cartesian skeleton of a graph serves as a bridge between factorizations under the direct product and factorizations under the Cartesian product. What is presented here are simply the essentials of Cartesian skeletons; greater detail can be found in [2]. We begin with a necessary introduction to the concept of R-thin graphs.

From figure 2.4, notice there is an automorphism of G that swaps the vertices u and v but leaves the others fixed. This automorphism is not induced by any pair of automorphisms of the factors, and clouds the relationship between factorizations of G and its automorphisms, both of which we spend a great deal of time on. The automorphism arises from the fact that $N(u) = N(v)$, so it is convenient to restrict distinct vertices from having identical neighborhoods.

For a graph G, Section 8.2 of [4] defines an equivalence relation R on $V(G)$, where xRy if and only if $N(x) = N(y)$ for $x, y \in V(G)$. A graph is then called **R-thin** if all of its R-equivalence classes contain a single vertex. Equivalently, no two distinct vertices of an R-thin graph have the same neighborhood. At one point, we will use the fact that the direct product of R-thin graphs is again R-thin, which follows immediately from Proposition 8.5 of [4].

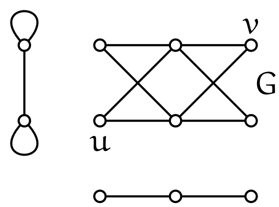

Figure 2.4: Non-example of an R-thin direct product G

The Cartesian skeleton of a graph begins with its Boolean square. The **Boolean square** of G, denoted G^s, has $V(G)$ as its vertices, and has edge set

$$E(G^s) = \{uv : N_G(u) \cap N_G(v) \neq \emptyset\},$$

that is, u and v are adjacent in G^s if and only if they share a neighbor in G. Note that u and v are not required to be distinct. Below is an example of the Boolean square for two factors, and for their direct product.

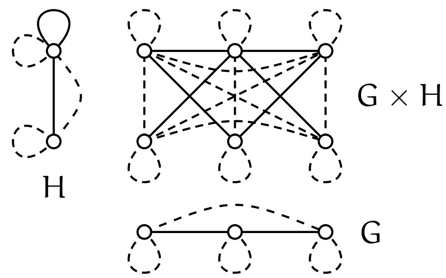

Figure 2.5: Graphs G, H, $G \times H$ (solid) and their Boolean squares (dotted)

From here, we form the **Cartesian skeleton** of G, $S(G)$, by removing all **dispensable** edges from the G^s. An edge $uv \in E(G^s)$ is dispensable if it is a loop, or if there exists $w \in V(G)$ such that both

(i) $N_G(u) \cap N_G(v) \subset N_G(u) \cap N_G(w)$ or $N_G(u) \subset N_G(w) \subset N_G(v)$, and

(ii) $N_G(u) \cap N_G(v) \subset N_G(v) \cap N_G(w)$ or $N_G(v) \subset N_G(w) \subset N_G(u)$

are true. The following figure is the same figure from above with all dispensable edges removed. We can now state the essential results about Cartesian skeletons that are necessary for our purposes.

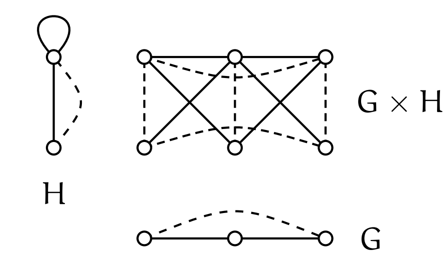

Figure 2.6: Graphs G, H, $G \times H$ (solid) and their Cartesian skeletons (dotted)

Proposition 2 (Prop. 2.7 of [2]). *Let G be a connected graph. Then G is non-bipartite if and only if $S(G)$ is connected. In particular, if G is bipartite then $S(G)$ has two components whose respective vertex sets are the partite sets of G.*

Proposition 3 (Prop. 2.5 of [2]). *Let G and H be R-thin graphs without isolated vertices. Then* $S(G \times H) = S(G) \,\square\, S(H)$.

Proposition 4 (Prop. 2.8 of [2])**.** *If* φ *is an isomorphism from* G *to* H*, then* φ *is also an isomorphism from* $S(G)$ *to* $S(H)$*.*

Chapter 3

Stacking

Here we introduce an operation on the vertex set of a product graph called **stacking**. For our purposes, consider a product $G \square H$ with prime factorization

$$G \square H = \underbrace{F_1 \square \cdots F_k}_{G} \square \underbrace{F_{k+1} \square \cdots \square F_n}_{H},$$

and let φ be an automorphism of $G \square H$. From Theorem 5, there is some permutation π and isomorphisms $\alpha_i : F_{\pi(i)} \to F_i$ such that

$$\varphi(v_1, v_2, \ldots, v_n) = \big(\alpha_1(v_{\pi(1)}), \alpha_2(v_{\pi(2)}), \ldots, \alpha_n(v_{\pi(n)})\big)$$

for $(v_1, v_2, \ldots, v_n) \in V(G \square H)$. In other words, φ restricts to isomorphisms between the factors of $G \square H$, and as a result may permute them. Ideally, φ does not permute factors of G with factors of H; when it does, stacking (defined below) eliminates this "mixing" to a certain extent.

It is convenient to view π as the product of disjoint cycles. This way, we reduce the problem to eliminating H-coordinates only on cycles which permute factors of both G and H. The following figure gives an example of such a mixing cycle:

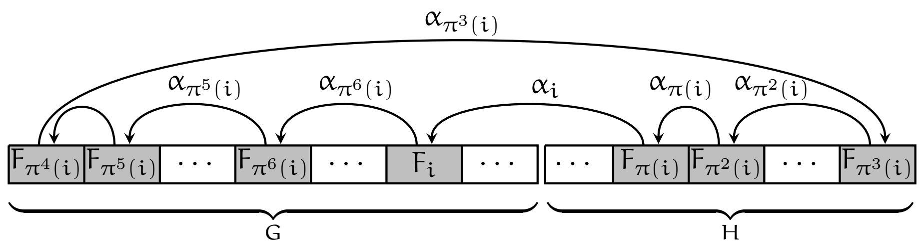

Figure 3.1: A typical mixing cycle

Notice in the figure that when φ is applied to some vertex v of $V(G) \times V(H)$, the new coordinates of $F_{\pi^4(i)}$, $F_{\pi^5(i)}$, and $F_{\pi^6(i)}$ are functions of coordinates from factors of G, while the new F_i-coordinate is a function of a coordinate from H. In this way, some of the G-coordinates of $\varphi(v)$ are dependent on coordinates of H. This dependence on H is what stacking removes.

Given an input vertex $(g_0, h_0) \in V(G \,\square\, H)$, a single iteration of the stacking operation involves the following steps:

1. Apply φ: $\varphi(g_0, h_0) = (g_1, h_1)$;

2. Replace the G coordinate with g_0: $(g_1, h_1) \to (g_0, h_1)$.

Repeating steps 1 and 2 produces a sequence of vertices

$$(g_0, h_0), (g_0, h_1), (g_0, h_2), \dots,$$

which reaches some target vertex (g_0, h_M), whose significance is shown in the next two paragraphs. The next figure follows the effect of each stacking iteration on a vertex $(v_1, v_2, \dots, v_3, \dots, v_4, \dots)$ of $G \,\square\, H$, with $v_i \in G$, where we are only showing the G-coordinates of this vertex which correspond to the coordinates on the cycle in Figure 3.1. The asterisks in Figure 3.2 represent coordinates from H along the cycle, whose only significance is that they are not elements of G. For notational convenience, we denote $\alpha_{\pi^3(i)}(v_1)$, $\alpha_{\pi^2(i)}\alpha_{\pi^3(i)}(v_1)$, and $\alpha_{\pi(i)}\alpha_{\pi^2(i)}\alpha_{\pi^3(i)}(v_1)$ by v_1', v_1'', and v_1''', respectively.

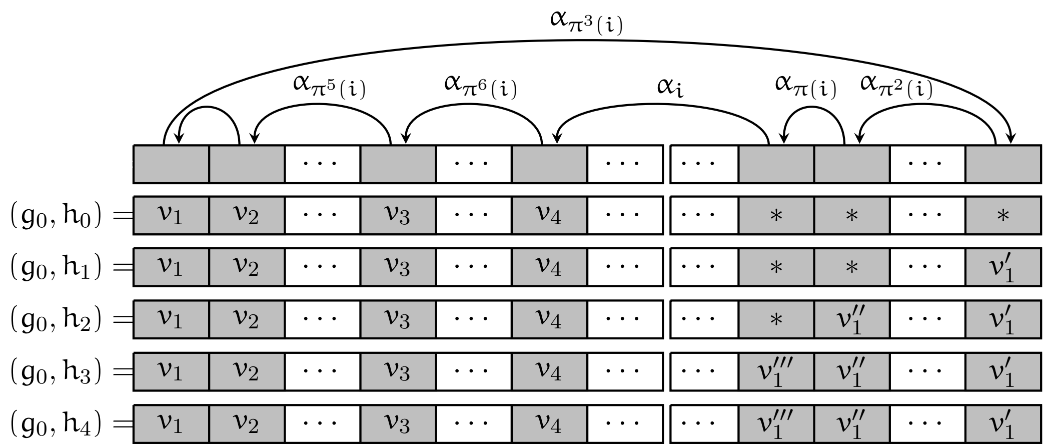

Figure 3.2: Following the stacking operation

At each stage, α_i maps a coordinate of H into the G-coordinates, but continuously "rewriting" the G-coordinates prevents this mixing. After three iterations, α_i is now sending the image of a G-*coordinate* back to the G-coordinates. Further, as illustrated by (g_0, h_4), more stacking iterations don't cause any additional changes to this cycle. For this reason, we may apply the stacking operation repeatedly to a given vertex as many times as necessary until all the cycles have independently eliminated their H-coordinates. So, after $M \geqslant n - k$ iterations, we reach the fully stacked vertex (g_0, h_M).

Applying φ to this fully stacked vertex (g_0, h_M) gives a new vertex whose G-coordinates are functions only of other G-coordinates. That is, if $\varphi(g_0, h_M) = (g, h) \in V(G \,\square\, H)$, then g is some function of *only* g_0, while h is some function of both g_0 and h_M. Moreover, the coordinates of g are each coordinates from G mapped by compositions of the α_j's which are each bijections, therefore this is a *bijective* correspondence between g_0 and g. We will frequently capitalize on this fact by defining the bijection $\theta : V(G) \to V(G)$ by $\varphi(g_0, h_M) = (g, h) = \big(\theta(g_0), f(g_0, h_M)\big)$, where f is simply some function from $V(G) \times V(H)$ to $V(H)$.

Chapter 4

Results

4.1 Context

The heart of this paper begins with the following proposition:

Proposition 5. *Let* A *and* B *be edge-transitive graphs. If at least one of them is also arc-transitive, then* $A \times B$ *is edge-transitive. If both are arc-transitive, then* $A \times B$ *is also arc-transitive.*

Proof. Assume without loss of generality that A is arc transitive and B is edge transitive. Let $(a_1, b_1)(a_1', b_1'), (a_2, b_2)(a_2', b_2')$ be arbitrary edges of $A \times B$. From the definition of the direct product, $a_1 a_1', a_2 a_2' \in E(A)$ and $b_1 b_1', b_2 b_2' \in E(B)$.

By the edge-transitivity of B, there is some $\beta \in \mathrm{Aut}(B)$ such that either

(i) $\beta(b_1) = b_2$ and $\beta(b_1') = b_2'$, or

(ii) $\beta(b_1) = b_2'$ and $\beta(b_1') = b_2$.

Case 1: Assume (i) holds. Then since A is arc-transitive, there is an automorphism $\alpha \in \mathrm{Aut}(A)$ such that $\alpha(a_1) = a_2$ and $\alpha(a_1') = a_2'$. From the definition of the direct product, the map $\varphi : V(A) \times V(B) \to V(A) \times V(B)$ defined by $\varphi(a, b) = \big(\alpha(a), \beta(b)\big)$ is an automorphism of $A \times B$. Further, $\varphi(a_1, b_1) = (a_2, b_2)$ and $\varphi(a_1', b_1') = (a_2', b_2')$, thus $A \times B$ is edge-transitive.

Case 2: Assume (ii) holds. In this case, there is another automorphism $\alpha' \in \mathrm{Aut}(A)$ with $\alpha'(a_1) = a_2'$ and $\alpha'(a_1') = a_2$. We then define φ as $\varphi(a, b) = \big(\alpha'(a), \beta(b)\big)$ to get an automorphism of $A \times B$ with $\varphi(a_1, b_1) = (a_2', b_2')$ and $\varphi(a_1', b_1') = (a_2, b_2)$. Thus, $A \times B$ is again edge transitive.

If it happened that B is arc transitive as well, then there would be an automorphism satisfying (i) and another satisfying (ii) as well, so by Case 1 and Case 2 together we get the pair of automorphisms which imply that $A \times B$ is arc-transitive as well. □

In [3] it was shown that the converse holds when A and B are both non-bipartite, but the proof relied on the unique prime factorization of non-bipartite graphs under the direct product. Little progress has been made in the class of bipartite graphs due to this limitation. This brings us to our main theorem.

4.2 Main Theorem

Theorem 6. *Let A and B be connected, bipartite, and R-thin graphs.*

1. *If $A \times B$ is edge transitive, then at least one of its factors is also edge-transitive.*

2. *$A \times B$ is arc-transitive if and only if both A and B are arc-transitive.*

Proof. Note that necessity in statement 2 follows immediately from Proposition 5; we show sufficiency. Along the way we will prove two claims: first, that at least one factor of $A \times B$ must have a reversible edge, and second, that a factor's reversible edge implies edge-transitivity for the other factor.

Let A and B be connected, bipartite, R-thin graphs, and assume that $A \times B$ is edge-transitive.

We begin by examining the Cartesian skeleton $S(A \times B)$. Recall from Proposition 3 that $S(A \times B) = S(A) \,\square\, S(B)$, and since A and B are connected and bipartite, Proposition 2 says their skeletons each have two components, say A_0, A_1 for $S(A)$, and B_0, B_1 for $S(B)$, whose vertex sets correspond to the partite sets of A and B. In this way,

$$\begin{aligned} S(A) \,\square\, S(B) &= (A_0 + A_1) \,\square\, (B_0 + B_1) \\ &= (A_0 \,\square\, B_0) + (A_0 \,\square B_1) + (A_1 \,\square\, B_0) + (A_1 \,\square\, B_1). \end{aligned}$$

Claim I: At least one of A, B has a reversible edge.

Let aa' and bb' be arbitrary edges of A and B respectively. We show that at least one is reversible. By the definition of the direct product, $(a, b)(a', b')$ and $(a', b)(a, b')$ are edges of $A \times B$. Since $A \times B$ is edge-transitive, there is an automorphism φ of $A \times B$ such that either:

i) $\varphi(a, b) = (a', b)$ and $\varphi(a', b') = (a, b')$, or

ii) $\varphi(a, b) = (a, b')$ or $\varphi(a', b') = (a', b)$.

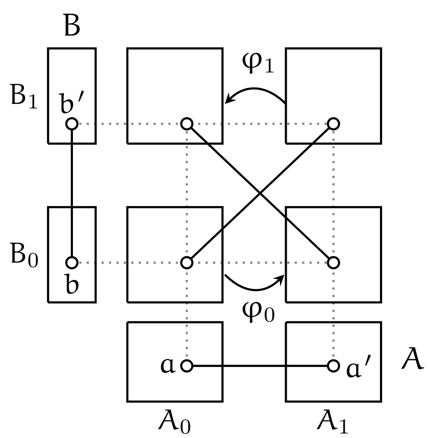

Figure 4.1: The setup for Claim I

Assume i) is true. By Proposition 4, φ is also an automorphism of $S(A \times B)$. Now since $\varphi(a, b) = (a', b)$ and $\varphi(a', b') = (a, b')$, φ restricts to the isomorphisms φ_0 :

$V(A_0 \mathbin{\square} B_0) \to V(A_1 \mathbin{\square} B_0)$, and $\varphi_1 : V(A_1 \mathbin{\square} B_1) \to V(A_0 \mathbin{\square} B_1)$. We may then assume that φ is, in particular, the automorphism that sends $V(A_1 \mathbin{\square} B_0)$ back to $V(A_0 \mathbin{\square} B_0)$ according to φ_0^{-1}, and $V(A_0 \mathbin{\square} B_1)$ back to $V(A_1 \mathbin{\square} B_1)$ according to φ_1^{-1}.

In this way, φ is an involution (φ^2 is the identity map) that swaps the partite sets of $A \times B$. From here, we use φ to define an automorphism of A.

Let (a, b_0) be an arbitrary vertex of $A_0 \mathbin{\square} B_0$. In this special case, φ_0 at most transposes factors, so we need only apply one iteration of stacking, giving $(a, b_1) \in V(A_1 \mathbin{\square} B_0)$. By construction, applying φ_0 once more gives a new vertex of $A_1 \mathbin{\square} B_0$ whose A_1-coordinate depends only on a; define $\theta_0 : V(A_0) \to V(A_1)$ by $\varphi_0(a, b_1) = \big(\theta_0(a), f(a, b_1)\big)$, where f is some function from $V(A_0 \mathbin{\square} B_0)$ to $V(B_0)$.

As a special note, recall that $\varphi_0(a, b) = (a', b)$, so replacing a' with a gives back (a, b). This means the stacking operation does not change the vertex (a, b), so in particular the fully stacked vertex (a, b_1) is actually (a, b). Consequently, $(a', b) = \varphi_0(a, b) = \varphi_0(a, b_1) = \big(\theta_0(a), f(a, b_1)\big)$, so $\theta_0(a) = a'$.

In the same way, define $\theta_1 : V(A_1) \to V(A_0)$ by stacking vertices of $A_1 \mathbin{\square} B_1$ using φ_1, where here $\theta_1(a') = a$.

As reasoned in the stacking introduction, θ_0 and θ_1 are bijections. Define $\theta : V(A) \to V(A)$ by

$$\theta(a) = \begin{cases} \theta_0(a) & \text{if } a \in V(A_0) \\ \theta_1(a) & \text{if } a \in V(A_1). \end{cases}$$

Then θ is a bijection, so it remains to show that it preserves adjacency.

So let uv be an arbitrary edge of A and fix an edge xy of B. Then $(u, x)(v, y)$ is an edge of $A \times B$, say with $(u, x) \in V(A_0) \times V(B_0)$ and $(v, y) \in V(A_1) \times V(B_1)$. Stack the endpoints of this edge simultaneously:

0. Given starting edge $(u, x)(v, y) \in E(A \times B)$

1. Apply φ to each end: $\varphi_0(u, x) = (u_1, x_1)$ and $\varphi_1(v, y) = (v_1, y_1)$, giving $(u_1, x_1)(v_1, y_1) \in E(A \times B)$

2. Replace u_1 and v_1: $(u_1, x_1)(v_1, y_1) \to (u, x_1)(v, y_1) \in E(A \times B)$

In this special case with φ as an involution, applying φ to the ends of the new edge once more produces the edge $\big(\theta_0(u), f(u, x)\big)\big(\theta_1(v), f'(v, y)\big)$ of $A \times B$, therefore $\theta_0(u)\theta_1(v) = \theta(u)\theta(v)$ is an edge of A.

Hence, θ is an automorphism of the graph A with $\theta(a) = \theta_0(a) = a'$ and $\theta(a') = \theta_1(a') = a$, therefore A has a reversible edge.

Note that had ii) been true, a symmetric argument gives that B has a reversible edge. Thus, Claim I is true.

Claim II: If B has a reversible edge, then A is edge-transitive.

Let a_0a_1 and $a_0'a_1'$ be arbitrary edges of A; our goal is to construct an automorphism θ of A that sends one to the other. Let bb' be the reversible edge of B, and $\gamma \in \mathrm{Aut}(B)$ the automorphism such that $\gamma(b) = b'$ and $\gamma(b') = b$. Then $(a_0, b)(a_1, b')$ and $(a_0', b)(a_1', b')$ are edges of $A \times B$. Since $A \times B$ is edge-transitive, there is some automorphism φ of $A \times B$ such that either

i) $\varphi(a_0, b) = (a_0', b)$ and $\varphi(a_1, b') = (a_1', b')$, or

ii) $\varphi(a_0, b) = (a_1', b')$ and $\varphi(a_1, b') = (a_0', b)$.

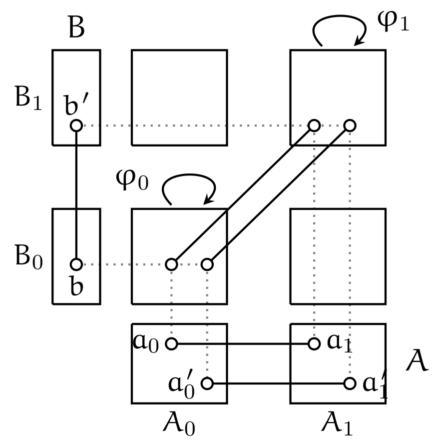

Figure 4.2: The setup for Case 1 Claim II

Case 1: Suppose that i) holds.

By Proposition 4, φ is also an automorphism of $S(A \times B)$. Since i) is true, φ restricts to the automorphisms $\varphi_0 : V(A_0 \square B_0) \to V(A_0 \square B_0)$ and $\varphi_1 : V(A_1 \square B_1) \to V(A_1 \square B_1)$. We use these automorphisms to define an automorphism of A.

First, prime factor $A_0 \square B_0$, say

$$A_0 \square B_0 = \underbrace{F_1 \square \cdots \square F_k}_{A_0} \square \underbrace{F_{k+1} \square \cdots \square F_n}_{B_0}.$$

Then for a given vertex (a, b_0) in $V(A_0 \square B_0)$, apply the stacking operation $M \geqslant n-k$ times with the function φ_0 to get the fully stacked vertex $(a, b_M) \in V(A_0 \square B_0)$. Applying φ_0 once more gives a new vertex in $A_0 \square B_0$ whose A_0 coordinate depends only on a; define $\theta_0 : V(A_0) \to V(A_0)$ by $\varphi_0(a, b_M) = (\theta_0(a), f(a, b_M))$.

Like before, note that in particular $\varphi_0(a_0, b) = (a_0', b)$, so replacing a_0' with a_0 gives back (a_0, b). This means the fully stacked vertex (a_0, b_M) is in fact (a_0, b). Consequently, $(a_0', b) = \varphi_0(a_0, b) = \varphi_0(a_0, b_M) = \big(\theta_0(a_0), f(a_0, b_M)\big)$, so $\theta_0(a_0) = a_0'$.

In the same way, define $\theta_1 : V(A_1) \to V(A_1)$ by stacking vertices in $V(A_1 \square B_1)$ using φ_1. In this case, $M' \geqslant n' - k'$ stacking iterations are sufficient, where we assume the prime factorization

$$A_1 \square B_1 = \underbrace{F'_1 \square \cdots \square F'_{k'}}_{A_1} \square \underbrace{F'_{k'+1} \square \cdots \square F'_{n'}}_{B_1}.$$

And, like with θ_0, we get that $\theta_1(a_1) = a'_1$.

As noted in the introduction to stacking, θ_0 and θ_1 are bijections. Define $\theta : V(A) \to V(A)$ by

$$\theta(a) = \begin{cases} \theta_0(a) & \text{if } a \in V(A_0) \\ \theta_1(a) & \text{if } a \in V(A_1). \end{cases}$$

Certainly θ is a bijection as well, so it remains to show that it preserves adjacency.

To this end, let aa' be an edge of A, and fix an edge xy of B. Then $(a, x)(a', y)$ is an edge of $A \times B$, say with $(a, x) \in V(A_0) \times V(B_0)$ and $(a', y) \in V(A_1) \times V(B_1)$. Stack the endpoints of this edge simultaneously:

0. Given starting edge $(a, x)(a', y) \in E(A \times B)$

1. Apply φ to each end: $\varphi_0(a, x) = (a_1, x_1)$ and $\varphi_1(a', y) = (a'_1, y_1)$, giving $(a_1, x_1)(a'_1, y_1) \in E(A \times B)$

2. Replace a_1 and a'_1: $(a_1, x_1)(a'_1, y_1) \to (a, x_1)(a', y_1) \in E(A \times B)$

In this way, we have done one iteration of stacking on each endpoint of the starting edge, and still have an edge between the resulting vertices. Letting $M'' = \min\{M, M'\}$, then M'' iterations of steps 1 and 2 gives the edge $(a, x_M)(a', y_M)$ of $A \times B$. One final application of φ to each endpoint gives $\big(\theta_0(a), f(a, x)\big)\big(\theta_1(a'), f'(a', y)\big) \in E(A \times B)$, so it follows that $\theta_0(a)\theta_1(a') = \theta(a)\theta(a')$ is an edge of A.

Consequently, θ is an automorphism of A with $\theta(a_0) = \theta_0(a_0) = a'_0$ and $\theta(a_1) = \theta_1(a_1) = a'_1$, so A is edge-transitive.

Case 2: Suppose that ii) holds.

Recall that $\varphi(a_0, b) = (a_1', b')$, $\varphi(a_1, b') = (a_0', b)$, and γ is the automorphism of B that reverses bb'. Define the map $\hat{\gamma} : V(A \times B) \to V(A \times B)$ by $(x, y) \mapsto (x, \gamma(y))$, so that $\hat{\gamma}$ is an automorphism of $A \times B$. Next, define $\hat{\varphi}$ as $\hat{\gamma} \circ \varphi$. Then $\hat{\varphi}$ is the composition of automorphisms of $A \times B$, so is again an automorphism. Note also that

$$\hat{\varphi}(a_0, b) = \hat{\gamma} \circ \varphi(a_0, b) = \hat{\gamma}(a_1', b') = (a_1', b)$$
$$\hat{\varphi}(a_1, b') = \hat{\gamma} \circ \varphi(a_1, b') = \hat{\gamma}(a_0', b) = (a_0', b').$$

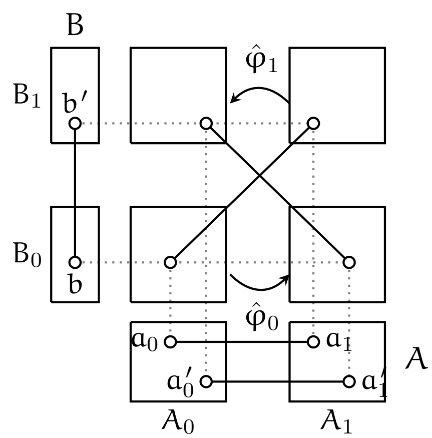

Figure 4.3: The function $\hat{\varphi}$

By Proposition 4, $\hat{\varphi}$ is also an automorphism of $S(A \times B)$. Call the restrictions of $\hat{\varphi}$ the isomorphisms $\hat{\varphi}_0 : V(A_0 \square B_0) \to V(A_1 \square B_0)$ and $\hat{\varphi}_1 : V(A_1 \square B_1) \to V(A_0 \square B_1)$. Like in the proof of Claim I, if $\hat{\varphi}$ is not already an involution, we may redefine it to restrict as $\hat{\varphi}_0^{-1}$ on $V(A_1 \square B_0)$ and $\hat{\varphi}_1^{-1}$ on $V(A_0 \square B_1)$ so that it is. From here, we use $\hat{\varphi}$ to define an automorphism of A.

So let (a, b_0) be an arbitrary vertex of $A_0 \square B_0$. Since $\hat{\varphi}_0$ is also an involution, it at most transposes factors. Thus, after one iteration of stacking on (a, b_0) using $\hat{\varphi}_0$ we get the fully stacked vertex $(a, b_1) \in V(A_0 \square B_0)$. Applying $\hat{\varphi}_0$ once more gives a new vertex in $V(A_1 \square B_0)$ whose A_1-coordinate depends only on the coordinates of a; define

$\theta_0 : V(A_0) \to V(A_1)$ by $\hat{\varphi}_0(a, b_1) = \big(\theta_0(a), f(a, b_1)\big)$, where f is some function from $V(A_0 \square B_0)$ to $V(B)$.

Since $\hat{\varphi}_0(a_0, b) = (a_1', b)$, replacing a_1' with a_0 gives (a_0, b), so again stacking leaves this endpoint of our original edge unchanged. As a result, $(a_1', b) = \hat{\varphi}_0(a_0, b) = \hat{\varphi}_0(a_0, b_1) = \big(\theta(a_0), f(a_0, b_1)\big)$, so $\theta_0(a_0) = a_1'$.

In the same way, stack vertices of $V(A_1 \square B_1)$ using $\hat{\varphi}_1$, and define $\theta_1 : V(A_1) \to V(A_0)$ by $\hat{\varphi}_1(a, b_1) = (\theta_1(a), f'(a, b_1))$. In this case, we get $\theta_1(a_1) = a_0'$.

As reasoned in the stacking introduction, θ_0 and θ_1 are bijections, and we define $\theta : V(A) \to V(A)$ by

$$\theta(a) = \begin{cases} \theta_0(a) & \text{if } a \in A_0 \\ \theta_1(a) & \text{if } a \in A_1, \end{cases}$$

so θ is a bijection as well. Verifying that θ preserves adjacency here is the same as at the end of the proof of Claim I, so θ is an automorphism of A sending $a_0 a_1$ to $a_1' a_0'$, thus A is edge-transitive.

If A has a reversible edge, a symmetric argument shows that B is edge-transitive.

Summary

By Claim I, at least one of A, B must have a reversible edge, so Claim II then implies that at least one of A, B is edge transitive, proving statement 1 of the Theorem. In the case that $A \times B$ is arc-transitive, both statements i) and ii) in the proof of Claim I are true, so both A and B have reversible edges. Claim II then implies A and B are both edge-transitive with reversible edges, hence both are arc-transitive by Proposition 1. This proves statement 2. $\square$

4.3 Remarks

Ideally, both factors could be shown to be edge-transitive given an edge-transitive product. If we were to assume that, say, B were not edge-transitive, the contrapositive of Claim I tells us A has no reversible edges. So if we fix an edge aa' of A, the argument in Claim I could be applied to aa' and *any* edge bb' of B. Since aa' can't be reversible, it must be that the arbitrary edge bb' of B *is*. Thus, every edge of B is reversible.

Further, B is connected, so for any pair of vertices b, b' of B, there is a path P connecting them. Say P is the sequence of edges $e_1, e_2, \ldots, e_k$. Then every edge e_i is reversible by some automorphism α_i, so the composition $\alpha_1 \circ \alpha_2 \circ \cdots \circ \alpha_k$ is an automorphism that relates b and b', so B is additionally vertex transitive.

So while it is not yet necessary that B is edge-transitive, this evidence strongly suggests it is. If that were the case, it would also be arc-transitive by Proposition 1. This suggests the following conjecture:

Conjecture 1. *Let A and B be connected, bipartite, R-thin graphs. Then $A \times B$ is edge-transitive if and only if A and B are both edge-transitive, and at least one is arc-transitive.*

Bibliography

[1] *A First Course in Graph Theory,* by G.Chartrand and P. Zhang, Dover Publications, Inc., Minneola, NY, 2012.

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

[3] R. Hammack, W. Imrich and S. Klavžar, Edge-transitive products, *Journal of Algebraic Combinatorics*, **43**(4) 837–850, 2016.

[4] *Handbook of Product Graphs,* second edition, by R. Hammack, W. Imrich and S. Klavžar, CRC Press, Boca Raton, FLA, 2011.