The Cartesian Product of Hypergraphs

Lydia Gringmann
BIOINFORMATICS GROUP, DEPARTMENT OF COMPUTER SCIENCE AND INTERDISCIPLINARY CENTER FOR BIOINFORMATICS UNIVERSITY OF LEIPZIG, HÄRTELSTRASSE 16-18, D-04107 LEIPZIG, GERMANY
E-MAIL: GLYDIA@BIOINF.UNI-LEIPZIG.DE

Marc Hellmuth
MAX PLANCK INSTITUTE FOR MATHEMATICS IN THE SCIENCES INSELSTRASSE 22, D-04103 LEIPZIG, GERMANY

BIOINFORMATICS GROUP, DEPARTMENT OF COMPUTER SCIENCE AND INTERDISCIPLINARY CENTER FOR BIOINFORMATICS UNIVERSITY OF LEIPZIG, HÄRTELSTRASSE 16-18, D-04107 LEIPZIG, GERMANY
E-MAIL: MARC@BIOINF.UNI-LEIPZIG.DE

Peter F. Stadler
BIOINFORMATICS GROUP, DEPARTMENT OF COMPUTER SCIENCE AND INTERDISCIPLINARY CENTER FOR BIOINFORMATICS UNIVERSITY OF LEIPZIG, HÄRTELSTRASSE 16-18, D-04107 LEIPZIG, GERMANY

MAX PLANCK INSTITUTE FOR MATHEMATICS IN THE SCIENCES INSELSTRASSE 22, D-04103 LEIPZIG, GERMANY

RNOMICS GROUP, FRAUNHOFER INSTITUT FÜR ZELLTHERAPIE UND IMMUNOLOGIE, DEUTSCHER PLATZ 5E, D-04103 LEIPZIG, GERMANY

DEPARTMENT OF THEORETICAL CHEMISTRY, UNIVERSITY OF VIENNA, WÄHRINGERSTRAßE 17, A-1090 WIEN, AUSTRIA

SANTA FE INSTITUTE, 1399 HYDE PARK RD., SANTA FE, NM87501, USA
E-MAIL: STUDLA@BIOINF.UNI-LEIPZIG.DE

ABSTRACT

We show that every simple, (weakly) connected, possibly directed and infinite, hypergraph has a unique prime factor decomposition with respect to the (weak) Cartesian product, even if it has infinitely many factors. This generalizes previous results for graphs and undirected hypergraphs to directed and infinite hypergraphs. The proof adopts the strategy outlined by Imrich and Žerovnik for the case of graphs and introduced the notion of diagonal-free grids as a replacement of the chord-free 4-cycles that play a crucial role in the case of graphs. This leads to a generalization of relation δ on the arc set, whose convex hull is shown to coincide with the product relation of the prime factorization.

Keywords: directed Hypergraph, Hypergraph, weak Cartesian Product, Prime Factor Decomposition, grid property

1. INTRODUCTION

Directed hypergraphs are the common generalization of both directed graphs and (undirected) hypergraphs. A *(directed) hypergraph* $H = (V, \mathcal{E})$ consists of a *vertex* set V and a family of *(directed) hyperedges* or *hyperarcs* $\mathcal{E}$. Each hyperedge $E \in \mathcal{E}$ is an ordered pair of sets of vertices $E = (T(E), H(E)) \neq (\emptyset, \emptyset)$. The sets $T(E) \subseteq V$ and $H(E) \subseteq V$ are called the *tail* and *head* of E, respectively. To avoid the risk of confusion we will sometimes write $V(H)$ and $\mathcal{E}(H)$ for the vertex set and the arc set of a hypergraph H.

A hypergraph $H = (V, \mathcal{E})$ is *undirected* if $T(E) = H(E)$ for all $E \in \mathcal{E}$. Most of the literature on hypergraphs is concerned with undirected hypergraphs [1], surveys on directed hypergraphs can be found in [7, 8]. Directed hypergraphs have been used as model for complex networks in biology and chemistry [15]. For instance a chemical reaction is naturally represented as a hyperedge E where $T(E)$ lists the educts and $H(E)$ the products of the chemical transformation.

We say a hypergraph is *finite* if its vertex set and its edge set is finite. A hypergraph that is not finite is said to be *infinite*. A hypergraph $H = (V, \mathcal{E})$ is called *simple*, if (i) $|T(E) \cup H(E)| > 1$ for all $E \in \mathcal{E}$ and (ii) there are no two distinct arcs $E, E' \in \mathcal{E}$ such that $T(E) = T(E')$ and $H(E) = H(E')$. A hypergraph is a directed graph if $|T(E)| = |H(E)| = 1$ for all $E \in \mathcal{E}$.

We will be concerned with products of hypergraphs, more precisely with the most fundamental question that arises in this context: Does a hypergraph have a unique decomposition into prime factors? The answer to this question depends on the definition of the product. In the case of graphs, the product structures are quite well understood [11].

Several notions of hypergraph products have been studied in the literature, usually for restricted versions of undirected hypergraphs, see e.g. [2, 6, 9, 20]. The most commonly studied hypergraph product is the direct product, which consists of the Cartesian product of vertex sets, and the Cartesian (set) products of the hyperedges [3, 5, 18, 17]. For direct products of N-systems, i.e., directed hypergraphs satisfying $|T(E)| = 1$ and $T(E) \subseteq H(E)$ for all $E \in \mathcal{E}$, a prime factor theorem was proved in [13].

In this contribution we will focus on the *Cartesian product* of finite and infinite directed hypergraphs with finitely or infinitely many factors. The Cartesian *graph* product was introduced by Gert Sabidussi [19], who showed that connected graphs have a unique Cartesian prime factor decomposition. This result was generalized by Wilfried Imrich [9] to simple undirected hypergraphs:

Definition. Let $\{H_i \mid i \in I\}$ be a family of (finite or infinite), but undirected hypergraphs. The Cartesian product $H = \Box_{i\in I} H_i$ is defined as follows:

$$V(H) = \bigtimes_{i\in I} V(H_i)$$

$$\mathcal{E}(H) = \{E \subseteq V(H) \mid p_j(E) \in \mathcal{E}(H_j) \text{ for exactly one } j \in I, \text{ and } |p_i(E)| = 1 \text{ for } i \neq j\},$$

where, for $j \in I$, $p_j : V(H) \to V(H_j)$ is the projection of the Cartesian product of the vertex sets into $V(H_j)$. The value $p_j(v)$ is also called *j-th coordinate* of vertex v.

The Cartesian product of undirected hypergraphs is also considered for example in [3, 4, 5] A factorization algorithm for so-called conformal hypergraphs, a rather small class of finite and connected hypergraphs, with respect to the Cartesian product, is described in [4].

While the Cartesian product hypergraph of finitely many connected hypergraphs is connected, whether they are finite or not, this does not hold for the product of infinitely many hypergraphs: In this case, there are vertices that differ in infinitely many coordinates and hence are not connected by a path of finite length. As in the case of graphs [19], an infinite connected hypergraph that has infinitely many Cartesian prime factors cannot be the Cartesian product of its factors, but a connected component of this Cartesian product [10]. This gives rise to the notion of a weak Cartesian product:

Definition. [10] Let $\{H_i \mid i \in I\}$ be a family of hypergraphs and let $a_i \in V(H_i)$ for $i \in I$. The *weak* Cartesian product $H = \Box_{i\in I}(H_i, a_i)$ of the "rooted" hypergraphs (H_i, a_i) is defined by

$$V(H) = \{v \in \bigtimes_{i\in I} V(H_i) \mid p_i(v) \neq a_i \text{ for at most finitely many } i \in I\}$$

$$\mathcal{E}(H) = \{E \subseteq V(H) \mid p_j(E) \in \mathcal{E}(H_j) \text{ for exactly one } j \in I, \text{ and } |p_i(E)| = 1 \text{ for } i \neq j\}.$$

For finite index sets I, the weak Cartesian product does not depend on the choice of the a_i and coincides with the usual Cartesian product. If I is infinite, it is the connected component of the Cartesian product containing $a = (a_i)_{i\in I}$. Every connected graph and undirected hypergraph has a unique representation as a weak Cartesian product [16, 10].

Here we extend and generalize these results further and show that every connected directed hypergraph has a unique prime factor decomposition with respect to the (weak)

Cartesian product. Instead of following the proof strategies of the classical papers, we adopt the approach of Imrich and Žerovnik [14] that constructs a product relation σ starting from simpler relations on the edge set E. In the case of graphs, the Square Property [12] plays a central role as technical device. The Grid Property, which is introduced here, serves as generalization of this construction. Together with a generalization of the relation δ, we arrive at our main results:

Theorem 1. Let γ be a convex equivalence relation on the arc set $\mathcal{E}(H)$ of a connected simple hypergraph H which satisfies the grid property. Then γ induces a factorization of H with respect to the weak Cartesian product.

Theorem 2. Every connected simple hypergraph has a unique representation as a weak Cartesian product.

Theorem 3. The product relation σ corresponding to the unique prime factor decomposition with respect to the weak Cartesian product of a connected simple hypergraph H equals the convex hull $\mathcal{C}(\delta)$ of the relation δ.

Since the Cartesian and the weak Cartesian product coincides for a finite number of factors, we also obtain the following corollaries.

Corollary. The prime factor decomposition of a connected hypergraph with finitely many factors with respect to the Cartesian product is unique in the class of simple hypergraphs.

Corollary. The prime factor decomposition of a connected hypergraph with respect to the Cartesian product is unique in the class of finite simple hypergraphs.

2. PRELIMINARIES

As far as possible, we follow the notation and terminology of Berge's classical book on hypergraphs [1], although our hypergraphs will in general be directed, $T(E) \neq H(E)$ for some $E \in \mathcal{E}$. For simplicity, we will refer to hyperarcs as arcs.

Two arcs $E, E' \in \mathcal{E}(H)$ of a hypergraph H are *incident* if $E \cap E' \neq \emptyset$, i.e., if there is a vertex that is contained in both arcs, independent of the directions. Two vertices $x, y \in V$ are *adjacent* if there is an arc E containing them, i.e., $x, y \in T(E) \cup H(E)$, again without regard of direction.

Concepts of sub-hypergraphs, paths, etc., are defined in the appendix to make this contribution self-consistent.

For the set product $V = \bigtimes_{i\in I} V_i$ we define the *projection* $p_j : V \to V_j$ by $(v_1, v_2, \dots, v_j, \dots) \mapsto v_j$. For subsets of V and ordered tuples of elements of V, the projection is defined element-wise. For example, for a hypergraph $H = (\bigtimes_{i\in I} V_i, \mathcal{E})$ we have

$$p_j(E) = p_j(T(E), H(E)) = (p_j(T(E)), p_j(H(E))) = (\bigcup_{v \in T(E)} p_j(v), \bigcup_{w \in H(E)} p_j(w)). \quad (1)$$

By abuse of notation, we will write

$$|p_j(E)| := |p_j(T(E)) \cup p_j(T(H))|$$

i.e., $|p_j(E)|$ refers to the cardinality of the union of the projections of head and tail of the arc E.

Definition. Let H_i, $i \in I$ be hypergraphs. The Cartesian product $\Box_{i \in I} H_i$ has the following vertex and arc sets:

(1) $V(\Box_{i \in I} H_i) = \times_{i \in I} V(H_i)$,
(2) and $E = (T(E), H(E))$ is an arc in $\Box_{i \in I} H_i$ if and only if there is a $j \in I$, such that
   (i) $p_j(E) = p_j((T(E), H(E))) = (T(p_j(E)), H(p_j(E))) \in \mathcal{E}(H_j)$ and
   (ii) $|p_i(E)| = 1$ for all $i \neq j$.

Figure 1 shows an example of a Cartesian product of two hypergraphs.

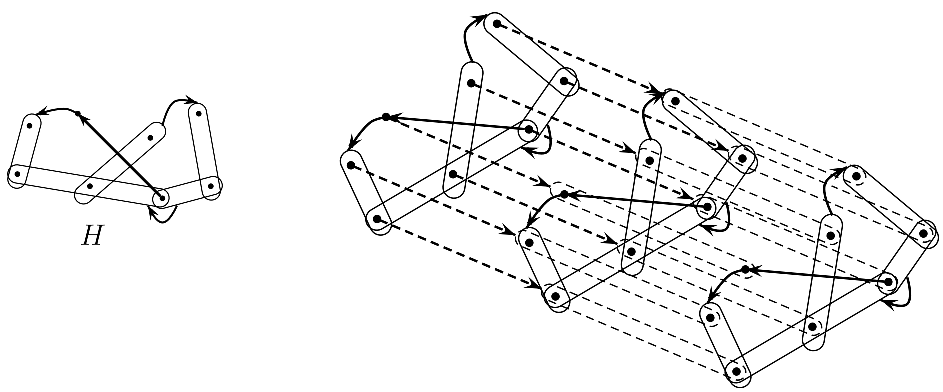

FIGURE 1. Hypergraph H and Cartesian product $H \Box E_{1,2}$, where $E_{1,2}$ is the hypergraph consisting of a single arc of three vertices such that $|T(E_{1,2})| = 1$ and $|H(E_{1,2})| = 2$.

Lemma 1. The Cartesian Product $H = \Box_{i \in I} H_i$ of hypergraphs H_i is an undirected hypergraph if and only if all of its factors are undirected.

Proof. The assertion follows directly from equ.(1). ∎

By construction, the Cartesian product is associative, commutative, distributive w.r.t. the disjoint union and has the trivial one vertex hypergraph K_1 without arcs as unit.

Lemma 2. The Cartesian product $H = \Box_{i=1}^{n} H_i$ of finitely many hypergraphs H_i is connected if and only if all of its factors H_i are connected.

Proof. Because of associativity and commutativity it suffices to show the assertion for two factors, hence let $H = H_1 \Box H_2$.

First, assume that H_1 and H_2 are connected. Let $v = (x, y)$ and $v' = (x', y')$ be two arbitrary vertices in $V(H)$. Consider a path $P_{xx'} = (E_1, \ldots, E_r)$ from x to x' in H_1 and a path $P_{yy'} = (F_1, \ldots, F_s)$ from y to y' in H_2. Then $(E_1 \times \{y\}, \ldots, E_r \times \{y\})$ is a path from (x, y) to (x', y) in H and $(\{x'\} \times F_1, \ldots, \{x'\} \times F_s)$ is a path from (x', y) to (x', y') in H. Hence, $(E_1 \times \{y\}, \ldots, E_r \times \{y\}, \{x'\} \times F_1, \ldots, \{x'\} \times F_s)$ is a path from v to v' in H.

W.l.o.g., suppose that H_1 is not connected, i.e., it can be written in the form $H_1 = H_1' + H_1''$. Since the Cartesian product is distributive w.r.t. the disjoint union, we have $H = H_1' \Box H_2 + H_1'' \Box H_2$. Hence H is the disjoint union of two hypergraphs, i.e., H is not connected. ∎

Let us now turn to the weak Cartesian product. As in the undirected case, the product of finitely many connected hypergraphs is connected, whether they are finite or not, but the product of infinitely many factors is not connected.

Definition. Let $\{H_i \mid i \in I\}$ be a family of hypergraphs and let $a_i \in V(H_i)$ for $i \in I$. The weak Cartesian product $H = \Box_{i \in I}(H_i, a_i)$ of the hypergraphs (H_i, a_i) rooted at a_i is given by

$$V(H) = \{v \in \bigtimes_{i \in I} V(H_i) \mid p_i(v) \neq a_i \text{ for at most finitely many } i \in I\}$$

$$\mathcal{E}(H) = \{E \subseteq V(H) \mid p_j(E) = p_j((T(E), H(E))) = (T(p_j(E)), H(p_j(E))) \in \mathcal{E}(H_j)$$
$$\text{for exactly one } j \in I, \text{ and } |p_i(E)| = 1 \text{ for } i \neq j\}.$$

We will, in the following write $\Box_{i \in I}^{a} H_i$ for $\Box_{i \in I}(H_i, a_i)$, where $a \in V(H)$ such that $p_i(a) = a_i$ for all $i \in I$.

Again, the weak Cartesian product does not depend on the $(a_i)_{i \in I}$ if I is finite. Moreover, it coincides with the Cartesian product if I is finite.

The partial hypergraph of H induced by all vertices of H which differ from $w \in V(H)$ exactly in the j-th coordinate is isomorphic to H_j. More formally

$$\langle \{v \in V(H) \mid p_k(v) = w_k \text{ for } k \neq j\} \rangle \simeq H_j.$$

We will call this partial hypergraph the H_j*-layer through* w and denote it by H_j^w. The isomorphism $H_j^w \to H_j$ is then the projection p_j. For $u \in V(H_j^w)$ we have $H_j^u = H_j^w$ and, moreover, $V(H_j^u) \cap V(H_j^w) = \emptyset$ if and only if $u \notin V(H_j^w)$.

Lemma 2 implies the following

Corollary. A weak Cartesian product $H = \Box^a_{i\in I} H_i$ is connected if and only if all of its factors H_i are connected. In this case, $H = \Box^a_{i\in I} H_i$ is the connected component of the Cartesian product $\Box_{i\in I} H_i$ that contains the vertex a.

Lemma 3. The Cartesian product $H = \Box_{i\in I} H_i$ of hypergraphs H_i is simple if and only if all of its factors H_i are simple.

Proof. Let E_{xj} denote an arc containing $x \in V$ and whose projection on the j-th factor is an arc in the j-th factor.

First let all $H_i,\ i \in I$ be simple and suppose H is not simple. We have to examine two cases: First, suppose H contains at least one loop on vertex x with coordinates $x_i,\ i \in I$. Hence there is a loop in some factor H_j, contradicting that H_j is simple. Thus, it must hold $|E| \geq 2$ for all $E \in \mathcal{E}(H)$. Second, let E_{xj} and E'_{yk} be two different arcs, such that $T(E_{xj}) = T(E'_{yk})$ and $H(E_{xj}) = H(E'_{yk})$. Hence, $x \in E'_{yk}$ and $j = k$. Thus, $p_j(T(E_{xj})) = p_k(T(E'_{yk}))$, as well as $p_j(H(E_{xj})) = p_j(H(E'_{yj}))$. Therefore H_j is not simple, a contradiction.

Now assume that (at least) one of the factors is not simple; w.l.o.g. say H_1. There are two possibilities: First, assume there is an arc $E \in \mathcal{E}(H_1)$ with $|E| = 1$, say $E = \{x_1\}$. Hence, $|E_{1x}| = |\{x\}| = 1$ for any $x \in V(H)$ with $p_1(x) = x_1$ and H would not be simple. Second, assume there are two different arcs $E_i, E_j \in \mathcal{E}(H_1)$, such that $T(E_i) = T(E_j)$ and $H(E_i) = H(E_j)$. Then, the definition of the Cartesian product implies that there are two arcs $E = E_i \times \{x\}$ and $E' = E_j \times \{x\}$ in H with $x \in \bigtimes_{l\in I\setminus\{1\}} V(H_l)$ such that $T(E) = T(E')$ and $H(E) = H(E')$. Hence, H is not simple. ∎

Note that the weak Cartesian product $\Box_{i\in I}(H_i, a_i)$ is a partial hypergraph of the Cartesian product $\Box_{i\in I} H_i$, that is induced by $V(\Box_{i\in I}(H_i, a_i))$. Hence we have

Corollary. The weak Cartesian product $H = \Box_{i\in I}(H_i, a_i)$ of hypergraphs H_i is simple if and only if all of its factors H_i are simple.

From here on we assume that all hypergraphs are simple.

3. UNIQUE PRIME FACTOR DECOMPOSITION

A hypergraph H is *prime* with respect to the (weak) Cartesian product if it cannot be represented as the (weak) Cartesian product of two nontrivial hypergraphs. A *prime factor decomposition (PFD)* of H is a representation as a Cartesian product $H = \Box_{i\in I} H_i$, or as a a weak Cartesian product $H = \Box^a_{i\in I} H_i$, resp., such that all factors H_i, $i \in I$, are prime and $H_i \not\simeq K_1$.

In order to show that every hypergraph has a unique representation as a weak Cartesian product, we follow the strategy of Imrich and Žerovnik [14] and characterize the so-called product relations defined on the arc set $\mathcal{E}$ of H. The advantage of this approach is that it does not require finiteness and hence also pertains to the weak Cartesian product of infinitely many factors.

Definition. A *product relation* is an equivalence relation on the arc set $\mathcal{E}(H)$ of a (weak) Cartesian product $H = \Box^a_{i\in I} H_i$ of (not necessarily prime) hypergraphs H_i such that, for $E, F \in \mathcal{E}(H)$, E and F are in relation γ, $E\gamma F$, if and only if there exists a $j \in I$, such that

$$|p_j(E)| > 1 \quad \text{and} \quad |p_j(F)| > 1,$$

and $|p_i(E)| = |p_i(F)| = 1$ holds for all $i \neq j$.

If all factors H_i are prime, we denote this relation by σ. Note that E and F are in relation σ if and only if their vertices differ in the same coordinate w.r.t. to the PFD. Let Σ_i, $i \in I$ be the equivalence classes of σ. By construction, every connected component of a partial hypergraph generated by the arcs of an equivalence class Σ_i is isomorphic to H_i. Furthermore, every union of σ-equivalence classes $\bigcup_{j\in J} \Sigma_j$, $J \subseteq I$ generates a partial hypergraph of H, whose connected components are isomorphic to $H_J := \Box^a_{j\in J} H_j$.

THE GRID PROPERTY

Definition. Let $\mathcal{G}$ be a collection of arcs of H of the form $E_a = (T(E_a), H(E_a))$, $a \in A$ and $F_b = (T(F_b), H(F_b))$, $b \in B$. We say that $\mathcal{G}$ is an $|A| \times |B|$-*grid* if, for all $a, a' \in A$ and $b, b' \in B$ with $a \neq a'$ and $b \neq b'$ the following two conditions are satisfied:

(i) E_a and F_b have exactly one vertex in common, i.e., $(T(E_a) \cup H(E_a)) \cap (T(F_b) \cup H(F_b)) = \{z_{ab}\}$, and

  (ia) $z_{ab} \in T(F_b)$ (resp. $H(F_b)$) if and only if $z_{ab'} \in T(F_{b'})$ (resp. $H(F_{b'})$) for all $b' \in B$, and

  (ib) $z_{ab} \in T(E_a)$ (resp. $H(E_a)$) if and only if $z_{a'b} \in T(E_{a'})$ (resp. $H(E_{a'})$) for all $a' \in A$, and

(ii) E_a and $E_{a'}$ have no common vertex for $a \neq a'$, i.e., $(T(E_a) \cup H(E_a)) \cap (T(E'_a) \cup H(E'_a)) = \emptyset$. Analogously, F_b and $F_{b'}$ are disjoint for $b \neq b'$.

An arc $D \in \mathcal{E}(H)$ satisfying $z_{ab} \in D$ and $z_{a'b'} \in D$ for all for $a, a' \in A$ and $b, b' \in B$ with $a \neq a'$ and $b \neq b'$, is a *diagonal* of the $|A| \times |B|$-grid $\mathcal{G}$.

The construction of the grid implies that E_a and F_b satisfy $|E_a| = |B|$ and $|F_b| = |A|$ for all $a \in A$, and $b \in B$.

Diagonal-free grids appear whenever two arcs of two hypergraphs are multiplied with respect to the Cartesian product. In this sense they generalize the chordless squares appearing as Cartesian products of arcs of undirected simple graphs. This suggests to generalize the relation δ [19, 14] in the following way:

Definition. Let H be a connected hypergraph. For $E, F \in \mathcal{E}(H)$ we say E and F are in *relation* δ, $E\delta F$, if one of the following conditions is satisfied:

(i) E and F have no vertex in common and form the opposite arcs of a 4-cycle.
(ii) E and F are incident to at least one common vertex and there is no grid without diagonals that contains both E and F.
(iii) $E = F$.

Note that whenever E and F share two or more vertices there is no $(|E| \times |F|)$-grid that contains E and F, and hence $E \delta F$.

Obviously, the δ is reflexive and symmetric. Its transitive closure δ^*, i.e., the smallest transitive relation containing δ, is therefore an equivalence relation.

Condition (ii) implies that any two incident arcs E and F with $E \not\delta F$ span an $(|E| \times |F|)$-grid without diagonals, whose arcs we denote by E, $\{E_a\}_{a \in A}$ and F, $\{F_b\}_{b \in B}$, where $E \delta E_a$ and $F \delta F_b$ for all $a \in A$ and $b \in B$, respectively.

Lemma 4. Two incident arcs that are not in relation δ span a unique $(|E| \times |F|)$-grid.

Proof. Suppose there exists another $(|E| \times |F|)$-grid consisting of arcs E, $\{E'_a\}_{a \in A}$ and F, $\{F'_b\}_{b \in B}$. Then there must be a $k \in A$ and an $l \in B$ such that $E'_k \notin \{E_a\}_{a \in A}$ and $F'_l \notin \{F_b\}_{b \in B}$, respectively. Hence there exists both an arc $E_r \in \{E_a\}_{a \in A}$ and and arc $F_s \in \{F_b\}_{b \in B}$ such that E'_k and E_r have common vertices and F'_l and F_s have common vertices. Thus there is a 4-cycle $E'_k E_r F_s F'_l$, where E'_k and F_s as well as E_r and F'_l are opposite arcs. Hence $E'_k \delta F_s$ and $E_r \delta F'_l$, and therefore $(E, F) \in \delta^*$. It follows that if E and F belong to distinct δ^*-equivalence classes, they span a unique $(|E| \times |F|)$-grid. ∎

This observation suggests the following definition:

Definition. Let γ be an equivalence relation on the arc set $\mathcal{E}(H)$ of a hypergraph H. We say γ has the *grid property* if any two adjacent arcs E and F of H with $E \not\gamma F$ span exactly one diagonal-free $|E| \times |F|$-grid.

Our discussion above implies that δ^* has the grid property.

Let γ be an arbitrary equivalence relation on the arc set of a hypergraph H that contains δ^*. For any two arcs E and F with $E \not\gamma F$ we also have $E \not\delta^* F$ and therefore, they span exactly one (diagonal-free) $|E| \times |F|$-grid. As a consequence, every equivalence relation γ that contains δ^* satisfies the grid property.

Lemma 5. Let H be a connected hypergraph and let γ be an equivalence relation on $\mathcal{E}(H)$ satisfying the grid property. Denote the equivalence classes of γ by Γ_i, $i \in I$. Then every vertex of $v \in V(H)$ is incident to an arc $E \in \Gamma_i$ for every $i \in I$.

Proof. Suppose that there is an equivalence class Γ_i of γ and a set of vertices that is not contained in any Γ_i-arc. By connectedness of H, there is a pair of vertices $u, v \in V(H)$ and an arc $E \in \mathcal{E}(H)$ with $\{u, v\} \subseteq T(E) \cup H(E)$ such that u belongs to a Γ_i-arc, say F, and there is no Γ_i-arc containing v. It follows that $E \notin \Gamma_i$, i.e., it must be contained in some other equivalence class Γ_k, $k \neq i$. By construction, E and F are two incident arcs belonging to different equivalence classes of γ and hence span a grid. Thus there must be a Γ_i-arc containing v, contradicting the assumption. ∎

Lemma 6. Let $H = \Box^a_{i\in I} H_i$ be a weak Cartesian product of prime hypergraphs H_i and let $E, F \in \mathcal{E}(H)$. If E and F are in relation δ, they are in relation σ.

Proof. Suppose first, that for the arcs E and F holds $E \cap F = \emptyset$ and there are arcs $E', F' \in \mathcal{E}(H)$ such that $\{E, E', F, F'\}$ is a 4-cycle. Moreover, we denote with $c(E)$ the coordinates where the vertices of the arc E differ. W.l.o.g. assume x_1 is common to E and E', x_2 is common to E and F', y_1 is common to F and E', and y_2 is common to F and F'. The coordinates varying within these arcs are denote by $c(E) = i$, $c(F) = j$, $c(E') = i'$, $c(F') = j'$; see Fig. 2.

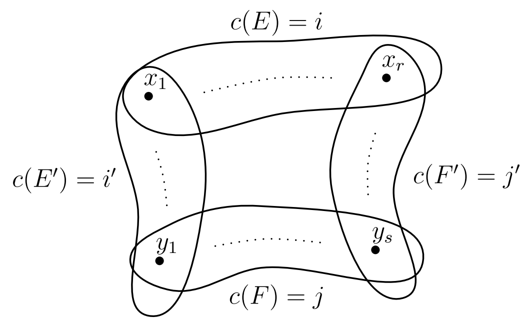

FIGURE 2. 4-cycle $EE'FF'$

Then we have:

$$p_k(x_1) = p_k(x_2) \qquad \text{for all } k \neq i \tag{2}$$
$$p_k(y_1) = p_k(y_2) \qquad \text{for all } k \neq j \tag{3}$$
$$p_k(x_1) = p_k(y_1) \qquad \text{for all } k \neq i' \tag{4}$$
$$p_k(x_2) = p_k(y_2) \qquad \text{for all } k \neq j'. \tag{5}$$

It follows from (2) and (5) that

$$p_k(x_1) = p_k(y_2) \qquad \text{for all } k \neq i, j' \tag{6}$$

and from (3) and (4)

$$p_k(x_1) = p_k(y_2) \qquad \text{for all } k \neq j, i'. \tag{7}$$

Therefore we have either $i = j$ and $i' = j'$ or $i = i'$ and $j' = j$. Assume $i \neq j$. Then the latter case must hold and we have $p_k(x_2) = p_k(x_1) = p_k(y_2)$ for all $k \neq i$, and since

$x_2 \neq y_1$ holds, $p_i(x_2) \neq p_i(y_1) = p_i(y_2)$ and $p_j(x_2) \neq p_j(y_2)$. Hence, x_2 and y_2 differ in more than one coordinate, thus they cannot lie in the same arc F', which contradicts the assumption, therefore, $i = j$ must hold, i.e. $c(E) = c(F)$, and therefore $E\sigma F$.

Now let E and F be incident arcs of a hypergraph H and assume that there is no $|E| \times |F|$-grid without diagonals containing them.

First, consider the case that E and F share more than a single vertex. Then there is an index $i \in I$ such that $|p_i(E)| > 1$ and in particular, $|p_i(E')| > 1$ holds for all $E' \subseteq T(E) \cup H(E)$ with $|E'| > 1$.

Since E and F have more than one vertex in common it follows that $|p_i((T(E) \cup H(E)) \cap (T(F) \cup H(F))| > 1$ and hence $|p_i(E)| > 1$ and $|p_i(F)| > 1$, and thus $E\sigma F$.

Now suppose that E and F share a single vertex v and assume that E and F are not in relation σ. Let $\hat{E} = T(E) \cup H(E) = \{v\} \cup \bigcup_{a\in A}\{x_a\}$ and $\hat{F} = T(F) \cup H(F) = \{v\} \cup \bigcup_{b\in B}\{y_b\}$. Furthermore set $j = c(F)$ and observe that $j \neq i$. For all $x_a \in E$, $a \in A$, and all $b \in B$ there exist vertices $z_{ab} \in V(H)$ such that

$$p_i(z_{ab}) = p_i(x_a) \tag{8}$$

$$p_k(z_{ab}) = p_k(y_b) \qquad \text{for all } k \neq i. \tag{9}$$

Using (9) and the fact, that $\{v, y_b\} \subseteq F$ for all $b \in B$ and $\{v, x_a\} \subseteq E$, we can conclude that the set $\hat{F}_a = \{x_a\} \cup \bigcup_{b\in B}\{z_{ab}\}$ satisfies

$$p_j(\hat{F}_a) = \{p_j(x_a)\} \cup \bigcup_{b\in B}\{p_j(z_{ab})\} = \{p_j(v)\} \cup \bigcup_{b\in B}\{p_j(y_b)\} = p_j(\hat{F}) \tag{10}$$

as well as

$$p_k(z_{ab}) = p_k(y_b) = p_k(v) = p_k(x_a) \quad \text{for all } k \neq i, j \quad \text{and for all } b \in B \tag{11}$$

Now we use (8) and (11) and obtain

$$p_k(z_{ab}) = p_k(x_a) \qquad \text{for all } k \neq i \quad \text{and for all } b \in B. \tag{12}$$

Analgously, the set $\hat{E}_b = \{y_b\} \cup \bigcup_{a\in A}\{z_{ab}\}$ satisfies

$$p_i(\hat{E}_b) = \{p_i(y_b)\} \cup \bigcup_{a\in A}\{p_i(z_{ab})\} = \{p_i(v)\} \cup \bigcup_{a\in A}\{p_i(x_a)\} = p_i(\hat{E}). \tag{13}$$

Again, we use (8) and the fact that $\{v, y_b\} \subseteq F$ for all $b \in B$.

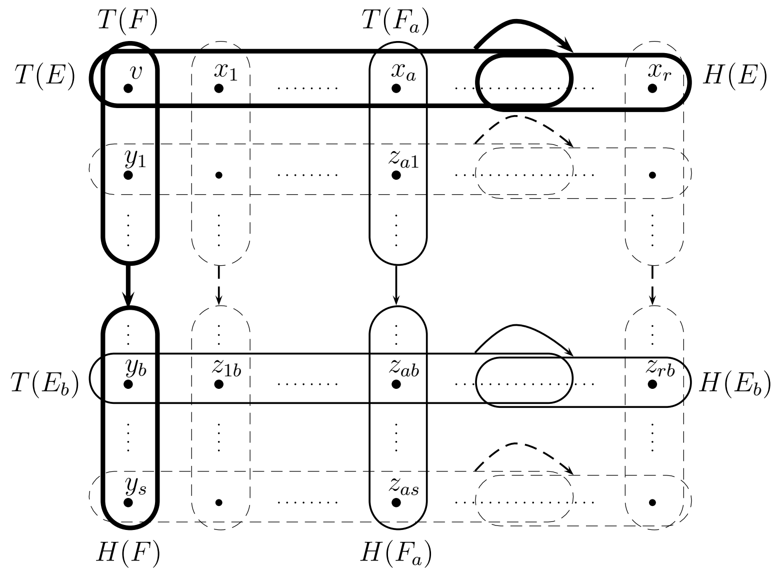

FIGURE 3. Arcs E and F (thick arcs) with $(E, F) \notin \sigma$ and the $|E| \times |F|$-grid they span

It will be useful to relabel v as z_{00}, $\hat{E}$ as $\hat{E}_0$, $\hat{F}$ as $\hat{F}_0$, the x_l as z_{l0} for all $l \in A$, and the y_l as z_{0l} for all $l \in B$. Then we have

$$\hat{E}_b \cap \hat{F}_a = \{z_{ab}\} \tag{14}$$

for all $a \in A_0 := A \cup \{0\}$ and $b \in B_0 := B \cup \{0\}$.

Since $z_{0b} \in T(F)$ (resp. $H(F)$) if and only if $p_j(z_{0b}) \in p_j(T(F))$ (resp. $p_j(H(F))$) for all $b \in B_0$, we can conclude by the definition of the (weak) Cartesian product and from equations (10) and (12) that for all $\hat{F}_a$, $a \in A$, there exists an arc $F_a = (T(F_a), H(F_a))$ with $|F_a| = |F|$ such that $\hat{F}_a = T(F_a) \cup H(F_a)$. More precisely, we have

$$\begin{aligned} T(F_a) &= \{z_{ab} \mid b \in B_0 \text{ and } z_{0b} \in T(F)\}, \text{ resp.} \\ H(F_a) &= \{z_{ab} \mid b \in B_0 \text{ and } z_{0b} \in H(F)\} \end{aligned} \tag{15}$$

for all $a \in A$.

An analogous argument can be made for $\hat{E}_b$: Since $z_{a0} \in T(E)$ (resp. $H(E)$) if and only if $p_i(z_{a0}) \in p_i(T(E))$ (resp. $p_i(H(E))$) for all $a \in A_0$, it follows from equations (9) and (13) that for all $\hat{E}_b$, $b \in B$, there exists an arc $E_b = (T(E_b), H(E_b))$ with $|E_b| = |E|$ such that $\hat{E}_b = T(E_b) \cup H(E_b)$. More precisely, we have

$$\begin{aligned} T(E_b) &= \{z_{ab} \mid a \in A_0 \text{ and } z_{a0} \in T(E)\}, \text{ resp.} \\ H(E_b) &= \{z_{ab} \mid a \in A_0 \text{ and } z_{a0} \in H(E)\} \end{aligned} \tag{16}$$

for all $b \in B$.

Equation (14) immediately implies that the intersection of more then two arcs is empty.

Consequently, whenever two adjacent arcs E and F are not in relation σ, they span an $|E| \times |F|$-grid with $|E| \times |F|$ vertices.

It remains to show that these grids have no diagonals. Therefore, we need to show that there is no $D \in \mathcal{E}(H)$ such that $\{z_{ab}, z_{a'b'}\} \subseteq D$ for all $a, a' \in \{0\} \cup A$, $b, b' \in \{0\} \cup B$ with $a \neq a'$ and $b \neq b'$. For $z_{ab}, z_{a'b'}$ we have:

$$p_i(z_{ab}) \overset{(8)}{=} p_i(z_{a0}) \neq p_i(z_{a'0}) \overset{(8)}{=} p_i(z_{a'b'}) \tag{17}$$

$$p_j(z_{ab}) \overset{(9)}{=} p_j(z_{0b}) \neq p_j(z_{0b'}) \overset{(9)}{=} p_j(z_{a'b'}) \tag{18}$$

The inequalities follow from the fact, that $\{z_{a0}, z_{a'0}\} \subseteq E_0$ and $\{z_{0b}, z_{0b'}\} \subseteq F_0$.

That means that z_{ab} and $z_{a'b'}$ differ in more than one coordinate, hence, they cannot be contained in the same arc, which completes the proof. ∎

Lemma 6 implies that $\delta \subseteq \sigma$ and since σ is an equivalence relation, $\delta^* \subseteq \sigma$ holds as well. Thus, the product relation σ has the grid property.

CONVEXITY

Lemma 7. Let $H = \Box^a_{i \in I} H_i$ be a weak Cartesian product of connected hypergraphs H_i. Then each $H_J = \Box^a_{j \in J} H_j$-layer is convex for any index set $J \subseteq I$.

Proof. It suffices to show that, whenever there is a path P between two arbitrary vertices u and v of the same H_J-layer H_J^u containing no arcs of this layer, then there exists a path Q which entirely lies in H_J^u and satisfies $|Q| < |P|$.

Suppose $P = (u = u_0, E_1, u_1, E_2, \ldots, u_{k-1}, E_k, u_k = v)$. Since u and v belong to the same H_j-layer, $p_l(u) = p_l(v)$ holds for all $l \in I \setminus J$. There must be an arc E_i of P such that E_i is contained in some H_J layer, by assumption different from H_J^u. Otherwise we would have $p_l(u) = p_l(v)$ for all $l \in J$, hence $p_l(u) = p_l(v)$ for all $l \in I$, i.e. $u = v$.

Let $\{E_{j_1}, \ldots, E_{j_r}\}$ be a subset of arcs of P, with $j_1, j_2, \ldots, j_r \in \{1, \ldots, k\}$, $j_1 < j_2 < \ldots < j_r$, that are in some H_J-layer different from H_J^u, and no arc is the copy of another. To be more precise, for each j_i there is a $k_i \in J$ with

$$p_{k_i}(E_{j_i}) \in \mathcal{E}(H_{k_i}) \qquad \text{and} \qquad p_{k_a}(E_{j_a}) \neq p_{k_b}(E_{j_b}) \quad \text{for } a \neq b \tag{19}$$

and j_r is maximal. Note that $a \neq b$ does not imply $k_a \neq k_b$.

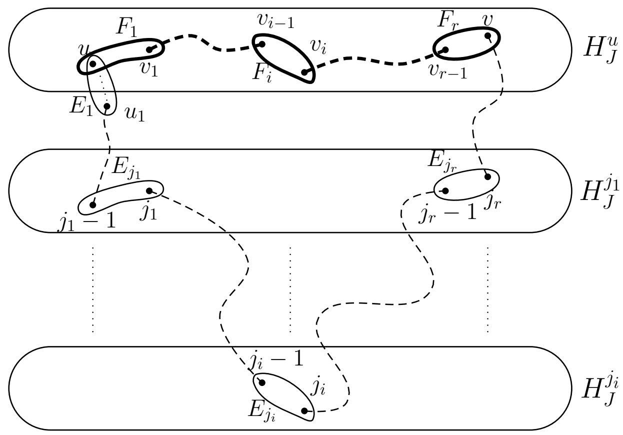

FIGURE 4. **Idea of the proof:** Path P and Path Q (thick) which we got by shifting the arcs E_{j_i} in the H_J^u-layer

By assumption, E_1 is not contained in any H_J-layer, thus $r < k$. Without loss of generality, we can assume that the E_{j_i} are not incident. In the following we will denote the vertices u_{j_i} by j_i. If we set $j_0 := u_0$, we have

$$p_l(j_{i-1}) = p_l(j_i - 1) \quad \text{for all } l \in J \text{ and all } i \in \{1, \dots, r\} \tag{20}$$

and since $j_i - 1, j_i \in E_{j_i}$,

$$p_l(j_i) = p_l(j_i - 1) \tag{21}$$

holds for all $l \neq k_i$, $k_i \in J$. Furthermore, for each j_i, $i \in \{1, \dots, r\}$, there exists a $v_i \in V(H)$, such that

$$p_l(v_i) = p_l(u) \quad \text{for all } l \in I \setminus J \tag{22}$$

$$p_l(v_i) = p_l(j_i) \quad \text{for all } l \in J \tag{23}$$

In particular, Equation (22) implies $v_i \in V(H_J^u)$ for all $i \in \{1, \dots, r\}$. It follows

$$p_l(v_{i-1}) \overset{(23)}{=} p_l(j_{i-1}) \overset{(20),(21)}{=} p_l(j_i) \overset{(23)}{=} p_l(v_i) \quad \text{for all } l \in J \setminus \{k_i\} \tag{24}$$

and by Equation (22) we have

$$p_l(v_{i-1}) = p_l(v_i) \quad \text{for all } l \in I \setminus J \tag{25}$$

hence by Equations (24) and (25)

$$p_l(v_{i-1}) = p_l(v_i) \quad \text{for all } l \neq k_i \tag{26}$$

In other words, any two vertices v_{i-1}, v_i lie in the same H_{k_i} layer for some $k_i \in J$.

Next we show that there are arcs $F_i \in \mathcal{E}(H_J^u)$ containing both v_{i-1} and v_i. From Equations (20) and (23) it follows

$$p_{k_i}(v_{i-1}) \overset{(23)}{=} p_{k_i}(j_{i-1}) \overset{(20)}{=} p_{k_i}(j_i - 1) \tag{27}$$

$$p_{k_i}(v_i) \overset{(23)}{=} p_{k_i}(j_i) \tag{28}$$

Thus we have by Equations (27), (28) and (19)

$$p_{k_i}(v_{i-1}),\ p_{k_i}(v_i) \in p_{k_i}(E_{j_i}) \tag{29}$$

Hence by Equations (22), (26) and (29), for each $i \in \{1, \ldots, r\}$ there exists an arc F_i in H_J^u, such that $v_{i-1}, v_i \in F_i$.

Now consider the vertex v_r. Since there is no more arc E_j, $j > j_r$, of P that is contained in any H_J-layer, j_r and v belong to the same $\widehat{H_J} := \Box_{i \in I \setminus J} H_i$-layer Therefore we have

$$p_l(v_r) \overset{(23)}{=} p_l(j_r) = p_l(v) \quad \text{for all } l \in J \tag{30}$$

and from the definition of v_r and the fact that u and v are in the same H_J-layer, it follows

$$p_l(v_r) \overset{(22)}{=} p_l(v) \quad \text{for all } l \in I \setminus J. \tag{31}$$

Therefore, we can conclude $v_r = v$, and we have found a path $Q = (u = v_0, F_1, v_1, \ldots, v_{r-1}, F_r, v_r = v)$ from u to v, whose arcs all lie entirely within H_J^u. Furthermore, we have $|Q| = r < k = |P|$, which completes the proof. ∎

Definition. An equivalence relation γ on the arc set $\mathcal{E}(H)$ of a connected hypergraph H with equivalence classes Γ_i, $i \in I$, is called *convex* if for any $J \subseteq I$ every connected component of the partial hypergraph generated by $\bigcup_{i \in J} \Gamma_i$ is convex.

By Lemma 7, the product relation σ is a convex relation. Moreover, any product relation must be convex and has to satisfy the grid property.

Lemma 8. Let γ be an equivalence relation on the arc set $\mathcal{E}(H)$ of a connected hypergraph H which satisfies the grid property. Let Γ be an equivalence class of γ. If all connected components of the partial hypergraph of H generated by Γ are convex, they are isomorphic.

Proof. Let H_Γ be the partial hypergraph generated by Γ with connected components C_i, $i \in I$ and let $\widehat{\Gamma}$ denote the union of all equivalence classes of γ, distinct from Γ, i.e.,

$\widehat{\Gamma} = \bigcup_{\Gamma' \neq \Gamma} \Gamma'$. It suffices to show that any two components C_1 and C_2 that are connected by a $\widehat{\Gamma}$-arc are isomorphic. We define a mapping $\varphi : V(C_1) \to V(C_2)$ such that $x \mapsto \varphi x$, whenever x and φx are connected by a $\widehat{\Gamma}$-arc. From the grid property and Lemma 5, it follows that for all $x \in V(C_1)$ there exists a $\varphi x \in V(C_2)$. The grid property ensures that adjacent vertices in C_1 have different images in C_2 and arcs in C_1 map onto arcs in C_2, such that $\varphi(E) = \varphi((T(E), H(E))) = (T(\varphi(E)), H(\varphi(E))) \in \mathcal{E}(C_2)$ for all $E \in \mathcal{E}(C_1)$. Convexity implies that non adjacent vertices in C_1 have different images in C_2 as well, i.e., the mapping φ is injective. On the other hand we can extend φ^{-1} to a mapping $\psi : V(C_2) \to V(C_1)$. Analogously, it follows that for all $y \in V(C_2)$ there is a ψy in $V(C_1)$, hence $\varphi^{-1} = \psi$, i.e., φ is bijective, and every arc in C_2 maps onto an arc in C_1, thus φ is an isomorphism between C_1 and C_2. ∎

Sometimes, the transitive closure of δ is already convex. If this is the case, then each path between two vertices of the same connected component of an equivalence class of δ^* must contain at least one arc of this equivalence class as a consequence of Lemma 7.

Lemma 9. Let γ be an equivalence relation on the arc set $\mathcal{E}(H)$ of a connected hypergraph H satisfying the grid property with only two equivalence classes Γ and $\widehat{\Gamma}$. Let H_Γ and $H_{\widehat{\Gamma}}$ be the partial hypergraphs generated by Γ and $\widehat{\Gamma}$, with connected components C_i, $i \in I$, and $\widehat{C}_j$, $j \in J$, respectively. Then

$$V(C_i) \cap V(\widehat{C}_j) \neq \emptyset \quad \text{for all } i \in I,\ j \in J.$$

In particular,

$$|V(C_i) \cap V(\widehat{C}_j)| = 1$$

holds if C_i and $\widehat{C}_j$ are convex.

Proof. Suppose there are components C_i, $\widehat{C}_j$ with $V(C_i) \cap V(\widehat{C}_j) = \emptyset$, such that they have minimal distance. Let $P = (v_0, E_1, v_1, E_2, \ldots, E_k, v_k)$ be a shortest path from C_i to $\widehat{C}_j$, such that $v_0 \in V(C_i)$ and $v_k \in V(\widehat{C}_j)$. Obviously, the first arc E_1 must lie in $\widehat{\Gamma}$ and the vertex v_1 is not in C_i, otherwise E_1 would be in Γ which contradicts the minimality of P. Lemma 5 implies that v_1 must be contained in a Γ-component, say C_k. Since the distance from C_k to $\widehat{C}_j$ is smaller than $|P|$, we have $V(C_k) \cap V(\widehat{C}_j) \neq \emptyset$. Let w be a vertex in $V(C_k) \cap V(\widehat{C}_j)$ and let P' be a path from v_1 to w in C_k. By repeated application of the grid property we obtain a vertex u in $V(C_i)$ connected to w by a $\widehat{\Gamma}$-arc. But then u must be in $V(\widehat{C}_j)$ and thus $|V(C_i) \cap V(\widehat{C}_j)| \geq 1$.

Now assume $|V(C_i) \cap V(\widehat{C}_j)| \geq 2$. Let $u, w \in V(C_i) \cap V(\widehat{C}_j)$. By connectivity we have a path Q from u to w in C_i and a path Q' from u to w in $\widehat{C}_j$ as well. Therefore either $|Q| > |Q'|$ or $|Q'| > |Q|$ or $|Q| = |Q'|$ holds. Hence either C_i or $\widehat{C}_j$ or both are not convex, and thus the second assertion holds. ∎

PROOF OF THE THEOREMS

We are now able to prove Theorem 1:

Proof of Theorem 1. First assume that γ has only two equivalence classes Γ and $\widehat{\Gamma}$ with connected components C_i, $i \in J$ and $\widehat{C}_j$, $j \in \widehat{J}$ respectively, of the generated partial hypergraphs. By Lemma 9 we can assign uniquely determined coordinates (i, j) to each vertex of H, whenever $\{v\} = V(C_i) \cap V(\widehat{C}_j)$, $i \in J$, $j \in \widehat{J}$. On the other hand for all such coordinates there exists a uniquely determined vertex in $V(H)$, since $|V(C_i) \cap V(\widehat{C}_j)| = 1$.

In the following we will identify each vertex of H with its coordinates. Obviously we have $V(C_i) = \{(i, j) \mid j \in \widehat{J}\}$ for all $i \in J$ and $V(\widehat{C}_j) = \{(i, j) \mid i \in J\}$ for all $j \in \widehat{J}$. Recall that Lemma 8 implies that the C_i are isomorphic for all $i \in J$. In particular $C_1 \simeq C_i$ holds for all $i \in J$. The isomorphism is given by the mapping

$$(1, j) \mapsto (i, j) \quad \text{for all} \quad j \in \widehat{J}.$$

If C_1 and C_i are connected by an arc, it is an isomorphism as in the proof of Lemma 8. If they are connected by a path, it is an isomorphism by induction on the length of the path. Analogously we have $\widehat{C}_1 \simeq \widehat{C}_j$ for all $j \in \widehat{J}$ given by the isomorphism

$$(i, 1) \mapsto (i, j) \quad \text{for all} \quad i \in J.$$

A set of vertices $E = \{(i_1, j_1), \ldots, (i_q, j_q), \ldots\}$, $i_1, \ldots, i_q, \ldots \in J$, $j_1, \ldots, j_q, \ldots \in \widehat{J}$, is an arc in H with $T(E) = \{(i_{t_1}, j_{t_1}), \ldots, (i_{t_r}, j_{t_r}), \ldots\}$ and $H(E) = \{(i_{h_1}, j_{h_1}), \ldots, (i_{h_s}, j_{h_s}), \ldots\}$ if and only if either

(i) it is in the same C_i, hence $i_1 = \ldots = i_q = \ldots = i$ and $E_1 = \{(1, j_1), \ldots, (1, j_q), \ldots\}$ is an arc in C_1, such that $T(E_1) = \{(1, j_{t_1}), \ldots, (1, j_{t_r}), \ldots\}$ and $H(E_1) = \{(1, j_{h_1}), \ldots, (1, j_{h_s}), \ldots\}$, or

(ii) it is in the same $\widehat{C}_j$, hence $j_1 = \ldots = j_q = \ldots = j$ and $E_2 = \{(i_1, 1), \ldots, (i_q, 1)\}$ is an arc in $\widehat{C}_1$, such that $T(E_2) = \{(i_{t_1}, 1), \ldots, (i_{t_r}, 1), \ldots\}$ and $H(E_2) = \{(i_{h_1}, 1), \ldots, (i_{h_s}, 1), \ldots\}$.

That is, H is isomorphic to $C_1 \Box \widehat{C}_1$.

Now define hypergraphs H_1 and H_2 by setting $V(H_1) = \{i : (i, 1) \in V(C_1)\}$ and $V(H_2) = \{j : (1, j) \in V(\widehat{C}_1)\}$. H_1 and H_2 are isomorphic to C_1 and $\widehat{C}_1$ by the isomorphic mappings $i \mapsto (i, 1)$ and $j \mapsto (1, j)$ respectively, thus $H = H_1 \Box H_2$.

Assume now that γ has arbitrarily many equivalence classes Γ_i, $i \in I$. Let γ_i be the equivalence relation with the two equivalence classes Γ_i and $\widehat{\Gamma}_i = \bigcup_{k \in I, k \neq i} \Gamma_k$ for arbitrary $i \in I$. As already shown, we get a factorization of H into two factors $H_i \Box \widehat{H}_i$ where H_i and $\widehat{H}_i$ belong to Γ_i and $\widehat{\Gamma}_i$, respectively. We will call the projection, more precisely the image of the projection, of a vertex v in $H_i \Box \widehat{H}_i$ into the factor H_i the i-th coordinate of v, denoted by v^i.

Clearly, we can assign coordinates to each vertex. If two vertices u, v have the same i-th coordinate, then, by convexity, there is no Γ_i-arc on any shortest path between them. Thus, if u and v have the same coordinates, there is no nontrivial shortest path between them, hence $u = v$. It follows that the assignment of coordinates to vertices of a connected hypergraph H is bijective.

A subset $E = \{v_1, \ldots, v_q, \ldots\}$ of $V(H)$ is an arc of H with $T(E) = \{v_{t_1}, \ldots, v_{t_r}, \ldots\}$ and $H(E) = \{v_{h_1}, \ldots, v_{h_s}, \ldots\}$ if and only if the v_l differ in the same coordinate, say the i-th, for all $l \in \{1, \ldots, q, \ldots\}$ and $E_i = \{v_1^i, \ldots, v_q^i, \ldots\}$ is an arc in H_i with $T(E_i) = \{v_{t_1}^i, \ldots, v_{t_r}^i, \ldots\}$ and $H(E_i) = \{v_{h_1}^i, \ldots, v_{h_s}^i, \ldots\}$.

Since H is connected, its vertices differ in at most finitely many coordinates, thus we have $H \simeq \Box_{i \in I}^a H_i$ for any $a \in V(H)$. ∎

The equivalence relation whose only equivalence class is the entire arc set of a hypergraph H is trivially convex and satisfies the grid property and is therefore a product relation. This relation always exists. By Theorem 1 we can conclude that any convex relation on the arc set of a connected hypergraph that satisfies the grid property is a product relation and induces a factorization of this hypergraph. The smallest convex relation satisfying the grid-property, if such a relation exists at all, therefore must induce a PFD with respect to the Cartesian product.

The next lemma is well known for undirected graphs. Its proof is one of the key application of the square-property [14]. This proof step can directly be generalized to the grid-property in the case of hypergraphs. Moreover, since the definition of a path (i.e., a weak path) coincides with the definition of a path in the undirected case, we can immediately state the next lemma for hypergraphs.

Lemma 10. Let γ_j, $j \in J$ be an arbitrary set of convex relations on the arc set $\mathcal{E}(H)$ of a hypergraph H containing δ. Then $\gamma = \bigcap_{j \in J} \gamma_j$ is convex.

It is clear that for arbitrary equivalence relations on the arc set of a hypergraph, which satisfy the grid-property, their intersection also has the grid-property. Therefore Lemma 10 implies that there is exactly one finest convex equivalence relation on the arc set $\mathcal{E}(H)$ of a hypergraph H satisfying the grid property, namely the intersection of all convex relations on $\mathcal{E}(H)$ containing δ, that is its *convex hull*, $\mathcal{C}(\delta)$.

Conversely, any product relation must be convex and contains δ. Thus we have proved the Theorems 2 and 3. We conclude, furthermore

Corollary. The PFD of a connected hypergraph with respect to the Cartesian product is unique in the class of finite simple hypergraphs.

ACKNOWLEDGMENTS

We thank Wilfried Imrich for many stimulating discussions on graph and hypergraph products.

References

[1] C. Berge. *Hypergraphs: Combinatorics of finite sets*. North-Holland, Amsterdam, NL, 1989.

[2] C. Berge and M. Simonovitis. The coloring numbers of the direct product of two hypergraphs. In *Hypergraph Seminar*, volume 411 of *Lecture Notes in Mathematics*, pages 21–33, Berlin, Heidelberg, 1974. Springer-Verlag.

[3] A. Bretto. Hypergraphs and the helly property. *Ars Comb.*, 78:23–32, 2006.

[4] A. Bretto, Y. Silvestre, and T. Vallée. Cartesian product of hypergraphs: properties and algorithms. In *4th Athens Colloquium on Algorithms and Complexity (ACAC 2009)*, volume 4 of *EPTCS*, pages 22–28, 2009.

[5] W Dörfler. Multiple Covers of Hypergraphs. *Annals N.Y. Acad. Sci.*, 319:169–176, 1979.

[6] W. Dörfler and D. A. Waller. A category-theoretical approach to hypergraphs. *Arch. Math.*, 34:185–192, 1980.

[7] G. Gallo, G. Longo, S. Pallottino, and S. Nguyen. Directed hypergraphs and applications. *Discrete Appl. Math.*, 42:177–201, 1993.

[8] G. Gallo and M. G. Scutellà. Directed hypergraphs as a modelling paradigm. *Rivista AMASES*, 21:97–123, 1998.

[9] W. Imrich. Kartesisches Produkt von Mengensystemen und Graphen. *Studia Sci. Math. Hungar.*, 2:285–290, 1967.

[10] W. Imrich. über das schwache Kartesische Produkt von Graphen. *J. Comb. Theory*, 11:1–16, 1971.

[11] W. Imrich and S. Klavžar. *Product Graphs: Structure and Recognition*. John Wiley & Sons, New York, 2000.

[12] W. Imrich, S. Klavžar, and D. F. Rall. *Graphs and Their Cartesian Product*. Topics in Graph Theory. A K Peters, Ltd., Wellesley, MA, 2008.

[13] W. Imrich and P. F. Stadler. A prime factor theorem for a generalized direct product. *Discussiones Math. Graph Th.*, 26:135–140, 2006.

[14] W. Imrich and J. Žerovnik. Factoring Cartesian-product graphs. *J. Graph Theory*, 18:557–567, 1994.

[15] S Klamt, U-U Haus, and F Theis. Hypergraphs and cellular networks. *PLoS Comput Biol*, 5:e1000385, 2009.

[16] D. J. Miller. Weak cartesian product of graphs. *Colloq. Math.*, 21:55–74, 1970.

[17] D. Mubayi and V. Rödl. On the chromatic number and independence number of hypergraph products. *J. Comb. Theory, Ser. B*, 97:151–155, 2007.

[18] R. Pemantle, J. Propp, and D. Ullman. On tensor powers of integer programs. *SIAM J. Discrete Math.*, 5:127–143, 1992.

[19] G. Sabidussi. Graph multiplication. *Math. Z.*, 72:446–457, 1960.

[20] X. Zhu. On the chromatic number of the product of hypergraphs. *Ars Comb.*, 34:25–31, 1992.

APPENDIX: BASIC DEFINITIONS

A *partial hypergraph* $H' \subseteq H$ of a hypergraph $H = (V, \mathcal{E})$ is a hypergraph $H' = (V', \mathcal{E}')$ with $V' \subseteq V$ and $\mathcal{E}' \subseteq \mathcal{E}$. A partial hypergraph $H' = (V', \mathcal{E}') \subseteq H = (V, \mathcal{E})$ is *generated by the arc set* $\mathcal{E}'$ if $V' = \cup_{E \in \mathcal{E}'} E$. It is *induced by the vertex set* V', $H' = \langle V' \rangle$, if $\mathcal{E}' = (E \in \mathcal{E} \mid E \subseteq V')$.

A *strong path* or *simple path* P *of length* k in $H = (V, \mathcal{E})$ is a sequence $P = [v_0, E_1, v_1, E_2, \ldots, E_k, v_k]$ of distinct vertices and arcs of H, such that $v_0 \in T(E_1)$, $v_k \in H(E_k)$ and $v_j \in H(E_j) \cap T(E_{j+1})$. A *weak path* P^w *of length* k in a hypergraph $H = (V, \mathcal{E})$ is a sequence $P^w = (v_0, E_1, v_1, E_2, \ldots, E_k, v_k)$ of distinct vertices and arcs of H, such that $v_0 \in T(E_1) \cup H(E_1)$, $v_k \in T(E_k) \cup H(E_k)$ and $v_j \in [T(E_j) \cup H(E_j)] \cap [T(E_{j+1}) \cup H(E_{j+1})]$. For the sake of convenience, we will refer to weak paths simply as paths. A hypergraph H is said to be *weakly connected* or simply *connected* for short, if each two vertices of H can be connected by a path. A path between two partial hypergraphs H' and H'' of a hypergraph H is a path in H between two vertices $v \in V(H')$ and $w \in V(H'')$.

The *distance* between two vertices of a hypergraph is the length of the shortest path connecting them. The distance $d_H(H', H'')$ between two partial hypergraphs H', H'' is the minimal length of a path between H' and H''. A partial hypergraph $H' \subseteq H$ is *convex* if all shortest paths of H between vertices of H' are contained in H'.

A 4-*cycle* in H is a partial hypergraph generated by arcs E_1, E_2, E_3, and E_4, such that $E_i \cap E_{i+1} \neq \emptyset$, where the subscripts are taken modulo 4.

A *homomorphism* from $H_1 = (V_1, \mathcal{E}_1)$ into $H_2 = (V_2, \mathcal{E}_2)$ is a mapping $\varphi : V_1 \to V_2$ such that $\varphi(E)$ is an arc in H_2 whenever E is an arc in H_1 with the property that $\varphi(T(E)) = T(\varphi(E))$ and $\varphi(H(E)) = H(\varphi(E))$. A mapping $\varphi : V_1 \to V_2$ is a *weak homomorphism* if arcs are mapped either on arcs or on vertices. A bijective homomorphism φ whose inverse function is also a homomorphism is called an *isomorphism.*