International Scholarly Research Network
ISRN Discrete Mathematics
Volume 2011, Article ID 806193, 22 pages
doi:10.5402/2011/806193

Research Article

Equivalence between Hypergraph Convexities

Francesco M. Malvestuto, Mauro Mezzini, and Marina Moscarini

Computer Science Department, Sapienza University of Rome, 00185 Rome, Italy

Correspondence should be addressed to Francesco M. Malvestuto, malvestuto@di.uniroma1.it

Received 25 August 2011; Accepted 29 September 2011

Academic Editors: G. Isaak and J. Tarhio

Let G be a connected graph on V. A subset X of V is all-paths convex (or ap-convex) if X contains each vertex on every path joining two vertices in X and is monophonically convex (or m-convex) if X contains each vertex on every chordless path joining two vertices in X. First of all, we prove that ap-convexity and m-convexity coincide in G if and only if G is a tree. Next, in order to generalize this result to a connected hypergraph H, in addition to the hypergraph versions of ap-convexity and m-convexity, we consider canonical convexity (or c-convexity) and simple-path convexity (or sp-convexity) for which it is well known that m-convexity is finer than both c-convexity and sp-convexity and sp-convexity is finer than ap-convexity. After proving sp-convexity is coarser than c-convexity, we characterize the hypergraphs in which each pair of the four convexities above is equivalent. As a result, we obtain a convexity-theoretic characterization of Berge-acyclic hypergraphs and of γ-acyclic hypergraphs.

1. Introduction

Convexity is a fundamental concept occurring in geometry, topology, and functional analysis, and the problem of computing convex hulls is at the core of many computer engineering applications, for instance, in robotics, computer graphics, or optimization (see page 125 in [1]). For the theory of abstract convex structures see [2].

This paper is motivated by the paper by Farber and Jamison [3] where structural properties of different notions of convexity in graphs and hypergraphs are stated. We focus on the following four notions of convexities in hypergraphs: m-convexity (for *monophonic convexity*) [4], c-convexity (for *canonical convexity*) [4], sp-convexity (for *simple-path convexity*) [3], and ap-convexity (for *all-paths convexity*), which are known to be related to each other by the following implications:

$$\text{ap-convex} \Longrightarrow \text{sp-convex} \Longrightarrow m\text{-convex} \qquad c\text{-convex} \Longrightarrow m\text{-convex}. \tag{1.1}$$

Note that hypergraph m-convexity generalizes graph m-convexity [3, 5], and both hypergraph sp-convexity and ap-convexity generalize graph ap-convexity [6, 7]. After proving the implication sp-convex $\Rightarrow$ c-convex, we characterize the hypergraphs in which each pair of the four convexities above is equivalent.

As an application, suppose that for two hypergraph convexities θ and θ' we have two algorithms A and A' to compute θ-convex hulls and θ'-convex hulls for hypergraphs and that the computational complexity of A is less than A'. If θ and θ' are equivalent in a class of hypergraphs then, for every hypergraph H belonging to this class, we can compute θ'-convex hulls using the algorithm A instead of A'.

As a consequence of our equivalence results, we obtain the following convexity-theoretic characterizations of the so-called "Berge-acyclic" hypergraphs and "γ-acyclic" hypergraphs [8]:

(i) a hypergraph is Berge-acyclic if and only if c-convexity and ap-convexity are equivalent in the hypergraph (see Theorem 6.9);

(ii) a hypergraph is Berge-acyclic if and only if m-convexity and ap-convexity are equivalent in the hypergraph (see Theorem 6.10);

(iii) a hypergraph is γ-acyclic if and only if c-convexity and sp-convexity are equivalent in the hypergraph (see Theorem 5.5);

(iv) a hypergraph is γ-acyclic if and only if m-convexity and sp-convexity are equivalent in the hypergraph (see Theorem 5.6).

It is worth noting that Berge-acyclic and γ-acyclic hypergraphs belong to a family of hypergraphs (α-, β-, γ-, and Berge-acyclic hypergraphs) which enjoy a number of theoretical and computational properties in database theory [8, 9], artificial intelligence [10, 11], and statistics [12, 13]. Moreover, c-convexity is a "convex geometry" in α-acyclic hypergraphs [4], and β-acyclicity characterizes those hypergraphs in which sp-convexity yields a "convex geometry" [3].

To achieve the above-mentioned results, we need several standard hypergraph-theoretic definitions (such as the four acyclicity types above) and additional notions which are introduced *ad hoc*. The outline of the paper is as follows. Section 2 contains basic definitions and results on graphs and hypergraphs. In Section 3 we review basic results on structural properties of m-convexity, c-convexity, sp-convexity, and ap-convexity in graphs and hypergraphs; moreover, we state some preliminary results. In Section 4 we characterize the hypergraphs in which m-convexity and c-convexity are equivalent. In Section 5 we prove that γ-acyclicity characterizes those hypergraphs in which c-convexity (or m-convexity) and sp-convexity are equivalent. In Section 6 we characterize the hypergraphs in which sp-convexity and ap-convexity are equivalent; moreover, we prove that Berge-acyclicity characterizes those hypergraphs in which c-convexity (or m-convexity) and ap-convexity are equivalent. Section 7 contains an open problem for future research.

2. Terminology and Notation

2.1. Graphs

Henceforth, we only consider graphs [14] with no loops and no multiple edges, which henceforth will be referred simply to as *graphs*. Let G be a graph. Two vertices of G are *adjacent* if they are joined by some edge of G. A nonempty subset X of $V(G)$ is a *clique* if every two

distinct vertices in X are adjacent. The *subgraph* of G *induced* by a nonempty subset X of $V(G)$ is the graph, denoted by $G\ [X]$, with vertex set X in which two distinct vertices are adjacent if and only if they are adjacent in G. The notation $G\ [V \setminus X]$ is abridged into $G \setminus X$.

A *path* is a sequence $p = (a_0, a_1, \ldots, a_k)$, $k \geq 1$, of distinct vertices such that a_{i-1} and a_i are adjacent for $1 \leq i \leq k$. The path p is said to *join* a_0 and a_k (or, equivalently, to be an a_0–a_k path) and to have *length* k; moreover, if $k > 1$, p is said to *pass through* each a_h, $1 \leq h \leq k-1$, and two vertices a_i and a_j on p are said to be *consecutive* if $|i-j| = 1$. By $V(p)$ we denote the vertex set $\{a_0, a_1, \ldots, a_k\}$.

The *distance* between two vertices u and v is the length of any minimum-length u–v path. A graph G is *distance hereditary* if, for every two vertices u and v of G and for every connected, induced subgraph G' containing u and v, the distances between u and v in G' and G are the same.

A *cycle* of *length* k, $k \geq 2$, is a sequence $c = (a_0, a_1, \ldots, a_{k-1}, a_0)$ where $(a_0, a_1, \ldots, a_{k-1})$ is a path, and the vertices a_0 and a_{k-1} are adjacent. Two vertices a_i and a_j on the cycle c are *consecutive* if either $|i-j| = 1$ or $|i-j| = k-1$. By $V(c)$ we denote the set of vertices $\{a_0, a_1, \ldots, a_{k-1}\}$. A *chord* of c is an edge joining two nonconsecutive vertices on c.

A graph is *chordal* if every cycle of length at least 4 has a chord. A graph is *strongly chordal* if it is chordal and, for every cycle c of even length, there are two nonconsecutive vertices at odd distance on c that are adjacent. A graph is *Ptolemaic* if it is distance hereditary and chordal.

Let G be a graph with at least two vertices. A vertex u of G is a *cut vertex* (or an "articulation point") of G if the number of connected components of $G \setminus \{u\}$ is greater than the number of connected components of G or, equivalently, there exist two vertices $v \neq u$ and $w \neq u$ in the connected component of G containing u such that every v–w path passes through u [14]. A *block* of G is a maximal connected partial graph of G containing no cut vertices. A block of G is *trivial* if it consists of two vertices and *nontrivial* otherwise. Finally, G is a *block graph* if the vertex set of every block of G is a clique.

Proposition 2.1 (see [14]). *Let G be a nontrivial block. For every three distinct vertices u, v, and w of G, there exists a v–w path that passes through u.*

Proposition 2.2 (see [14]). *Let G be a connected graph, and let B be a block of G. If w is not a vertex of B, then B contains a cut vertex u of G such that, for every vertex $v \neq u$ of B, every v–w path passes through u.*

Let G be a graph. Two connected vertices are *separated* by a subset S of $V(G)$ if they are in two distinct components of the induced graph $G \setminus X$. A nonempty subset X of $V(G)$ is *nonseparable* if $G[X]$ is connected and no two vertices in X are separated by a clique of G. The *prime components* of G are the subgraphs of G induced by maximal nonseparable sets.

2.2. Hypergraphs

Generalizing notions and convexities from graphs to hypergraphs is not always straightforward, because there are often several nonequivalent ways to do this and different terminologies. This is true also for notions that hypergraph convexities are based on. For example, "simple paths" in [3] are called "chordless chains" in [15], "simple circuits" in [3] are called "chordless cycles" in [15], and "weak β-cycles" in [3] if they are of length at least 3, "nest" vertices in [3] are called "simple" vertices in [15].

The following basic definitions are taken from [16].

A (*generic*) *hypergraph* is a (possibly empty) set H of nonempty sets; the elements of H are the (*hyper*) *edges* of H and their union is the *vertex set* of H, denoted by $V(H)$.

A hypergraph is *trivial* if it has only one edge and *nontrivial* otherwise. A *partial* (*sub*) *hypergraph* of hypergraph H is any subset of H. A hypergraph is *simple* if no edge is contained in another edge. The *reduction* of a hypergraph H is the partial hypergraph of H whose edges are the maximal (with respect to set-inclusion) edges of H.

Let X be a nonempty subset of $V(H)$. The *subhypergraph* of H*induced* by X is the hypergraph, denoted by $H[X]$, whose edges are exactly the maximal (with respect to set-inclusion) edges of the hypergraph $\{A \cap X : A \in H\} \setminus \{\emptyset\}$. The notation $H\ [V(H) \setminus X]$ is abridged into $H \setminus X$.

A *partial edge* is a nonempty vertex set that is contained in some edge. Two vertices are *adjacent* if they belong together in some edge. A nonempty subset X of $V(H)$ is a *clique* if every two distinct vertices in X are adjacent. A hypergraph is *conformal* if every clique is a partial edge.

A *path* is a sequence $p = (a_0, A_1, a_1, \ldots, A_k, a_k)$, $k \geq 1$, where the a_i's are pairwise distinct vertices, the A_i's are pairwise distinct edges, and $\{a_{i-1}, a_i\} \subseteq A_i$ for $1 \leq i \leq k$. The path p is said to *join* a_0 and a_k (or, equivalently, to be an a_0–a_k path), to have length k and, if $k > 1$, to *pass through* each a_h, $1 \leq h \leq k-1$. Moreover, two vertices a_i and a_j on the path p are *consecutive* if $|i-j| = 1$. Finally, by $V(p)$ we denote the vertex set $\{a_0, a_1, \ldots, a_k\}$, and by $H(p)$ we denote the partial hypergraph $\{A_1, \ldots, A_{k-1}, A_k\}$ of H.

A hypergraph is *connected* if any two vertices are joined by a path. A nonempty subset X of $V(H)$ is *connected* if the induced subhypergraph $H\ [X]$ is connected. Note that if H is a graph, then the subgraph of H induced by X equals the reduction of $H[X]$.

The *connected components* of H are the subhypergraphs of H induced by maximal connected subsets of $V(H)$.

The following definitions of an "articulation set" and of a "block" in a hypergraph are the natural generalizations of the notions of a cut vertex and of a block in a graph [8, 17].

Let H be a reduced hypergraph. A *separator* is a nonempty subset S of $V(H)$ such that there exist two connected vertices of H that are in two distinct components of the induced subhypergraph $H \setminus S$. A separator of H is an *articulation set* if it is the intersection of two edges of H. A nonempty subset X of $V(H)$ is *nonseparable* if $H[X]$ is connected and has no articulation set. A *block* of H is the reduction of the subhypergraph of H induced by a maximal nonseparable set.

Example 2.3. Let $H = \{A_1, A_2, A_3, A_4, A_5\}$, where $A_1 = \{1,2\}$, $A_2 = \{1,3\}$, $A_3 = \{2,4\}$, $A_4 = \{3,4,5\}$, and $A_5 = \{4,5,6\}$. The hypergraph H is shown in Figure 1. The set $\{4,5\}$ is the only articulation set of H, and the blocks of H are shown in Figure 2.

Finally, with a hypergraph H we can associate two graphs: the "2-section" of H and the "incidence graph" of H, which are defined as follows.

The 2-*section* (also called "adjacency graph" or "underlying graph" or "primal graph", or "Gaifman graph") of H is the graph with vertex set $V(H)$ in which two vertices are adjacent if and only if they are adjacent in H. We denote the 2-section of H by $H_{[2]}$.

The *incidence graph* of H is the bipartite graph with bipartition $(V(H), H)$, where $a \in V(H)$ and $A \in H$ are joined by an edge if and only if $a \in A$. We denote the incidence graph of H by $G(H)$, and the *size* of H is the number of vertices and edges of $G(H)$ [18]. Note that, if H is connected, then the size of H is $O(e)$ where e is the number of edges of $G(H)$.

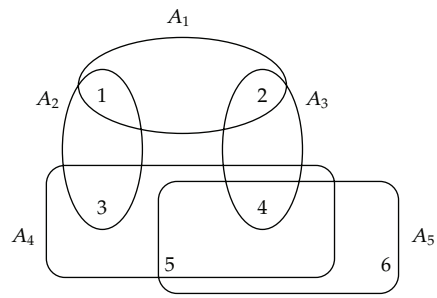

Figure 1

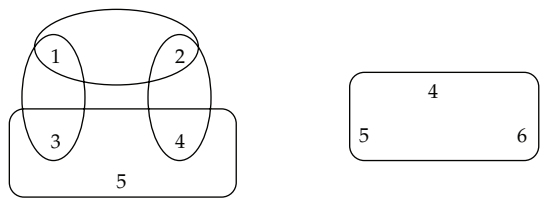

Figure 2

2.3. Acyclicity

Fagin [8] introduced four notions of acyclicity for hypergraphs which are now recalled and, in the next sections, will be proven to be closely related to hypergraph convexities.

A *cycle* (also called a "circuit" [3]) is a sequence $c = (a_0, A_1, a_1, \ldots, A_{k-1}, a_{k-1}, A_k, a_0)$, $k \geq 2$, where $(a_0, A_1, a_1, \ldots, A_{k-1}, a_{k-1})$ is a path, $\{a_0, a_{k-1}\} \subseteq A_k$ and $A_k \neq A_h$ for $1 \leq h \leq k-1$. The cycle c is said to have *length* k; moreover, two vertices a_i and a_j on c are *consecutive* if either $|i-j| = 1$ or $|i-j| = k-1$. By $V(c)$ we denote the set of vertices $\{a_0, a_1, \ldots, a_{k-1}\}$, and by $H(c)$ we denote the partial hypergraph $\{A_1, \ldots, A_{k-1}, A_k\}$ of H.

A γ-*cycle* is a cycle c of length at least 3 such that at most one vertex in $V(c)$ belongs to three or more edges of $H(c)$.

A β-*cycle* (a "weak β-cycle" in [8]) is a cycle c of length at least 3 such that every vertex in $V(c)$ belongs to exactly two edges of $H(c)$.

A hypergraph is *Berge-acyclic* if it contains no cycles, γ-*acyclic* if it contains no γ-cycles, and β-*acyclic* if it contains no β-cycles (or, equivalently, if every partial hypergraph is α-acyclic).

A reduced hypergraph is α-*acyclic* if all its blocks are trivial hypergraphs. A hypergraph is α-*acyclic* precisely if its reduction is α-acyclic.

It is well known [8] that the following implications on hypergraphs hold:

$$\text{Berge-acyclic} \Longrightarrow \gamma\text{-acyclic} \Longrightarrow \beta\text{-acyclic} \Longrightarrow \alpha\text{-acyclic} \tag{2.1}$$

but none of their reverse implications holds in general.

Several characterizations of Berge-acyclicity, γ-acyclicity, β-acyclicity, and α-acyclicity exist, and the following is based on the 2-section of a hypergraph.

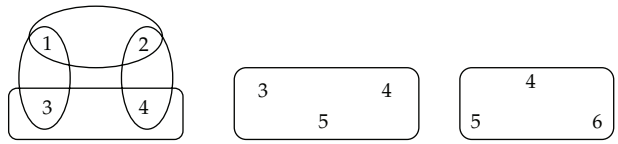

Figure 3

Proposition 2.4 (see [19]). *A hypergraph is Berge-acyclic if and only if it is conformal and its 2-section is a block graph. A hypergraph is γ-acyclic if and only if it is conformal and its 2-section is a Ptolemaic graph. A hypergraph is β-acyclic if and only if it is conformal and its 2-section is a strongly chordal graph. A hypergraph is α-acyclic if and only if it is conformal and its 2-section is a chordal graph.*

By Proposition 2.4, β-acyclic hypergraphs are the same as "totally balanced" hypergraphs [3] and as "totally decomposable" hypergraphs in [12]. Finally, note that α-acyclic hypergraphs are called "acyclic" hypergraphs in [4, 9, 11, 17, 18] and "decomposable" hypergraphs in [12, 13, 20, 21].

Before closing this subsection, we mention two α-acyclic hypergraphs which in some sense represent the "superstructure" of a graph and of a hypergraph.

The *prime hypergraph* of graph G is the (reduced) hypergraph whose edges are precisely the vertex sets of the prime components of G.

Proposition 2.5 (see [21]). *The prime hypergraph of a graph is a reduced, α-acyclic hypergraph.*

A nonempty subset X of $V(H)$ is a *compact set* of H if $H[X]$ is connected and no partial edge of $H[X]$ is a separator. Note that if X is a compact set, then $H[X]$ has no articulation set; but the reverse does not hold in general (see Example 2.3). A *compact component* of H is the reduction of the subhypergraph of H induced by a maximal compact set. The *compact hypergraph* of H is the (reduced) hypergraph whose edges are precisely the vertex sets of the compact components of H.

Proposition 2.6 (see [22, 23]). *The compact hypergraph of a hypergraph is a reduced, α-acyclic hypergraph. Moreover, a hypergraph is α-acyclic if and only if its reduction equals its compact hypergraph.*

Example 2.3 (continued). The compact components of H are shown in Figure 3 and the compact hypergraph of H is $\{\{1,2,3,4\},\ \{3,4,5\},\ \{4,5,6\}\}$.

3. Convexities in Graphs and Hypergraphs

In this section, we recall the definitions and basic results on some convexities in graphs and hypergraphs. Moreover, we state some preliminary results.

Let H be a connected hypergraph. A set $\theta(H)$ of subsets of $V(H)$ is a *convexity space* [24, 25] if

(i) the empty set, the singletons, and $V(H)$ belong to $\theta(H)$,

(ii) $\theta(H)$ is closed under set intersection,

(iii) every set in $\theta(H)$ is connected.

The sets in a convexity space $\theta(H)$ are called the *θ-convex sets* of H. For a subset X of $V(H)$, the *θ-convex hull* of X is the minimal (with respect to set inclusion) θ-convex set of H that includes X.

3.1. Graph Convexities

In this section, we recall the definitions and basic results on monophonic convexity (*m-convexity*) and all-paths convexity (ap-*convexity*) on a connected graph [14, 17, 24–26]. Let G be a connected graph. By $m(G)$ and $\mathrm{ap}(G)$ we denote the m-convexity space and the ap-convexity space on G, respectively.

A *chord* of a path p is an edge of G that joins two nonconsecutive vertices on p. A path is *chordless* (or "induced" or "minimal") if it has no chords. A subset X of $V(G)$ is *m-convex* if, for every chordless path p joining two vertices in X, each vertex on p belongs to X.

The following result provides a known characterization of m-convex sets. Let X be a subset of $V(H)$; an X–X *path* is a path $p = (a_0,\ A_1,\ a_1, \ldots,\ A_k,\ a_k)$ with $V(p) \cap X = \{a_0, a_k\}$.

Theorem 3.1 (see [5]). *Let G be a connected graph. A subset X of $V(G)$ is m-convex if and only if, for every two distinct vertices u and v in X, u and v are joined by an X–X path, then u and v are adjacent in G.*

Let G be a connected graph with n vertices and m edges. Dourado et al. [25] gave an algorithm for computing the m-convex hull of a subset X of $V(G)$ which runs in $O(mn^2|X|^2)$ time. A better algorithm was given by Kannan and Changat [27], which runs in $O(mn)$ time.

A subset X of $V(G)$ is ap-*convex* if, for every path p joining two vertices in X, each vertex on p belongs to X. It is easily seen that if G is a tree, then the ap-convex hull of a subset X of $V(G)$ can be computed in $O(m)$ time simply by deleting the leaves of G that are not in X.

We now state a characterization of those graphs G on which $\mathrm{ap}(G) = m(G)$.

Theorem 3.2. *Let G be a connected graph.The equality* $\mathrm{ap}(G) = m(G)$ *holds if and only if G is a tree.*

Proof. Since every chordless path is trivially a path, the inclusion $\mathrm{ap}(G) \subseteq m(G)$ is obvious. If G is a tree, then trivially one has $\mathrm{ap}(G) = m(G)$. Assume that G is not a tree. To prove that $\mathrm{ap}(G) \neq m(G)$, consider any nontrivial block of G, say B. Let u and v be two adjacent vertices of B. The path (u, v) is the only chordless u–v path in G and, hence, $\{u, v\}$ is m-convex. On the other hand, $\{u, v\}$ is not ap-convex since, by Propositions 2.1 and 2.2, its ap-convex hull is $V(B)$, which proves that $\mathrm{ap}(G) \neq m(G)$. □

By Theorem 3.2, if G is a tree, then the m-convex hull of a subset X of $V(G)$ coincides with the ap-convex hull of X and, hence, can be computed in $O(m)$ time.

3.2. Hypergraph Convexities

Let H be a connected hypergraph. In this section we recall the definitions and basic results on monophonic convexity (*m-convexity*), canonical convexity (*c-convexity*), simple-path convexity (sp-*convexity*), and all-paths convexity (ap-*convexity*) on H. By $m(H)$, $c(H)$, $\mathrm{sp}(H)$, and $\mathrm{ap}(H)$ we denote the m-convexity space, the c-convexity space, the sp-convexity space, and the ap-convexity space on H, respectively.

3.2.1. *m*-Convexity

Hypergraph m-convexity [4] was not defined in terms of paths but using the hypergraph-theoretic version of Theorem 3.1, that is, a subset X of $V(H)$ is *m-convex* if, for every two distinct vertices u and v in X, u and v are joined by an X–X path, then u and v are adjacent in H.

We now recall a useful characterization of m-convex sets. To this end, we need further definitions.

Let X be a subset of $V(H)$; two edges A and B of H are *connected outside* X, written $A \equiv_X B$, if

(i) $A = B$ or

(ii) $(A \cap B) \setminus X \neq \emptyset$ or

(iii) there exists an edge C of H such that $(A \cap C) \setminus X \neq \emptyset$ and $B \equiv_X C$.

The edge relation $\equiv_X$ is an equivalence relation; the classes of the resultant partition of H will be referred to as the *X-components* of H. The hypergraph H is *X-connected* if H has exactly one X-component. Note that H is X-connected if the graph $H_{[2]} \setminus X$ is connected and no edge of H is completely contained in X. For an X-component K of H, we call the set $X \cap V(K)$ the *boundary* of X with K.

Example 3.3. Consider again the hypergraph $H = \{A_1, A_2, A_3, A_4, A_5\}$ of Example 2.3 (see Figure 1). For $X = \{1, 3, 4\}$, the X-components of H are

$K_1 = \{A_1, A_3\}$ and the boundary of X with K_1 is $\{1, 4\}$,

$K_1 = \{A_2\}$ and the boundary of X with K_2 is $\{1, 3\}$,

$K_1 = \{A_4, A_5\}$ and the boundary of X with K_3 is $\{3, 4\}$.

Theorem 3.4 (see [4]). *Let H be a connected hypergraph. A subset X of $V(H)$ is m-convex if and only if either $X = \emptyset$ or the boundary of X with every X-component of H is a clique of H.*

As we noted above, in contrast with m-convexity in graphs, the original definition of m-convexity in hypergraphs was given without having recourse to any path type. We now prove that m-convexity in hypergraphs can be related to the following generalization of the notion of a chordless path in a graph, which (is different from that given in [15] and) is defined as follows.

A *chord* of a path p is a pair of nonconsecutive vertices on p which are adjacent in H. A path in H is *chordless* if it has no chords.

Theorem 3.5. *Let H be a connected hypergraph. A subset X of $V(H)$ is m-convex if and only if, for every chordless path p joining two vertices in X, each vertex on p belongs to X.*

Proof. (*if*) Assume that X contains $V(p)$ for every chordless path p joining two vertices in X. By Theorem 3.4, it is sufficient to prove that the boundary of X with every X-component of H is a clique of H. Suppose, by contradiction, that there exists an X-component K of H such that the boundary of X with K contains two vertices u and v that are not adjacent in H. Let Y be the boundary of X with K. Since u and v are not adjacent in H, u and v are not adjacent

in K and, since K is X-connected, there exists a u–v path p in K of length at least 2 such that $Y \cap V(p) = \{u, v\}$. Let p' be a u–v path of minimum length in $K(p)$. Of course, $V(p') \subseteq V(p)$ and p' is a chordless path; moreover, one has $Y \cap V(p') = \{u, v\}$. Since u and v are not adjacent in K, p' must be of length at least 2 and, hence, there exists a vertex w in $V(p') \cap V(K)$ that does not belong to Y. Since $X \cap V(p') = Y \cap V(p') = \{u, v\}$, w does not belong to X so that X does not contain $V(p')$ (contradiction).

(*only if*) Assume that X is m-convex. By Theorem 3.4, the boundary of X with every X-component of H is a clique of H. Suppose, by contradiction, that there exist two vertices u and v in X and a chordless u–v path $p = (a_0, A_1, a_1, \ldots, A_k, a_k)$ such that $V(p) \setminus X \neq \emptyset$. Let $i = \min\{h > 1 : a_h \notin X\}$ and let $j = \min\{h > i : a_h \in X\}$. Of course, both a_{i-1} and a_j belong to the boundary of X with the X-component of H containing a_i and, since the boundary of X with every X-component of H is a clique, a_{i-1} and a_j are adjacent. Since a_{i-1} and a_j are nonconsecutive on p, the pair $\{a_i, a_j\}$ is a chord of p (contradiction). □

Finally, let H be a connected hypergraph with n vertices and m edges. An algorithm for computing m-convex hulls in a hypergraph H is the "monophonic-closure algorithm" [4] which runs in $O(mn)$ time if the prime hypergraph of $H_{[2]}$ is given. Since the time needed to construct the prime hypergraph of $H_{[2]}$ is $O(en)$ [21], where e is the number of edges of $H_{[2]}$, and since $e = O(n^2)$, the time complexity of the monophonic-closure algorithm is $O(mn + n^3)$. On the other hand, it is easy to check that $(a_0, A_1, a_1, \ldots, A_k, a_k)$ is a chordless path in H if and only if $(a_0, a_1, \ldots, a_k)$ is a chordless path in $H_{[2]}$, which implies that $m(H) = m(H_{[2]})$ so that, given $H_{[2]}$, the m-convex hull of any vertex set in H can be computed in $O(ne)$ time, that is, in $O(n^3)$ time using the Kannan-Changat algorithm.

3.2.2. *c-Convexity*

A subset X of $V(H)$ is *c-convex* if the boundary of X with every X-component of H is a partial edge of H. Let H be a connected hypergraph with n vertices and m edges. It is proven in [4] that c-convex hulls can be computed using the Maier-Ullman algorithm [28], which runs in $O(m^4 n)$ time. A more efficient algorithm is the "canonical-closure algorithm" [4], which runs in $O(mn)$ time if the compact hypergraph of H is given. Since the time needed to construct the compact hypergraph of H is $O(m^3 n)$ [23], the time complexity of the canonical-closure algorithm is $O(m^3 n)$. Note that if H is α-acyclic, then the time complexity of the algorithm reduces to $O(mn)$; however, we can do better using the "selective-reduction algorithm" [18] which is linear in the size of H.

3.2.3. *sp-Convexity*

A path p in H is *simple* [3] if every two nonconsecutive vertices on p are not adjacent in the partial hypergraph $H(p)$; equivalently, a path p in H is simple if p is a chordless path in $H(p)$. A subset X of $V(H)$ is sp-*convex* if, for every simple path p joining two vertices in X, each vertex on p belongs to X. Let H be a connected hypergraph with n vertices and m edges. An efficient algorithm for computing sp-convex hulls was given in [29], which runs in $O(mn^3 e)$, where e is the number of edges of the incidence graph of H. Since $e \leq mn$, the time complexity of the algorithm is $O(m^2 n^4)$. However, if H is β-acyclic, then using the Anstee-Farber algorithm [15], the sp-convex hull of vertex set X can be computed in $O(mn^2)$ time simply by deleting the "nest" vertices of H that are not in X (see [12]).

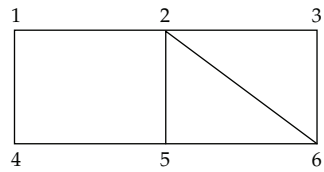

Figure 4

3.2.4. ap-Convexity

Let H be a connected hypergraph. The convexity space $\mathrm{ap}(H)$ is defined in the same way as in Section 3.1, that is, a subset X of $V(H)$ is ap-*convex* if, for every path p joining two vertices in X, each vertex on p belongs to X. Again one always has $\mathrm{ap}(H) \subseteq \mathrm{sp}(H)$. In Section 6 we will give an efficient algorithm for computing ap-convex hulls.

4. c-Convexity versus m-Convexity

Since every partial edge is a clique, one always has $c(H) \subseteq m(H)$. In this section we characterize the class of hypergraphs H for which $c(H) = m(H)$. First of all, observe that if H is conformal, then every clique is a partial edge so that, by Theorem 3.4, every m-convex set of H is also c-convex so that $c(H) = m(H)$. We will see that conformality is not a necessary condition for $c(H) = m(H)$.

A clique X of H is a *boundary clique* if there exists an X-component K of H such that X equals the boundary of X with K. A hypergraph H is *weakly conformal* if every boundary-clique of H is a partial edge of H. Of course, every conformal hypergraph is weakly conformal. The following is an example of a weakly conformal hypergraph that is not conformal.

Example 4.1. Consider the (hyper)graph H of Figure 4. The only cliques of H that are not boundary cliques are the two cliques with cardinality 3, namely, the sets $\{2, 3, 6\}$ and $\{2, 5, 6\}$. Since each clique of H with cardinality less than 3 is a partial edge of H, each boundary clique is a partial edge and, hence, H is weakly conformal.

Theorem 4.2. *Let H be a connected hypergraph. The equality $c(H) = m(H)$ holds if and only if H is weakly conformal.*

Proof. (*if*) Let X be any m-convex set. Let K be any X-component of H, and let Y be the boundary of X with K. By the very definition of m-convexity, Y is a clique; moreover, K is also a Y-component of H and the boundary of Y with K is Y itself. Therefore, Y is a boundary clique of H. Since H is weakly conformal, Y is a partial edge of H. It follows that the boundary of X with every X-component of H is a partial edge of H and, hence, X is c-convex.

(*only if*) Let X be any boundary clique of H. Since X is a clique, from the very definition of m-convexity it follows that X is m-convex and, since $c(H) = m(H)$ by hypothesis, X is c-convex. From the very definition of c-convexity, it follows that the boundary of X with every X-component of H is a partial edge. Finally, since X is a boundary clique of H, there exists an X-component K of H such that X is the boundary of X with K. So, X is a partial edge of H. It follows that H is weakly conformal. □

5. sp-Convexity versus c-Convexity and m-Convexity

In this section we characterize the class of hypergraphs H for which $\mathrm{sp}(H) = c(H)$ and the class of hypergraphs H for which $\mathrm{sp}(H) = m(H)$.

5.1. Equivalence between sp-Convexity and c-Convexity

Let H be a connected hypergraph. We first prove that $\mathrm{sp}(H) \subseteq c(H)$. To achieve this, we need the following two technical lemmas.

Lemma 5.1. *If p is a u–v path in H and u and v are not adjacent in $H(p)$, then there exists in $H(p)$ a simple u–v path p' of length at least 2 and with $V(p') \subseteq V(p)$.*

Proof. Let $(a_0, A_1, a_1, \ldots, A_k, a_k)$ be a u–v path. Let $i(1) = \max\{h : a_h$ and a_0 are adjacent in $H(p)\}$ and let $A_{j(1)}$ be an edge of $H(p)$ that contains both $a_{i(1)}$ and a_0. Since u and v are not adjacent in $H(p)$, one has $i(1) < k$. Consider the u–v path $(a_0, A_{j(1)}, a_{i(1)}, \ldots, A_k, a_k)$. Let $i(2) = \max\{h > i(1) : a_h$ and $a_{i(1)}$ are adjacent in $H(p)\}$, and let $A_{j(1)}$ be any edge of $H(p)$ that contains both $a_{i(2)}$ and $a_{i(1)}$. If $i(2) = k$, then the u–v path $(a_0, A_{j(1)}, a_{i(1)}, A_{j(2)}, a_k)$ is simple since $a_k \notin A_{j(1)}$ and $a_0 \notin A_{j(2)}$. If $i(2) < k$, then let $i(3) = \max\{h > i(2) : a_h$ and $a_{i(2)}$ are adjacent in $H(p)\}$, and let $A_{j(3)}$ be any edge of $H(p)$ that contains both $a_{i(3)}$ are $a_{i(2)}$. If $i(3) = k$, then the u–v path $(a_0, A_{j(1)}, a_{i(1)}, A_{j(2)}, a_{i(2)}, A_{j(3)}, a_k)$ is simple since $a_k \notin A_{j(2)}$ and $a_{i(1)} \notin A_{j(3)}$, and so on. □

The next lemma characterizes sp-convex sets.

Lemma 5.2. *A subset X of $V(H)$ is* sp*-convex if and only if either $|X| \leq 1$ or, for every two distinct vertices u and v in X, there exists no X–X path joining u and v in the partial hypergraph H_{uv} of H obtained by deleting the edges that contain both u and v.*

Proof. (*only if*) Assume that X is sp-convex and $|X| > 1$. Suppose, by contradiction, that there exist two vertices u and v in X and an X–X path joining u and v in H_{uv}. By construction of H_{uv}, the vertices u and v are not adjacent in H_{uv} so that by Lemma 5.1, there exists a simple X–X path p' of length at least 2 joining u and v in H_{uv}. Since p' is also a simple path in H and $V(p')$ is not contained in X, X is not sp-convex (contradiction).

(*if*) If $|X| \leq 1$ then X is trivially sp-convex. Assume that $|X| > 1$ and, for every two distinct vertices u and v in X, there exists no X–X path joining u and v in H_{uv}. Suppose, by contradiction, that X is not sp-convex. Then, there exist two vertices u and v in X and a simple u–v path p in H with $V(p) \setminus X \neq \emptyset$. Let w be a vertex in $V(p) \setminus X$, and let p' be the subpath of p such that w is on p' and p' is an X–X path. Of course, p' is a simple path in H and has length at least 2. Let x and y be the vertices (in X) that are joined by p'. Since p' is a simple path in H, no edge of $H(p')$ contains both x and y, so that p' is also a path in H_{xy}. To sum up, p' is an X–X path that in H_{xy} joins the vertices x and y in X so that X is not sp-convex (contradiction). □

Example 5.3. Let $H = \{A_1, A_2, A_3, A_4, A_5\}$, where $A_1 = \{1\}$, $A_2 = \{1,2\}$, $A_3 = \{1,2,3\}$, $A_4 = \{3\}$, $A_5 = \{3,4\}$, and $A_6 = \{3,4,5\}$. The hypergraph H is shown in Figure 5. Let $X = \{1,3,4\}$. Consider the three vertex pairs in X. For the vertex pair $\{1,3\}$, the partial hypergraph H_{13} of H is $\{A_1, A_2, A_4, A_5, A_6\}$ and, since 1 and 3 belong to distinct connected components of H_{13}, there exists no X–X path joining 1 and 3 in H_{13}. For the vertex pair $\{1,4\}$, the partial

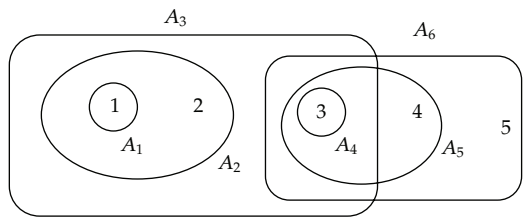

Figure 5

hypergraph H_{14} of H is H itself and, since every path joining 1 and 4 in H_{14} passes through 3, there exists no X–X path joining 1 and 4 in H_{14}. For the vertex pair $\{3,4\}$, the partial hypergraph H_{34} of H is $\{A_1, A_2, A_3, A_4\}$ and, since 4 is not a vertex of H_{14}, there exists no X–X path joining 3 and 4 in H_{34}. By Lemma 5.2, the set X is sp-convex, which is confirmed by the fact that the only simple paths joining two vertices in X are $(1, A_3, 3)$, $(1, A_3, 3, A_5, 4)$, $(1, A_3, 3, A_6, 4)$, $(3, A_5, 4)$, and $(3, A_6, 4)$.

We first prove the inclusion $\mathrm{sp}(H) \subseteq c(H)$.

Theorem 5.4. *Let H be a connected hypergraph. Every* sp*-convex set of H is c-convex.*

Proof. Suppose, by contradiction, that there exists an sp-convex set X of H that is not c-convex. Then, there exists an X-component K of H such that the boundary of X with K (i.e., the set $X \cap V(K)$) is not a partial edge of H. Of course, $|X \cap V(K)| \geq 2$ and the boundary of X with K is not a partial edge of K. Let A' be an edge of K such that for every edge A of K, either $X \cap A \subseteq A'$ or there exists $u \in (X \cap A') \setminus A$. Since the boundary of X with K is not a partial edge of K, $X \cap A'$ is a proper subset of $X \cap V(K)$; therefore, there exist two vertices u_0 and v such that $u_0 \in X \cap A'$ and $v \in (X \cap V(K)) \setminus A'$. Let A'' be any edge of K that contains v. Since K is an X-component of H, $A' \equiv_X A''$, and, hence, there exists an X–X path $p_0 = (u_0, A', v_1, \ldots, A_k, v)$ in K of length at least 2 (i.e., $k \geq 2$) with $A_k = A''$. If no edge of $H(p_0)$ contains both u_0 and v, then X is not sp-convex by Lemma 5.2 and a contradiction arises. Otherwise, let $A_{h(1)}$ be the first edge on p_0 that contains both u_0 and v. Since $v \in X \cap A_{h(1)}$ and $v \notin A'$, $X \cap A_{h(1)}$ is not a subset of A' so that there exists $u_1 \in (X \cap A') \setminus A_{h(1)}$. Consider the X–X path $p_1 = (u_1, A', v_1, \ldots, A_{h(1)}, v)$. If no edge of $H(p_1)$ contains both u_1 and v, then X is not sp-convex by Lemma 5.2 and a contradiction arises. Otherwise, let $A_{h(2)}$ be the first edge on p_1 that contains both u_1 and v. Since $v \in X \cap A_{h(2)}$ and $v \notin A'$, $X \cap A_{h(2)}$ is not a subset of A' so that there exists $u_2 \in (X \cap A') \setminus A_{h(2)}$. Consider the X–X path $p_2 = (u_2, A', w_0, \ldots, A_{h(2)}, v)$. If no edge of $H(p_2)$ contains both u_2 and v, then X is not sp-convex by Lemma 5.2 and a contradiction arises, and so on. Thus, ultimately one obtains an X–X path p of length at least 2 that joins two vertices u and v in X and is such that no edge of $H(p)$ contains both u and v. By Lemma 5.2, X is not sp-convex and a contradiction arises. □

We now characterize the class of hypergraphs H for which $\mathrm{sp}(H) = c(H)$.

Theorem 5.5. *Let H be a connected hypergraph. The equality* $\mathrm{sp}(H) = c(H)$ *holds if and only if H is γ-acyclic.*

Proof. (*if*) Assume that H is γ-acyclic and suppose, by contradiction, that $\mathrm{sp}(H) \neq c(H)$. Let X be a c-convex set that is not sp-convex. Then, there exist two vertices u and v in X and a simple

u–v path p (of length at least 2) such that $V(p) \setminus X \neq \emptyset$. Let $p = (a_0, A_1, a_1, A_2, \ldots, a_{k-1}, A_k, a_k)$, let h be such that $a_h \notin X$, let $i = \max\{l : 0 \leq l < h \text{ and } a_l \in X\}$, and let $j = \min\{l : h < l \leq k \text{ and } a_l \in X)$. Consider the subpath $p' = (a_i, A_{i+1}, a_{i+1}, \ldots, A_j, a_j)$ of p. Since p is a simple path, p' is a simple path of length at least 2, and $V(p') \cap X = \{a_i, a_j\}$ by construction. Let K be the X-component of H such that $H(p') \subseteq K$. Then, both a_i and a_j belong to the boundary of X with K. Since X is c-convex, the boundary of X with K is contained in an edge of H, say A. Since p' is a simple path of length at least 2, and A contains both a_i and a_j, A is not an edge of $H(p')$ so that $c = (a_i, A_{i+1}, a_{i+1}, \ldots, A_j, a_j, A, a_i)$ is a cycle; moreover, since p' is a path of length at least 2, c has length at least 3. Distinguish two cases depending on whether or not $V(p') \subseteq A$.

Case 1 ($V(p') \subseteq A$). Then, $c' = (a_i, A_{i+1}, a_{i+1}, A_{i+2}, a_{i+2}, A, a_i)$ is a cycle of length 3 and, since only the vertex a_{i+1} belongs to the three edges of c',c' is a γ-cycle (contradiction).

Case 2 ($V(p') \setminus A \neq \emptyset$). Then, there exists in $V(p')$ a vertex $a_r \notin A$ for some r, $i < r < j$. Let $i^* = \max\{l : i \leq l < r \text{ and } a_l \in A\}$ and $j^* = \min\{l : r < l \leq j \text{ and } a_l \in A\}$. Then $c' = (a_{i^*}, A_{i^*+1}, a_{i^*+1}, \ldots, a_{j^*-1}, A_{j^*}, a_{j^*}, A, a_{i^*})$ is a cycle of length at least 3 and, since every vertex in $V(c')$ belongs to exactly two edges of $H(c')$, c' is a γ-cycle (contradiction).

(*only if*) Assume that every c-convex set of H is also sp-convex and suppose, by contradiction, that H is not γ-acyclic. Let $c = (a_0, A_1, a_1, A_2, \ldots, a_{k-1}, A_k, a_0)$, $k \geq 3$, be a γ-cycle. Distinguish two cases depending on whether or not each vertex in $V(c)$ belongs to exactly two edges of $H(c)$.

Case 1. Each vertex in $V(c)$ belongs to exactly two edges of $H(c)$. Thus, c is a γ-cycle. Let $X = \{a_1, a_2\}$. Since $X \subseteq A_2$, for each X-component K of H, $X \cap V(K)$ is a partial edge of H and, hence, X is c-convex. On the other hand, since each vertex in $V(c)$ belongs to exactly two edges of $H(c)$, $(a_2, \ldots, A_k, a_0, A_1, a_1)$ is simple path of length at least 2 and, since $a_0 \notin X$, X is not sp-convex (contradiction).

Case 2. There exists a vertex in $V(c)$ that belongs to more than two edges of $H(c)$. Without loss of generality, let it be a_0. Since c is a γ-cycle, each a_h, $h \neq 0$, belongs to exactly two edges of $H(c)$. Let A_i be any edge of $H(c) \setminus \{A_1, A_k\}$ containing a_0, and let $X = \{a_{i-1}, a_i\}$. Since $X \subseteq A_i$, for each X-component K of H, $X \cap V(K)$ is a partial edge of H and, hence, X is c-convex. Let

$$r = \max\{h : 1 \leq h < i,\ a_0 \in A_h\} \qquad s = \min\{h : i < h \leq k,\ a_0 \in A_h\}. \tag{5.1}$$

It is easy to see that $(a_i, A_{i+1}, a_{i+1}, \ldots, A_s, a_0, A_r, a_r, \ldots, A_{i-1}, a_{i-1})$ is a simple path of length at least 2 and, since $a_0 \notin X$, X is not sp-convex (contradiction). □

Let H be a connected hypergraph with n vertices and m edges. If H is γ-acyclic, then H is β-acyclic and, hence, sp-convex hulls can be computed in $O(mn^2)$ using the Anstee-Farber algorithm. On the other hand, if H is γ-acyclic, then H is α-acyclic and, hence, c-convex hulls can be computed in linear time using the Tarjan-Yannakakis algorithm. By Theorem 5.5, sp-convex hulls can be computed in linear time, that is, more efficiently than using the Anstee-Farber algorithm.

5.2. Equivalence between sp-Convexity and m-Convexity

Note that, by Theorem 3.5 and by the fact that every chordless path in H is a simple path in H, one always has $\mathrm{sp}(H) \subseteq m(H)$. The following is another convexity-theoretic characterization of γ-acyclic hypergraphs.

Theorem 5.6. *A connected hypergraph H is γ-acyclic if and only if* $\mathrm{sp}(H) = m(H)$.

Proof. (*only if*) By hypothesis, H is γ-acyclic so that $\mathrm{sp}(H) = c(H)$ by Theorem 5.5. Moreover, by Proposition 2.4, H is conformal and, hence, weakly conformal, so that $c(H) = m(H)$ by Theorem 4.2. To sum up, $\mathrm{sp}(H) = c(H) = m(H)$.

(*if*) By hypothesis, $\mathrm{sp}(H) = m(H)$. Since $\mathrm{sp}(H) \subseteq c(H)$ by Theorem 5.5 and $c(H) \subseteq m(H)$, one has $\mathrm{sp}(H) = c(H)$ so that H is γ-acyclic again by Theorem 5.5. □

By Theorem 5.6, m-convex hulls can be computed in linear time more efficiently than using the monophonic-closure algorithm.

6. ap-Convexity versus sp-Convexity, c-Convexity, and m-Convexity

Let H be a connected hypergraph. In this section we characterize the three classes of hypergraphs H for which $\mathrm{ap}(H) = \mathrm{sp}(H)$, $\mathrm{ap}(H) = c(H)$, and $\mathrm{ap}(H) = m(H)$. To achieve this, we first give a polynomial algorithm for computing ap-convex hulls.

6.1. Computing ap-Convex Hulls

We represent H by its incidence graph $G(H)$.

Remark 6.1. For every two vertices u and v of H, every u–v path in H is a u–v path in $G(H)$ and *vice versa*; moreover, every cycle in H is a cycle in G and *vice versa*.

To avoid ambiguities, we call the vertices and edges of $G(H)$ the *nodes* and *arcs* of $G(H)$, respectively. A node v of $G(H)$ is a *vertex-node* or an *edge-node* depending on whether $v \in V(H)$ or $v \in H$. Moreover, we call *cutpoints* the cut vertices of $G(H)$; furthermore, a cutpoint of $G(H)$ is a *vertex-cutpoint* or an *edge-cutpoint* depending on whether it is a vertex-node or an edge-node. Note that, if a is a vertex-cutpoint of $G(H)$, then either the induced subhypergraph $H \setminus \{a\}$ is not connected (see the vertex-node 3 in Figure 6) or the singleton $\{a\}$ is an edge of H (see the vertex-node 1 in Figure 6); moreover, if A is an edge-cutpoint of $G(H)$, then either the partial hypergraph $H \setminus \{A\}$ is not connected (see the edge-node A_3 in Figure 6) or there exist one or more vertices of H that belong to A and to no other edge of H (see the edge-node A_6 in Figure 6). Our algorithm works with the "block-cutpoint tree" of $G(H)$, which is defined as follows. Let T be the bipartite graph whose nodes are the cutpoints and blocks of $G(H)$ and where (v, B) is an arc if the cutpoint v of $G(H)$ is a node of the block B of $G(H)$. A node of T is a *block-node* if it is a block of $G(H)$ and a *cutpoint-node* otherwise. It is well known [14] that T is a tree, which is called the *block cut-vertex tree* of $G(H)$. We also label each block-node B of T by the vertex set $V(B)$.

Example 6.1. Consider again the hypergraph $H = \{A_1, A_2, A_3, A_4, A_5, A_6\}$ of Example 5.3 (see Figure 5). The incidence graph $G(H)$ of H is shown in Figure 6, in which the cutpoints of $G(H)$ are circled.

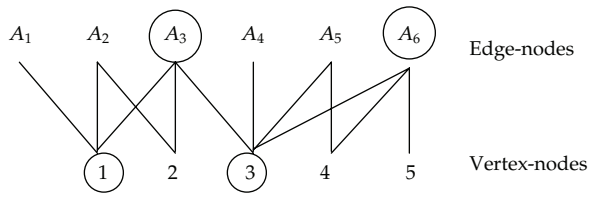

Figure 6

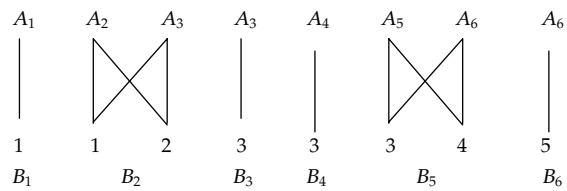

Figure 7

Note that the induced subhypergraph $H \setminus \{1\}$ is connected and the induced subhypergraph $H \setminus \{3\}$ is not connected; moreover, the partial hypergraph $H \setminus \{A_3\}$ is not connected and the partial hypergraph $H \setminus \{A_5\}$ is connected.

The blocks of $G(H)$ are reported in Figure 7.

The block-cutpoint tree T of $G(H)$ is shown in Figure 8.

The six block-nodes of T are labeled as follows:

$$\begin{array}{cc} \textit{block-node} & \textit{label} \\ B_1 & V(B_1) = \{1\} \\ B_2 & V(B_2) = \{1,2\} \\ B_3 & V(B_3) = \{3\} \\ B_4 & V(B_4) = \{3\} \\ B_5 & V(B_5) = \{3,4\} \\ B_6 & V(B_6) = \{5\}. \end{array} \tag{6.1}$$

Note that each leaf of T is a block-node. Moreover, if H is not a trivial hypergraph and T is a one-point tree, then the node of T is a nontrivial block of $G(H)$. Finally, if B is a trivial block of $G(H)$ with $H(B) = \{A\}$ and $V(B) = \{u\}$, then the block-node B of T is not a leaf if and only if u and A are both cutpoints of $G(H)$.

Algorithm 1 constructs the ap-convex hull (Y) of any subset X of $V(H)$.

Example 6.1 (continued). When we apply Algorithm 1 with input $X = \{1,3\}$, the tree T' resulting from the pruning of T is shown in Figure 9. So, the output of Algorithm 1 is $Y = \{1,2,3\}$.

When we apply Algorithm 1 with input $X = \{1,4\}$, the tree T' resulting from the pruning of T is shown in Figure 10. So, the output of Algorithm 1 is $Y = \{1,2,3,4\}$.

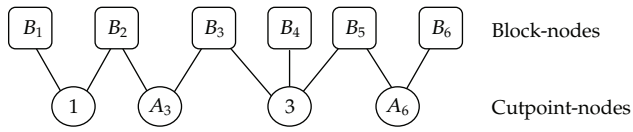

Figure 8

(1) If $|X| \leq 1$ then set $Y := X$ and Exit.
(2) Prune T by repeatedly removing a leaf if
  (i) either it is a cutpoint-node
  (ii) or it is a block-node, say B, and
    either $X \cap V(B) = \emptyset$
    or $|X \cap V(B)| = 1$ and the vertex in $X \cap V(B)$ is the cutpoint-node adjacent to B.
  Let T' be the resultant tree.
(3) Set $Y := \emptyset$. For each block-node B of T', set $Y := Y \bigcup V(B)$.

Algorithm 1

Remark 6.2. Each time a block-node B is deleted during the pruning process, either $X \cap V(B) = \emptyset$ or $|X \cap V(B)| = 1$ and the vertex in $X \cap V(B)$ belongs to $V(B')$ for some undeleted block-node B'. Therefore, one has that X is a subset of Y.

Fact 1. Each leaf of T' is a block-node and, if B is a leaf of T', then $V(B) \cap X \neq \emptyset$. If $|X| > 1$ and T' is a one-point tree, then the block-node of T' is a nontrivial block of $G(H)$.

Fact 2. Let B be a block-node of T' that is a trivial block of $G(H)$ with $H(B) = \{A\}$ and $V(B) = \{u\}$. If $|X| > 1$, then T' is not a one-point tree and, furthermore,

(i) if $u \in X$ then A is a node of T' adjacent to B (see the cutpoint-node A_3 and the block- node B_3 in Figure 7) for, otherwise, the block-node B would be a leaf of T' and u would be the only node of T' adjacent to B and, since $V(B) \cap X = \{u\}$, the leaf B of T' would have been deleted at step (1),

(ii) if $u \notin X$ then B is not a leaf of T' (see the block-node B_3 in Figure 10) for, otherwise, since $V(B) \cap X = \emptyset$, the block-node B would have been deleted at step (1).

Fact 3. If a cutpoint-node of T' is a vertex-node of $G(H)$, say u (see the cutpoint-node 3 in Figure 8), then there exist two leaves B and B' of T' such that every B–B' path in T' passes through the node u. Furthermore, since $V(B) \cap X \neq \emptyset$ (by Fact 1), there exists a vertex $v \neq u$ in $X \cap V(B)$ for, otherwise, $X \cap V(B) = \{u\}$ and, since u is the node adjacent to the leaf B, the leaf B of T' would have been deleted at step 1. Analogously, there exists a vertex $v' \neq u$ in $X \cap V(B')$. Finally, every v–v' path in $G(H)$ passes through u.

Theorem 6.3. *Let H be a connected hypergraph and let X be a subset of $V(H)$. Algorithm 1 correctly computes the* ap-*convex hull of X.*

Proof. Let Y be the output of Algorithm 1 and let $\langle X \rangle_{\text{ap}}$ denote the ap-convex hull of X. If $|X| \leq 1$, then $Y = X$ (see step (2)) and $\langle X \rangle_{\text{ap}} = X$, which proves the statement. Assume that $|X| > 1$. We first prove that $\langle X \rangle_{\text{ap}} \subseteq Y$ and, then, $Y \subseteq \langle X \rangle_{\text{ap}}$.

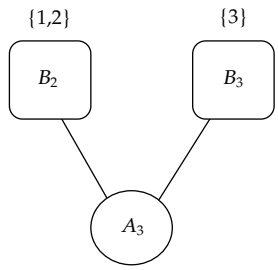

Figure 9

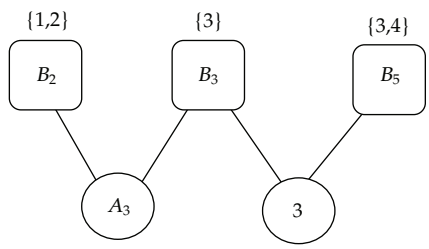

Figure 10

Proof of $\langle X\rangle_{\mathrm{ap}} \subseteq Y$. By Remark 6.2, it is sufficient to prove that Y is ap-convex. If $Y = V(H)$ then the statement is trivially true. Assume that $Y \neq V(H)$ and let u be any vertex in $V(H)\setminus Y$. Since $u \notin Y$, u belongs to $V(B)$ for some block-node B of T that has been deleted during the pruning process. Let B^* be the last of such block-nodes and let n be the cutpoint-node adjacent to B^* that is in every path joining B^* and every node of the tree T' resulting from the pruning process. Note that if n is a vertex-cutpoint of $G(H)$ then, since B^* is the block-node with $u \in V(B^*)$ that is last deleted, one has $n \neq u$ even if u is cutpoint-node of T (see Figure 11). Therefore, every path in $G(H)$ joining u to any vertex in Y must pass through n and, hence, no path in $G(H)$ joining two vertices in Y can pass through u and, by Remark 6.1, the same holds in H, which proves that Y is ap-convex. □

Proof of $Y \subseteq \langle X\rangle_{\mathrm{ap}}$. Consider again the tree T' resulting from the pruning process. By Fact 1, each leaf of T' is a block-node. Let L be the set of the leaves of T'. Thus, for each $B \in L$, $|X \cap V(B)| \geq 1$ and, if $|X \cap V(B)| = 1$, then the vertex in $X \cap V(B)$ is not the cutpoint-node adjacent to B in T'. Let $X' = X \cap (\bigcup_{B\in L} V(B))$ and let $\langle X'\rangle_{\mathrm{ap}}$ denote the ap-convex hull of X'. We now prove that $Y \subseteq \langle X'\rangle_{\mathrm{ap}}$. Consider the subgraph G' of $G(H)$ formed by the block-nodes of T'. First of all, observe that any cutpoint of G' lies on some path joining two vertex-nodes in X'. It follows that, for every trivial block B, the vertex in $V(B)$ lies on some path joining two vertex-nodes in X'; moreover, for every nontrivial block B of G', there exist at least two nodes of B that lie on some path joining two vertex-nodes in X' and, by Proposition 2.1, each node of B lies on some path joining two vertex-nodes in X'. Therefore, every vertex-node of G' lies on some path joining two vertex-nodes in X', that is, $V(G') \subseteq \langle X'\rangle_{\mathrm{ap}}$. By step (2) of Algorithm 1, $V(G') = Y$ and, hence, $Y \subseteq \langle X'\rangle_{\mathrm{ap}}$. Finally, since $X' \subseteq X$, one has $\langle X'\rangle_{\mathrm{ap}} \subseteq \langle X\rangle_{\mathrm{ap}}$ so that $Y \subseteq \langle X\rangle_{\mathrm{ap}}$. □

□

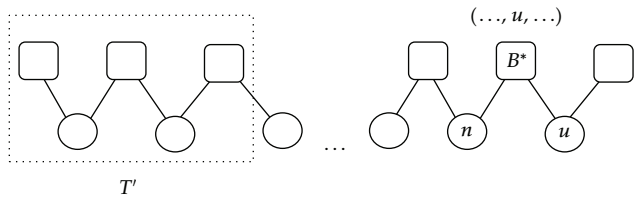

Figure 11

The following result, which is an immediate consequence of Theorem 6.3, will be used in the sequel.

Corollary 6.4. *Let H be a connected hypergraph. For every block B of $G(H)$, the subsets of $V(B)$ that are* ap*-convex are precisely the empty set, the singletons and $V(B)$.*

6.2. Equivalence between ap-Convexity and sp*-Convexity*

Recall from Section 3.2.4 that one always has $\mathrm{ap}(H) \subseteq \mathrm{sp}(H)$. We will prove that $\mathrm{ap}(H) = \mathrm{sp}(H)$ if and only if the incidence graph $G(H)$ of H satisfies the following two conditions:

(C1) every edge-cutpoint of $G(H)$ is a node of only trivial blocks of $G(H)$;

(C2) for every nontrivial block B of $G(H)$, and for every $X \subseteq V(B)$ with $1 < |X| < |V(B)|$, there exist two distinct vertex-nodes u and v in X and an X–X path joining u and v in the induced subgraph B_{uv} of B obtained by deleting the edge-nodes that are adjacent to both u and v.

Note that every graph satisfies (C1) owing to the fact that every edge contains exactly two vertices.

Remark 6.5. If (C1) holds then, for every nontrivial block B of $G(H)$, one has $|V(B)| \geq 3$ for, otherwise (i.e., if $|V(B)| = 2$), B would contain an edge-cutpoint of $G(H)$ (see Figure 12) owing to the fact that distinct edge-nodes of $G(H)$ have distinct sets of adjacent vertex-nodes.

In order to characterize the class of hypergraphs H for which $\mathrm{ap}(H) = \mathrm{sp}(H)$, we first restate Lemma 5.2 as follows.

Lemma 6.6. *Let G be the incidence graph of H. A subset X of $V(H)$ is* sp*-convex if and only if either $|X| \leq 1$ or, for every two distinct vertex-nodes u and v in X, there exists no X–X path joining u and v in the induced graph G_{uv} of $G(H)$ obtained by deleting the edge-nodes adjacent to both u and v.*

Corollary 6.7. *Let G be the incidence graph of H. If condition (C2) holds then, for every block B of $G(H)$, the subsets of $V(B)$ that are* sp*-convex are precisely the empty set, the singletons and $V(B)$.*

Proof. The empty set and the singletons are trivially sp-convex. Moreover, by Lemma 6.6 and condition (C2), for every nontrivial block of $G(H)$, no subset X of $V(B)$ with $1 < |X| < |V(B)|$ is sp-convex. □

Theorem 6.8. *Let H be a connected graph. The equality $\mathrm{ap}(H) = \mathrm{sp}(H)$ holds if and only if the incidence graph $G(H)$ of H satisfies both conditions (C1) and (C2).*

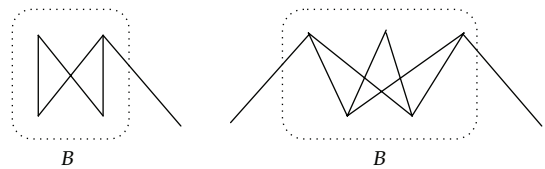

Figure 12

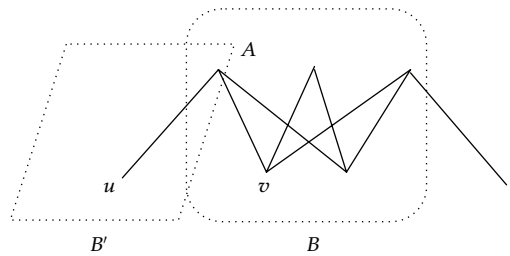

Figure 13

Proof. (*only if*) We first prove

(i) if $\mathrm{ap}(H) = \mathrm{sp}(H)$ then (C1)

and, then,

(ii) if $\mathrm{ap}(H) = \mathrm{sp}(H)$ then (C2).

Proof of (i). Suppose, by contradiction, that condition (C1) does not hold. Then, there exists an edge-cutpoint A of $G(H)$ that is a node of a nontrivial block B of $G(H)$. Let $v \in A \cap V(B)$, let $u \in A \setminus V(B)$, and let B' be the block containing both A and u (see Figure 13).

By Theorem 6.3, the ap-convex hull of $\{u, v\}$ is $V(B) \bigcup V(B')$, which is a proper superset of $\{u, v\}$ as B is a nontrivial block of $G(H)$. On the other hand, since (u, A, v) is the only simple u–v path in H, the set $\{u, v\}$ is sp-convex, which contradicts the hypothesis $\mathrm{ap}(H) = \mathrm{sp}(H)$. □

Proof of (ii). Let B be any nontrivial block of $G(H)$. Note that, since (C1) holds, by Remark 6.5 one has $|V(B)| \geq 3$. Let X be any subset of $V(B)$ with $1 < |X| < |V(B)|$. By Corollary 6.4, the set X is not ap-convex and, since $\mathrm{ap}(H) = \mathrm{sp}(H)$, X is not sp-convex. By Lemma 6.6, there exist two distinct vertex-nodes u and v in X and an X–X path joining u and v in G_{uv}. Since $X \subseteq V(B)$, by Proposition 2.2 every path joining u and v in G_{uv} is also a path in B_{uv}; therefore, there exists an X–X path joining u and v in B_{uv}, which proves that condition (C2) holds.

(*if*) Let X be any subset of $V(H)$ with $|X| > 1$, let $\langle X \rangle_{\mathrm{ap}}$ and $\langle X \rangle_{\mathrm{sp}}$ denote the ap-convex hull and the sp-convex hull of X, respectively. Of course, $X \subseteq \langle X \rangle_{\mathrm{sp}} \subseteq \langle X \rangle_{\mathrm{ap}}$. Therefore, it is sufficient to prove that $\langle X \rangle_{\mathrm{ap}} \subseteq \langle X \rangle_{\mathrm{sp}}$. By Theorem 6.3, $\langle X \rangle_{\mathrm{ap}}$ is the output of Algorithm 1, that is, if T' is the tree resulting from the pruning of the block-cutpoint tree of $G(H)$, then $\langle X \rangle_{\mathrm{ap}}$ is the union of the sets $V(B)$ for all block-nodes B of T'. So, we need to prove that, for every block-node B of T', one has $V(B) \subseteq \langle X \rangle_{\mathrm{sp}}$. Distinguish two cases depending on whether or not B is a trivial block of $G(H)$. □

Case 1. B is a trivial block of $G(H)$. Let u be the unique vertex in $V(B)$. If $u \in X$ then $u \in \langle X \rangle_{\mathrm{sp}}$ since $X \subseteq \langle X \rangle_{\mathrm{sp}}$. Consider now the case that $u \notin X$. Since $X \cap V(B) = \emptyset$, Fact 1, the block-node

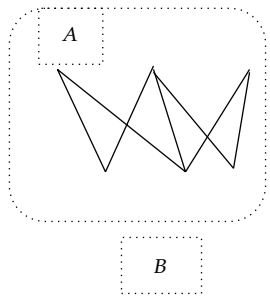

Figure 14

B of T' is not a leaf of T'. By Fact 2, u is one of the nodes of T' adjacent to the block-node B. By Fact 3, there exist two vertices v and v' in X such that every v–v' path in $G(H)$ passes through u. Since every v–v' path in H is a v–v' path in $G(H)$ (see Remark 6.1), every simple v–v' path in H passes through u so that u belongs to $\langle\{v, v'\}\rangle_{\mathrm{sp}}$ and, hence, to $\langle X\rangle_{\mathrm{sp}}$, which proves that $V(B) \subseteq \langle X\rangle_{\mathrm{sp}}$.

Case 2. B is a nontrivial block of $G(H)$. Consider first the case that T' is a one-point tree; thus, B is the only node of T' and $X \subseteq V(B)$. Since $|X| \geq 2$ and condition (C2) holds, $\langle X\rangle_{\mathrm{sp}} = V(B)$ by Corollary 6.7. Consider now the case that T' is not a one-point tree. Since B is a nontrivial block, by (C1) the cutpoint-nodes of T' that are adjacent to the block-node B are vertex-nodes of G. Let us distinguish two cases depending on whether or not B is a leaf of T'.

Subcase 1. B is a leaf of T'. Let u be the cutpoint-node of T' that is adjacent to B; thus, $u \in V(B)$. By Fact 3, there exist a vertex $v \neq u$ in $X \cap V(B)$ and a vertex v' in $X \cap V(B)$, where is another leaf of T', such that every v–v' path in $G(H)$ passes through u. Since every v–v'path in H is a v–v' path in $G(H)$ (see Remark 6.1), every simple v–v' path in H passes through u so that u belongs to $\langle\{v, v'\}\rangle_{\mathrm{sp}}$ and, hence, to $\langle X\rangle_{\mathrm{sp}}$. Since both u and v belong to $V(B) \cap \langle X\rangle_{\mathrm{sp}}$, one has $|V(B) \cap \langle X\rangle_{\mathrm{sp}}| \geq 2$. If one had $V(B) \setminus \langle X\rangle_{\mathrm{sp}} \neq \emptyset$ then, since condition (C2) holds, by Corollary 6.7 the set $V(B) \cap \langle X\rangle_{\mathrm{sp}}$ would not be sp-convex and, hence, there would exist a vertex in $V(B) \setminus \langle X\rangle_{\mathrm{sp}}$ that lies on some simple path joining two vertices in $V(B) \cap \langle X\rangle_{\mathrm{sp}}$ which contradicts the fact that the set $\langle X\rangle_{\mathrm{sp}}$ is sp-convex. Therefore, one has $V(B) \subseteq \langle X\rangle_{\mathrm{sp}}$.

Subcase 2. B is not a leaf of T'. Let u_1 and u_2 be two cutpoint-nodes of T' that are adjacent to B; thus, $\{u_1, u_2\} \in \mathrm{V}(B)$. By Fact 3, for each $i = 1, 2$ there exist two distinct vertices v_i and v_i' in X such that every v_i–v_i' path in $G(H)$ passes through u_i. Since every v_i–v_i' path in H is a v_i–v_i' path in $G(H)$ (see Remark 6.1), every simple v_i–v_i' path in H passes through u_i so that u_i belongs to $\langle\{v_i, v_i'\}\rangle_{\mathrm{sp}}$ and, hence, to $\langle X\rangle_{\mathrm{sp}}$. Since both u_1 and u_2 belong to both $V(B)$ and $\langle X\rangle_{\mathrm{sp}}$, one has $|V(B) \cap \langle X\rangle_{\mathrm{sp}}| \geq 2$. If one had $V(B) \setminus \langle X\rangle_{\mathrm{sp}} \neq \emptyset$, then, again by Corollary 6.7, the set $V(B) \cap \langle X\rangle_{\mathrm{sp}}$ would not be sp-convex and, hence, there would exist a vertex in $V(B) \setminus \langle X\rangle_{\mathrm{sp}}$ that lies on some simple path joining two vertices in $V(B) \cap \langle X\rangle_{\mathrm{sp}}$ which contradicts the fact that $\langle X\rangle_{\mathrm{sp}}$ is an sp-convex set. Therefore, one has $V(B) \subseteq \langle X\rangle_{\mathrm{sp}}$. □

6.3. Equivalence between ap-Convexity and c-Convexity

The next result characterizes the class of hypergraphs H for which $\mathrm{ap}(H) = c(H)$.

Theorem 6.9. *Let H be a connected hypergraph. The equality* $\mathrm{ap}(H) = c(H)$ *holds if and only if H is Berge-acyclic.*

Proof. By Remark 6.1, H is Berge-acyclic if and only if $G(H)$ is a tree.

(*only if*) Assume that $\mathrm{ap}(H) = c(H)$ and suppose, by contradiction, that H is not Berge-acyclic. Since $G(H)$ is not a tree, there exists a nontrivial block of $G(H)$, say B. Then there exists an edge-node A of B such that $1 < |A| < |V(B)|$ (see Figure 14). By Corollary 6.4, A is not ap-convex; but, since A is an edge of H, A is c-convex (contradiction).

(*if*) Assume that H is Berge-acyclic. Then, $G(H)$ is a tree and, hence, every block of $G(H)$ is trivial so that trivially conditions (C1) and (C2) of Theorem 6.8 hold and, hence, $\mathrm{ap}(H) = \mathrm{sp}(H)$. Moreover, since every Berge-acyclic hypergraph is γ-acyclic, H is γ-acyclic so that $\mathrm{sp}(H) = c(H)$ by Theorem 5.5. To sum up, one has $\mathrm{ap}(H) = \mathrm{sp}(H) = c(H)$, which proves the statement. □

If H is Berge-acyclic, then H is α-acyclic and, hence, c-convex hulls can be computed in linear time. By Theorem 6.9, ap-convex hulls can be computed in in linear time more efficiently than using Algorithm 1.

6.4. Equivalence between ap-Convexity and m-Convexity

The next theorem generalizes Theorem 3.2.

Theorem 6.10. *Let H be a connected hypergraph. The equality* $\mathrm{ap}(H) = m(H)$ *holds if and only if H is Berge-acyclic.*

Proof. If H is Berge-acyclic, then, by Proposition 2.4, H is conformal and, by Theorem 4.2, $c(H) = m(H)$. By Theorem 6.9, $\mathrm{ap}(H) = c(H)$ and, hence, $\mathrm{ap}(H) = m(H)$. On the other hand, if $\mathrm{ap}(H) = m(H)$ then, since $\mathrm{ap}(H) \subseteq c(H) \subseteq m(H)$, one has $\mathrm{ap}(H) = c(H)$ so that, by Theorem 6.9, H is Berge-acyclic. □

7. Future Research

We have considered the four notions of hypergraph convexity: m-convexity, c-convexity, sp-convexity, and ap-convexity, and for each pair of these convexities, we have characterized the class of hypergraphs in which the two convexities are equivalent. Another important notion of graph convexity is geodetic convexity (or g-convexity) [3]: a vertex set X in a graph is g-convex if, for every geodesic (i.e., minimum-length path) p joining two vertices in X, each vertex on p belongs to X. Of course, since every geodesic is a chordless path, every m-convex set is g-convex, that is, g-convexity is finer than m-convexity; moreover, g-convexity and m-convexity are equivalent in distance-hereditary graphs [3]. An open problem is the characterization of graphs (and, more in general, of hypergraphs) in which g-convexity and m-convexity are equivalent.

References

[1] J.-D. Boissonnat and M. Yvinec, *Algorithmic Geometry*, Cambridge University Press, 1999.
[2] M. J. L. Van de Vel, *Convexity Structures*, North-Holland, Amsterdam, 1997.
[3] M. Farber and R. E. Jamison, "Convexity in graphs and hypergraphs," *SIAM Journal on Algebraic and Discrete Methods*, vol. 7, pp. 433–444, 1986.

[4] F. M. Malvestuto, "Canonical and monophonic convexities in hypergraphs," *Discrete Mathematics*, vol. 309, no. 13, pp. 4287–4298, 2009.
[5] P. Duchet, "Convex sets in graphs, II. Minimal path convexity," *Journal of Combinatorial Theory, Series B*, vol. 44, no. 3, pp. 307–316, 1988.
[6] M. Changat, S. Klavžar, and H. M. Mulder, "The all-paths transit function of a graph," *Czechoslovak Mathematical Journal*, vol. 51, no. 2, pp. 439–448, 2001.
[7] E. Sampathkumar, "Convex sets in graphs," *Indian Journal of Pure and Applied Mathematics*, vol. 15, pp. 1065–1071, 1984.
[8] R. Fagin, "Degrees of acyclicity for hypergraphs and relational database schemes," *Journal of the ACM*, vol. 30, no. 3, pp. 514–550, 1983.
[9] G. Gottlob, N. Leone, and F. Scarcello, "The complexity of acyclic conjunctive queries," *Journal of the ACM*, vol. 48, no. 3, pp. 431–498, 2001.
[10] R. Dechter, "Constraint networks," in *Encyclopaedia of Artificial Intelligence*, pp. 276–285, Wiley, New York, NY, USA, 2nd edition, 1992.
[11] G. Gottlob, N. Leone, and F. Scarcello, "A comparison of structural CSP decomposition methods," *Artificial Intelligence*, vol. 124, no. 2, pp. 243–282, 2000.
[12] S. J. Haberman, *The Analysis of Frequency Data*, University of Chicago Press, Chicago, Ill, USA, 1974.
[13] S. L. Lauritzen, *Graphical Models*, Clarendon Press, Oxford, UK, 1996.
[14] F. Harary, *Graph Theory*, Addison-Wesley, New York, NY, USA, 1969.
[15] R. P. Anstee and M. Farber, "Characterizations of totally balanced matrices," *Journal of Algorithms*, vol. 5, no. 2, pp. 215–230, 1984.
[16] P. Duchet, "Hypergraphs," in *Handbook of Combinatorics. Volume I*, R. L. Graham, M. Grotschel, and L. Lovász, Eds., pp. 381–432, North-Holland, 1995.
[17] C. Beeri, R. Fagin, D. Maier, and M. Yannakakis, "On the desirability of acyclic database schemes," *Journal of the ACM*, vol. 30, no. 3, pp. 479–513, 1983.
[18] R. E. Tarjan and M. Yannakakis, "Simple linear-time algorithms to test chordality of graphs, test acyclicity of hypergraphs, and selectively reduce acyclic hypergraphs," *SIAM Journal on Computing*, vol. 13, no. 3, pp. 566–579, 1984.
[19] A. D'Atri and M. Moscarini, "On hypergraph acyclicity and graph chordality," *Information Processing Letters*, vol. 29, no. 5, pp. 271–274, 1988.
[20] S. L. Lauritzen, T. P. Speed, and K. Vijayan, "Decomposable graphs and hypergraphs," *Journal of the Australian Mathematical Society. Series A*, vol. 36, pp. 12–29, 1984.
[21] H. G. Leimer, "Optimal decomposition by clique separators," *Discrete Mathematics*, vol. 113, no. 1–3, pp. 99–123, 1993.
[22] F. M. Malvestuto and M. Moscarini, "A fast algorithm for query optimization in universal-relation databases," *Journal of Computer and System Sciences*, vol. 56, no. 3, pp. 299–309, 1998.
[23] F. M. Malvestuto and M. Moscarini, "Decomposition of a hypergraph by partial-edge separators," *Theoretical Computer Science*, vol. 237, no. 1-2, pp. 57–79, 2000.
[24] M. Changat and J. Mathew, "On triangle path convexity in graphs," *Discrete Mathematics*, vol. 206, no. 1–3, pp. 91–95, 1999.
[25] M. C. Dourado, F. Protti, and J. L. Szwarcfiter, "Algorithmic aspects of monophonic convexity," *Electronic Notes in Discrete Mathematics*, vol. 30, pp. 177–182, 2008.
[26] F. Ardila and E. Maneva, "Pruning processes and a new characterization of convex geometries," *Discrete Mathematics*, vol. 309, no. 10, pp. 3083–3091, 2009.
[27] B. Kannan and M. Changat, "Hull numbers of path convexities on graphs," in *Proceedings of the International Instructional Workshop on Convexity in Discrete Structures*, vol. 5 of *Lecture Notes Series*, pp. 11–23, Ramanujan Mathematical Society, 2008.
[28] D. Maier and J. D. Ullman, "Connections in acyclic hypergraphs," *Theoretical Computer Science*, vol. 32, no. 1-2, pp. 185–199, 1984.
[29] F. M. Malvestuto, M. Mezzini, and M. Moscarini, "Computing simple-path convex hulls in hypergraphs," *Information Processing Letters*, vol. 111, no. 5, pp. 231–234, 2011.