ON THE WEAK RECONSTRUCTION OF CARTESIAN-PRODUCT GRAPHS

Wilfried Imrich
Institut für Mathematik und Angewandte Geometrie
Montanuniversität Leoben, A-8700 Leoben
Austria

Janez Žerovnik [1]
University of Maribor, Faculty of Mechanical Engineering
Smetanova 17, 62000 Maribor
Slovenia

AMS Subject Classification (1991): 05C
Keywords: reconstruction problem, Cartesian product, composite graphs.

Abstract

In this paper we reconstruct nontrivial connected Cartesian product graphs from single vertex deleted subgraphs. We show that all one-vertex extensions of a given connected graph H, finite or infinite, to a nontrivial Cartesian product are isomorphic.

---

[1]And Inštitut za matematiko, fiziko in mehaniko, Jadranska 19, Ljubljana. Part of this work was done while this author was visiting the Institut für Mathematik und Angewandte Geometrie at the Montanuniversität Leoben, Austria. Support by the Research Council of Slovenia and the Austrian Ministry of Science and Research is gratefully acknowledged.

1 Introduction

In [7] S. M. Ulam asked the question whether a graph G is uniquely determined up to isomorphism by its deck, which is the set of all graphs $G \setminus x$ obtained from G by deleting a vertex x and all edges incident to it. While the conjecture is false for infinite graphs it still is open for finite graphs. When reconstructing a class of graphs, the problem of reconstruction partitions naturally into two subproblems, namely recognition: showing that membership in the class is determined by the deck and weak reconstruction: showing that no two nonisomorphic members of the class have the same deck. Many partial results have been found. For example, Dörfler [2] proved the validity of Ulam's conjecture for finite nontrivial Cartesian product graphs, i.e. graphs which are the Cartesian product of at least two nontrivial factors. In this paper we extend the work of Dörfler [2] by showing that both the recognition and the weak reconstruction problem can be solved from a single vertex–deleted subgraph for nontrivial, connected Cartesian product graphs.

We consider extensions of finite and infinite connected graphs to Cartesian products. In most cases, it will not be possible to extend a given connected graph H to a nontrivial Cartesian product. However, if such extensions exist, they are all isomorphic (Theorem 1). In fact, unless H has a special structure, there is exactly one such extension.

2 Definitions

All graphs considered in this paper are finite or infinite undirected graphs without loops or multiple edges. If G is a graph, we shall write $V(G)$ or V for its vertex set and $E(G)$ or E for its edge set. $E(G)$ shall be considered as a set of unordered pairs $\{x, y\}$ of distinct vertices of G.

Considering G as $V(G) \cup E(G)$, we shall often write $x \in G$ for $x \in V(G)$ and $e \in G$ for $e \in E(G)$.

Let G_ι, $\iota \in I$, be a set of graphs. Then the *Cartesian product* $G = \Box_{\iota \in I} G_\iota$ is defined as follows:

(i) $V(G)$ is the Cartesian product of the vertex sets of the factors. In other words, $V(G)$ is the set of functions $x : \iota \mapsto x_\iota \in V(G_\iota)$ of I into $\cup_{i \in I} V(G_\iota)$.

(ii) $E(G)$ consists of all unordered pairs $\{x, y\}$ of distinct vertices of G for which there exists a $\kappa \in I$ such that $\{x_\kappa, y_\kappa\} \in E(G_\kappa)$ and $x_\iota = y_\iota$ for $\iota \in I \setminus \{\kappa\}$.

For two factors G, H we obtain the usual Cartesian product $G \Box H$. It is commutative and associative in an obvious way, having the trivial graph as a unit.

Common examples of Cartesian products are squares, the skeletons of cubes and n–cubes, prisms (Cartesian products of n–gons by an edge) or the square lattice as the product of two infinite paths.

The product of finitely many graphs is connected if and only if every factor is. However, a product of infinitely many nontrivial graphs must be disconnected because it contains vertices differing in infinitely many coordinates. No two such vertices can be connected by a path of finite length, because every edge connects vertices differing in exactly one coordinate.

This gives rise to the notion of the so-called weak Cartesian product:

Let G_ι, $\iota \in I$ be a set of connected graphs and $a \in V(\Box_{\iota \in I} G_\iota)$. Then the *weak Cartesian product*

$$G = \Box^a_{\iota \in I} G_\iota$$

is the connected component of $G = \Box_{\iota \in I} G_\iota$ containing a. We note that $\Box^a G_\iota = \Box^b G_\iota$ if and only if a and b differ in at most finitely many coordinates.

The weak infinite dimensional cube whose vertices are all 0–1 sequences with only finitely many ones and where two vertices are connected by an edge if they differ in exactly one coordinate is a weak Cartesian product of countably many edges.

Now we define two relations σ and δ on $E(G)$ and state several known properties of σ and δ.

Definition Let $G = \Box^a_{\iota \in I} G_\iota$. Call $e = \{u, v\}$ a G_κ–edge if e connects vertices u and v such that $u_\iota = v_\iota$ for all $\iota \neq \kappa$ and $\{v_\kappa, u_\kappa\} \in E(G_\kappa)$. We then say that two edges e, f are in the relation $\sigma(\Box^a_{\iota \in I} G_\iota)$ if there is an ι such that e and f are G_ι edges.

It is easy to see that $\sigma(\Box^a_{\iota \in I} G_\iota)$ is an equivalence relation. It is known that among all relations induced by a Cartesian product representations of G there exists a finest one [6, 5]. This relation is unique and will be denoted by σ_G or simply σ. It determines the so-called prime factor decomposition in an obvious way. Since σ is unique the prime factor decomposition is unique up to isomorphisms and the order of factors. A graph G for which σ_G has only one equivalence class is called *prime*. Otherwise we call the graph *composite*. In this paper, by a Cartesian product graph we will always mean a composite graph. Let G be arbitrary graph and assume we have a decomposition of G, $G = \Box G_\iota$. Then $G^a_\kappa = \{v \mid v_\iota = a_\iota, \iota \neq \kappa\}$ is the *κ–layer* through the vertex $a \in G$.

Definition Let $e, f \in E(G)$. We say e and f are in the relation δ if one of the following conditions is satisfied

(1) e and f are the opposite edges of a chordless square.

(2) e and f are adjacent and there is no chordless square spanned by e and f.

(3) e=f.

Clearly δ is reflexive and symmetric. Hence, its transitive closure $\delta^\star$ is an equivalence relation.

Sometimes we shall write $\delta = \delta_G$ to avoid confusion if there is more than one graph in consideration.

By the definition of δ any pair of adjacent edges which belong to distinct $\delta^\star$ equivalence classes span a square. We say that the relation $\delta^\star$ has the *square property*. It is easy to see that the square property also holds for any equivalence relation containing $\delta^\star$.

In the algorithm developed in [3] for factoring a graph into prime graphs with respect to Cartesian multiplication one starts with the relation $\delta^\star$ and then unifies equivalence classes until the new relation is the Cartesian product relation σ. Hence, $\delta^\star \subseteq \sigma$ and equivalence classes of σ are unions of equivalence classes of $\delta^\star$.

We further note that each vertex in a connected graph G is incident to at least one edge of each class in $\delta^\star$ (Lemma 1 of [3]).

A *product square* in a Cartesian product graph is a square for which the two pairs of opposite edges belong to distinct equivalence classes of the relation σ. Sometimes we say that edges of the same σ class are colored with the same color. Hence, the edges of a product square are alternately colored by two colors.

Let H be a subgraph of G. Then H is *convex* (in G) if all shortest G–paths between two vertices of H are already in H. It is easy to see that X is convex in Z if X is convex in Y and Y is convex in Z.

The complete graph on n vertices is denoted by K_n, the path of length n (with n edges) by P_n and the cycle on n vertices by C_n.

We shall also need the concept of star graphs S_a. They are defined as follows: The vertex set of S_a consists of a central vertex c_0 of degree a and of a vertices $c_1, c_2, \ldots, c_a$ adjacent to a.

$G \setminus x$ denotes the subgraph of G induced by the vertex set $V(G) \setminus \{x\}$ and $\simeq$ denotes graph isomorphism, i.e. $G_1 \simeq G_2$ means that G_1 and G_2 are isomorphic.

3 Uniqueness of Reconstruction

Let G be a finite or infinite connected Cartesian product graph and let x be any vertex of G. We shall prove that, given the graph $G \setminus x$, it is possible to reconstruct G uniquely up to isomorphism.

Theorem 1 *Let G_1 and G_2 be finite or infinite connected Cartesian product graphs. If the one vertex deleted subgraphs $G_1 \setminus x$ and $G_2 \setminus y$, where $x \in G_1$ and $y \in G_2$, are isomorphic, then $G_1 \simeq G_2$.*

In other words, if H is an arbitrary finite or infinite connected graph and if $G_1 = H \cup x \cup E_x$ and $G_2 = H \cup y \cup E_y$ are one–vertex extensions of H, such that $E_x = \{\{x, z\} \mid \{x, z\} \in E(G_1)\}$ and $E_y = \{\{y, z\} \mid \{y, z\} \in E(G_2)\}$, then G_1 and G_2 are isomorphic, provided they are Cartesian product graphs.

Note that $G \setminus x$ is connected if G is a connected Cartesian product graph.

We say the extension is *unique*, if x and y have the same neighbors in H. More formally, let $N(x) = \{z \mid \{z, x\} \in E_x\}$ and $N(y) = \{z \mid \{z, y\} \in E_y\}$ be the neighborhoods of the new vertices x and y, respectively. Then the extension is unique, if $N(x) = N(y)$.

For the proof we shall consider two cases. In the first H contains no product square of G_1 or G_2. We shall see that this is only possible if G_1 (and G_2) is a product of two stars. Then the extension is unique unless $H = C_8$ (and $G_1 \simeq G_2 \simeq S_2 \Box S_2$). From the assumption that H has no product square, it will also follow that there is no square in H. Thus, if there exist squares in H, at least one of them must be a product square. This will be the second case. In this case the reconstruction is unique if G has more than two (nontrivial) factors or no K_2 factor. If $G = K_2 \Box P$, where P is a prime graph, then the reconstruction need not be unique.

3.1 Case 1, product of stars

Proposition 1 *Assume G is a connected Cartesian product graph and $G \setminus x$ contains no product square of G. Then there are $a, b \in \mathbf{N}$ such that*

$$G = S_a \Box S_b.$$

Furthermore, if $a \neq 2$ or $b \neq 2$ then the extension is unique. If $a = b = 2$ then there are two isomorphic extensions.

Proof Assume $G = H_1 \Box H_2$ is any factoring of G and $G \setminus x$ contains no product square.

Let $x = (x_1, x_2)$. If there is an edge in $H_1 \setminus x_1$ then $(H_1 \setminus x_1) \Box H_2 \subset G \setminus x$ and, since H_2 is nontrivial and connected, there are product squares in $G \setminus x$. Hence, there can be no edge in $H_1 \setminus x_1$. Since H_1 is connected, H_1 must be a star and x_1 must be the central vertex of H_1. By the same arguments, H_2 is a star with central vertex x_2.

Now, let G be a product of two stars S_a, S_b and let $x = (x_1, x_2)$, where x_1 and x_2 are the central vertices of S_a and S_b, respectively.

Then the vertices of $G = S_a \Box S_b$ have the following degree sequence:

degree	a+b	a+1	b+1	2
no. of vertices	1	b	a	a.b

After deleting the central vertex, we have the degree sequence:

degree	a	b	2
no. of vertices	b	a	a.b

If the degree sequence of $G \setminus x$, where G contains no product square, agrees with the sequence in the table for some cardinals a and b, then G can be easily reconstructed. If not, then G cannot be a Cartesian product graph.

Note that if $a \neq 2$ or $b \neq 2$ the reconstruction is unique. The only ambiguity occurs when $a = b = 2$. In this case $G \setminus x$ is a cycle on eight vertices and there are two (isomorphic) ways of obtaining the extension $S_2 \Box S_2$. $\Box$

3.2 Case 2

In this section we assume there is a product square in $G \setminus x$. We prove that this implies that the reconstruction is unique up to isomorphism and characterize the graphs for which several isomorphic reconstructions are possible.

Let $\mathcal{S}$ be the set of convex subgraphs of $G \setminus x$ which are Cartesian product graphs. $\mathcal{S}$ is nonempty since the product square that exists by assumption is convex. The set $\mathcal{S}$ is partially ordered by inclusion. Below we shall characterize the elements of $\mathcal{S}$ and all maximal elements of $\mathcal{S}$.

Proposition 2 *Let $G = \Box^x_{\iota \in I} G_\iota$ be a prime factor decomposition of the finite or infinite connected graph G with respect to the (weak) Cartesian product. Then the elements of $\mathcal{S}$ are of the form $H = \Box^y_{\iota \in I} H_\iota$, where y is a vertex of finite distance from x in G and H_ι is convex in G_ι for all ι. Moreover, the maximal elements of $\mathcal{S}$ are of the form $H = \Box^y_{\iota \in I} H_\iota$, where y is a neighbor of x in G and there exists an index κ such that H_κ is a maximal subgraph of $G_\kappa \setminus x_\kappa$ which is convex in G_κ and $H_\iota = G_\iota$ for $\iota \neq \kappa$.*

For the proof of the Proposition we need a Lemma.

Lemma 1 *Let H be any convex subgraph of $G \setminus x$. Let $G = G_1 \Box G_2$ be any factorization of G and let $H_1 \subseteq G_1$ and $H_2 \subseteq G_2$ be the two projections of the graph H into the factors of G. Then either $H = H_1 \Box H_2$ or $x \in H$ and H is prime.*

Proof If either of the projections H_1 or H_2 has only one vertex, the assertion is trivially true.

Hence, we shall assume that both projections have at least two vertices.

Claim $H_1 \Box H_2 = H$ if $x \notin H_1 \Box H_2$.

Proof Let $y = (y_1, y_2) \in (H_1 \Box H_2) \setminus H$ be arbitrarily chosen. We have to prove that $y \in H$ if $x \notin H_1 \Box H_2$.

For any vertex a, let H_i^a be the H_i–layer of $H_1 \Box H_2$ containing the vertex a. Let $z = (z_1, y_2) \in H$ be a vertex of H of minimal G–distance from y in the layer H_1^y, in symbols,

$$d_G(z, y) = \min_{u \in H, u_2 = y_2} d_G(u, y).$$

Furthermore, let $w = (y_1, w_2) \in H$ be a vertex of H of minimal G–distance from y in the layer H_2^y, in symbols,

$$d_G(w, y) = \min_{u \in H, u_1 = y_1} d_G(u, y).$$

Since, by definition of z and w,

$$d_G(z, w) = d_G(z, y) + d_G(y, w),$$

y is on a shortest path in G between two vertices of H. If $d_G(z, y) = d_{G \setminus x}(z, y)$ and $d_G(w, y) = d_{G \setminus x}(w, y)$, then $y \in H$ unless x is either on all shortest paths from z to y in G or on all shortest paths from w to y. This implies that for any $y \in H_1 \Box H_2 \setminus H$ either $y_1 = x_1 \in H_1$ and $x_2 \in H_2^y$ or $x_1 \in H_1^y$ and $y_2 = x_2 \in H_2$. In both cases we would have $x \in H_1 \Box H_2$, hence $y \in H$. $\Box$

Thus, we can assume x lies in the product $H_1 \Box H_2$. We shall prove that in this case H cannot be a Cartesian product graph.

First we show that there is at least one vertex $c = (c_1, c_2)$ in H, such that $c_1 = x_1$ and c_2 is adjacent to x_2 in H_2. Since $x_1 \in H_1$, there must be at least one vertex, say $v = (v_1, v_2)$, in H with $v_1 = x_1$. If v is of distance 1 from x in G then we set $c = v$. If v is of distance $d > 1$ from x in G, then we construct a vertex of distance $d - 1$ from x in H_2^x as follows. Since H_1 has at least two vertices, there is a vertex $y = (y_1, y_2)$ where y_1 is adjacent to x_1 in H_1. There is a shortest path from y to v that meets (y_1, v_2). Let $w = (w_1, w_2)$ with $w_1 = y_1$ and $w_2 = v_2$. Then w is on a shortest path from y to v and thus in H by convexity. Let u_2 be a neighbor of w_2 in H_2 on any shortest path from w_2 to x_2 in H_2. Again $u = (u_1, u_2) \in H$ because $u \in H_1 \Box H_2$, $u_1 = w_1 \neq x_1$ and $u_2 \neq x_2$. As $u, v, w \in H$ convexity of H implies $z = (x_1, u_2) \in H$. So we have found a vertex $z \in H$ of distance $d - 1$ from x with $z_1 = x_1$. Repeating this construction we eventually get a vertex of distance one from x as needed.

By symmetry there exists a vertex $b = (b_1, b_2)$ with $b_2 = x_2$ and b_1 adjacent to x_1 in H_1.

Since the edges $\{x, b\}$ and $\{x, c\}$ are in distinct σ equivalence classes in G, there is a unique vertex $a \neq x$ in G with coordinates $a = (a_1, a_2) = (b_1, c_2)$. By convexity of H in $G \setminus x$, a must be in H.

In $G \setminus x$, and therefore also in H, there is no square spanned by the vertices $\{a, b, c\}$. Hence, the edges $\{a, b\}$ and $\{a, c\}$ are in relation δ_H. Since $\delta_H \subseteq \sigma_H$, $\{a, b\}$ and $\{a, c\}$ are in the same equivalence class, say E_1, of σ_H.

Assume that H is a Cartesian product graph. We will now show that this leads to a contradiction. In a Cartesian product, edges of every equivalence class of σ meet every vertex. Hence, there must be a neighbor of a, say d, such that the edge $\{a, d\}$ is not in the same σ_H equivalence class as the two edges $\{a, b\}$ and $\{a, c\}$. Say $\{a, d\} \in E_2 \neq E_1$. Since H is a subgraph of $G \setminus x$, the coordinates of $d = (d_1, d_2)$ in G satisfy either $d_1 = a_1$ and d_2 is adjacent to a_2 in H_2 or $d_2 = a_2$ and d_1 is adjacent to a_1 in H_1.

Without loss of generality we assume that $d_1 = a_1$ and d_2 is adjacent to a_2. By the square property, there must be two vertices, say f and e, which 'close' the two squares over $\{a, b, d\}$ and $\{a, c, d\}$. We shall thus associate f with $\{a, b, d\}$ and e with $\{a, c, d\}$ in the sequel. Since H is a subgraph of $G \setminus x$, the coordinates of e and f in G must be $e = (e_1, e_2)$, where $e_1 = c_1 = x_1$ and $e_2 = d_2$, and $f = (f_1, f_2)$, where $f_1 = a_1 = b_1 = d_1$ and f_2 is adjacent to both d_2 and b_2 in H_2. By construction, the edges $\{a, b\}$, $\{a, c\}$, $\{d, e\}$ and $\{d, f\}$ must be in the equivalence class E_1 and the edges $\{a, d\}$, $\{c, e\}$ and $\{b, f\}$ must be in equivalence class E_2. Call g the vertex (e_1, f_2) in $G \setminus x$. Since $e - g - f$ is a shortest path in $G \setminus x$ between e and f and H is convex, $g \in H$. By definition of δ the edges of the square $\{d, e, g, f\}$ are all in the class E_1. Since the edges $\{b, f\}$ and $\{f, g\}$ are in distinct σ_H equivalence classes, there must be a second common neighbor of b and g in H and thus G. In G, the unique second common neighbor of these two vertices is x, hence there can be no second common neighbor of b and g in H, since H is a subgraph of $G \setminus x$.

From this contradiction we conclude that H must be prime if $x \in H_1 \Box H_2$. $\Box$

Remark The convexity of Cartesian product subgraphs is essential for Lemma 1. There are examples of nonconvex Cartesian product subgraphs for which the lemma does not hold. For the smallest example we know see Fig.1. Surprisingly, there even are examples where such a nonconvex Cartesian product has more vertices than any convex Cartesian product subgraph (see Fig.2).

For $H = G \setminus x$ the lemma proves $G \setminus x$ is prime if G is a Cartesian product graph. Clearly, $G \setminus x$ is a convex subgraph of itself, the projections of $G \setminus x$ are G_1 and G_2, and $G \setminus x \neq G$. Hence, $G \setminus x$ must be prime. We formulate this as a proposition before continuing with the proof of Proposition 2.

Proposition 3 *Let G be a Cartesian product graph and $x \in G$. Then $G \setminus x$ is prime.*

Let $H = H_1 \Box H_2$ be convex subgraph of $G_1 \Box G_2$. Then it is easy to see that H_1 must be convex in G_1 and H_2 must be convex in G_2. On the other hand, if H_1 and H_2 are convex in G_1 and G_2, respectively, this implies that $H = H_1 \Box H_2$ must be convex in $G_1 \Box G_2$. To show this, let P be any shortest path between a pair of vertices of $H = H_1 \Box H_2$. Consider the projections P_1 and P_2 of P into G_1 and G_2. By convexity of H_1 and H_2, $P_1 \subseteq H_1$ and $P_2 \subseteq H_2$, hence $P = P_1 \Box P_2 \subseteq H_1 \Box H_2$. We formulate this observation as a Lemma.

Lemma 2 *$H = H_1 \Box H_2$ is convex in $G_1 \Box G_2$ if and only if H_1 is convex in G_1 and H_2 is convex in G_2.*

Proof of Proposition 2 Knowing that the elements of $\mathcal{S}$ must be of the form $H = \Box_{\iota \in I} H_\iota$, where $H_\iota \subseteq G_\iota$ and $x_\kappa \notin H_\kappa$ for some index κ, it is easy to characterize the maximal elements with respect to set inclusion. Now consider subgraphs of $G_\kappa \setminus x_\kappa$ containing H_κ and convex in G_κ. Let B be a maximal such graph. Since $H = \Box_{\iota \in I} H_\iota \in \mathcal{S}$, $B \Box (\Box_{\iota \in I, \iota \neq \kappa} G_\iota)$ is also in $\mathcal{S}$ and is clearly maximal. □

Lemma 3 *Any maximal convex Cartesian product subgraph of $G \setminus x$ can be uniquely extended to a maximal connected Cartesian product subgraph G^* of $G \setminus x$. It is of the form $C \Box (\Box_{\iota \in I, \iota \neq \kappa} G_\iota)$, where C is a connected component of $G_\kappa \setminus x_\kappa$.*

Proof Let H be a maximal convex Cartesian product subgraph of $G \setminus x$. By Proposition 2, $H = B \Box (\Box_{\iota \in I, \iota \neq \kappa} G_\iota)$, where B is a maximal subgraph of

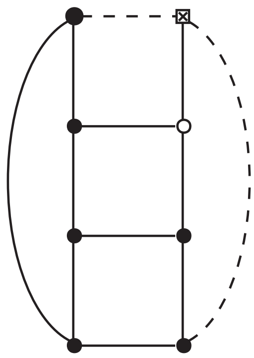

Figure 1: Example of a nonconvex Cartesian product subgraph.

$G_\kappa \setminus x_\kappa$ which is convex in G_κ. Let C be the connected component of $G_\kappa \setminus x_\kappa$ containing B. Clearly, $B\Box(\Box_{\iota\in I,\iota\neq\kappa}G_\iota) \subseteq C\Box(\Box_{\iota\in I,\iota\neq\kappa}G_\iota)$. $\Box$

A construction of G^* from $B\Box(\Box_{\iota\in I,\iota\neq\kappa}G_\iota)$ is not difficult. Note that any vertex v of distance 1 from $B\Box(\Box_{\iota\in I,\iota\neq\kappa}G_\iota)$ can only be connected to $B\Box(\Box_{\iota\in I,\iota\neq\kappa}G_\iota)$ by an G_κ-edge. In G, this gives rise (by repeated application of the square property) to a copy of $\Box_{\iota\in I,\iota\neq\kappa}G_\iota$, namely $\{v_\kappa\}\Box(\Box_{\iota\in I,\iota\neq\kappa}G_\iota)$. In $G \setminus x$, there are two possible cases. If $x_\kappa = v_\kappa$ then $\{v_\kappa\}\Box(\Box_{\iota\in I,\iota\neq\kappa}G_\iota)$ is not completely in $G \setminus x$. However, this is not possible because G^* would not be prime by Proposition 3. Otherwise, (if $x_\kappa \neq v_\kappa$) the whole copy $\{v_\kappa\}\Box(\Box_{\iota\in I,\iota\neq\kappa}G_\iota)$ is in $G \setminus x$. We can extend the graph $B\Box(\Box_{\iota\in I,\iota\neq\kappa}G_\iota)$ as follows. Take a vertex of distance 1. Using the square property, add new vertices. If the resulting graph is not prime, let $B = B\cup\{v_\kappa\}$, otherwise reject the new vertices. Correctness of this decision follows from Proposition 3(The implementation of the construction may label rejected edges so that each edge is considered at most once by the algorithm.)

Note that any connected component of $G_1 \setminus x_1$ must contain at least one neighbor of x_1.

The next step in the proof of the uniqueness of the reconstruction is the following:

Lemma 4 *Let $G = G_1\Box G_2$ be any factorization of G and let C be a connected component of $G_1 \setminus x_1$. Let $G^* = C\Box G_2$. Then the set of vertices of*

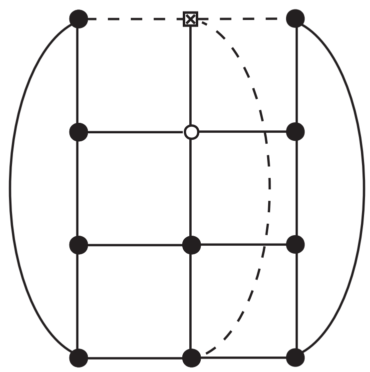

Figure 2: The maximal Cartesian product subgraph is not convex

distance 1 from G^* *in* $G \setminus x$ *induces the graph* $(\{x_1\}\Box G_2) \setminus x$.

Proof Let v be any vertex of $G \setminus x$ of distance 1 from G^*. Then there is a vertex $u = (u_1, u_2) \in G^*$ such that $v_2 = u_2$ and $\{v_1, u_1\}$ is an edge of G_1. Since C is a connected component of $G_1 \setminus x_1$, the only edges possible are of the form $\{x_1, u_1\}$, i.e. $v_1 = x_1$, and $v \in (\{x_1\}\Box G_2) \setminus x$. On the other hand, any vertex of $G \setminus x$ with coordinates $z = (x_1, z_2)$ is of distance 1 from the vertex (u_1, v_1). $\Box$

Let G be a Cartesian product graph with prime factor decomposition $G = \Box_{\iota \in I} P_\iota$. Then a subgraph G^* of Lemma 3 will be referred to in the sequel as a canonical maximal product subgraph of $G \setminus x$. It is important to note that it can be constructed solely from the knowledge of $G \setminus x$ by first choosing a maximal convex product subgraph of $G \setminus x$ and then extending it to a maximal connected product subgraph of $G \setminus x$. By Lemma 3 such a canonical maximal product subgraph is of the form $C \Box G_2$, where C is a connected component of $G_\kappa \setminus x_\kappa$ for some $\kappa \in I$ and $G_2 = \Box_{\iota \in I \setminus \kappa} P_\iota$.

Noting that C need not be prime we observe that the prime factor decomposition of G^* contains all prime factors of G except one and that the product of the other prime factors of G^* is C, i.e. a subgraph of the missing prime factor of G. Our problem reduces to correctly sorting out the prime factors of G and to extend C.

We distinguish two cases:

- There is a canonical maximal Cartesian product subgraph $G^* = C \Box G_2$ in $G \setminus x$ such that G_2 has more than 2 vertices.

  In this case the number of vertices of distance 1 from G^* is at least 2.

- Any canonical maximal Cartesian product subgraph in $G \setminus x$ is of the form $G^* = C \Box K_2$.

  In this case there is only one vertex of distance 1 from G^* in $G \setminus x$.

Lemma 5 *Assume there is a canonical maximal Cartesian product subgraph G^* in $G \setminus x$ such that the number of vertices of distance 1 from G^* is at least 2. Then the reconstruction of G is unique.*

Proof Let $G = \Box_{\iota \in I} P_\iota$ be the prime factor decomposition (PFD) of G. We know that G^* is of the form $H_\kappa \Box (\Box_{\iota \neq \kappa} P_\iota)$ where H_κ is a connected component of $P_\kappa \setminus x_\kappa$.

Set $G_1 = P_\kappa$, $C = H_\kappa$ and $G_2 = \Box_{\iota \neq \kappa} P_\iota$. Then $G = G_1 \Box G_2$ and $G^* = C \Box G_2$. We shall thus write $x = (x_1, x_2)$, where $x_1 = p_{G_1}(x)$ and $x_2 = p_{G_2}(x)$. Clearly $x_1 = x_\kappa$. By Lemma 4 $(\{x_1\} \Box G_2) \setminus x$ is the set of neighbors of G^* in $G \setminus x$ and hence known.

For a vertex $u \in (\{x_1\} \Box G_2) \setminus x$ let $N(u)$ be the set of neighbors of u in $C \Box G_2$. Furthermore, let M be the union of all $N(v)$, $v \in \{x_1\} \Box G_2$. Clearly M forms a subproduct of G. We claim that this subproduct is of the form $N \Box (G_2 \setminus x_2)$, where $N \simeq N(v)$ for all v and that there is only one way to extend this subgraph to $N \Box G_2$. More precisely, there is exactly one possible subset N in C such that $N \Box G_2$ is a Cartesian product subgraph of $C \Box G_2$ induced exactly by the union of neighborhoods $N(v)$ for all $v \in \{x_1\} \Box G_2$. The extension is determined by the set of vertices $N(x)$, i.e. the vertices in G^* which have to be connected to the new vertex x. We have to show that $N(x)$ is uniquely determined. In order to do this we should keep in mind that $C = H_\kappa$ need not be prime. Thus, let $H_\kappa = \Box_{\mu \in M} Q_\mu$ be the PFD of H_κ. Since the P_ι are all prime we thus obtain the PFD

$$(\Box_{\mu \in M} Q_\mu) \Box (\Box_{\iota \in I, \iota \neq \kappa} P_\iota)$$

of G^*. The coloring of the edges of G^* with respect to this decomposition is thus a refinement of the coloring induced by the PFD of $G = \Box_{\iota \in I} P_\iota$. Of course, this latter coloring is yet unknown to us.

Since there are at least two sets of neighbors, say $N(u)$ and $N(v)$, and since the graph $C \Box G_2$ is connected, the shortest paths from $N(u)$ to $N(v)$ (in general, between vertices of different neighborhood sets $N(w), w \in \{x_1\} \Box G_2$) are exactly of the colors with which the G_2-layers in $C \Box G_2$ are colored. If G_2 is not prime, then there is more than one 'G_2–color' in the prime factor decomposition of the maximal convex Cartesian product decomposition of $C \Box G_2$, but clearly none of these colors can appear in any of the sugraphs induced by the $N(v)$'s. The G_2–colors determine the only possible subset $N(x)$ needed as follows: $N(x)$ is the set of vertices in $C \Box G_2$, which are connected to some $N(v)$ by G_2–colored edges.

At this stage of the proof we have thus concretely determined G_2 and all $P_\iota, \iota \neq \kappa$, as well as C and N. Thus we also know how x is connected with G^*. We still have to determine P_κ. If $P_\kappa \setminus \{x_1\}$ has only one connected component, then we must already have $(\{x_1 \cup C\} \Box G_2) \setminus x \simeq G^*$ and P_κ is the constructed graph $\{x_1\} \cup C$. Otherwise, there must be another connected component C_2 of $P_\kappa \setminus \{x_1\}$ giving rise to another maximal Cartesian product graph G_2^* of $G \setminus x$, disjoint from G^*. By the same reasoning as above we can determine C_2 as well as all other connected components of $P_\kappa \setminus \{x_1\}$ and hence P_κ. $\Box$

Lemma 6 *Assume there is a canonical maximal Cartesian product subgraph G^* in $G \setminus x$ such that the number of vertices of distance 1 from G^* is 1. Then all the possible reconstructions are isomorphic.*

Proof Now there is only one vertex v of distance one from $G^* = C \Box G_2$ in $G \setminus x$. Let $G^* = \Box H_i$ be the prime factoring of G^*. There is at least one K_2 factor in G^*, but there may be more. Now we can take any K_2 factor of G^*, such that the projection of the neighborhood $N(v)$ on that K_2 factor has only one vertex. Call such a factor *free*. If there are more such K_2 factors in G^* they define different sets of neighbors of the vertex x in G^*. (Note that now there is no shortest path between different neighbor sets $N(v)$ and thus no color, i.e direction in $C \Box G_2$, determined by $\{x_1\} \Box G_2 \setminus x$.) But in all cases the resulting graph is isomorphic to $K_2 \Box (C \cup x)$.

More formally, G^* can be written as a product $K_2^s \Box R$, where $N(v) \subset R^v$ and such that s, the number of free K_2 factors, is maximal. Note that K_2^s is the product of s edges, i.e. an n-dimensional hypercube, and that R^v, the

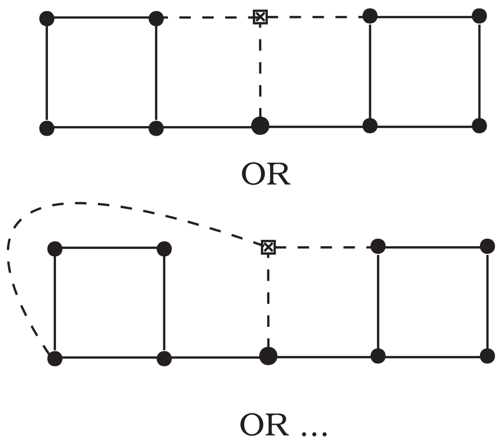

Figure 3: Nonunique extensions, example $G = K_2 \Box P_4$.

R-layer through v, is the induced subgraph of G^* spanned by those vertices which differ from v only in the R-coordinate. Then the vertex x can be connected to G^* in s different, but isomorphic ways. An example is given in Fig. 3.

If there are more connected components of $G_1 \setminus x_1$, then for each connected component D the reconstruction is essentially as before. Again we have a subgraph $D \Box K_2$, and a set of neighbors $N(v)$ of v in $D \Box K_2$. Depending on the number of free K_2 factors, there may be more isomorphic ways of connecting the vertex x to the subgraph $D \Box K_2$. $\Box$

Acknowledgement We wish to thank Johann Hagauer for many fruitful discussions.

References

[1] F. AURENHAMMER, J. HAGAUER, W. IMRICH: Cartesian graph factorization at logarithmic cost per edge, Comput. Complexity **2** (1992) 331-349.

[2] W. DÖRFLER: Some Results on the Reconstruction of Graphs, Colloq. Math. Soc. János Bolyai, 10, Keszthely, Hungary, 1973, 361-383.

[3] J. FEIGENBAUM, J. HERSCHBERGER, A. A. SCHÄFFER: A Polynomial Time Algorithm for Finding the Prime Factors of Cartesian–Product Graphs, Discrete Appl. Math. **12** (1985) 123-138.

[4] J. FEIGENBAUM, R. HADDAD: On Factorable Extensions and Subgraphs of Prime Graphs, SIAM J. Disc. Math **2** (1989) 197-218.

[5] W. IMRICH: Embedding Graphs into Cartesian Products, Graph Theory and Applications: East and West, Annals of the New York Academy of Sciences, vol.**576**, (1989) 266-274.

[6] G. SABIDUSSI: Graph Multiplication, Math. Z. **72** (1960) 446-457.

[7] S. M. ULAM: A Collection of Mathematical Problems, Wiley, New York, 1960, p.29.