Gamma Acyclicity, Annotated Relations, and Consistency Witness Functions

Albert Atserias
Universitat Politècnica de Catalunya, Barcelona, Spain
Centre de Recerca Matemàtica, Bellaterra, Spain

Phokion G. Kolaitis
University of California Santa Cruz, CA, USA
IBM Research, San Jose, CA, USA

Abstract

During the early days of relational database theory it was realized that "acyclic" database schemas possess a number of desirable properties. In fact, three different notions of "acyclicity" were identified and investigated during the 1980s, namely, α-acyclicity, β-acyclicity, and γ-acyclicity. Much more recently, the study of α-acyclicity was extended to annotated relations, where the annotations are values from some positive commutative monoid. The recent results about α-acyclic schemas and annotated relations give rise to results about β-acyclic schemas and annotated relations, since a schema is β-acyclic if and only if every sub-schema of it is α-acyclic. Here, we study γ-acyclic schemas and annotated relations. Our main finding is that the characterization of γ-acyclic schemas in terms of monotone sequential join expression extends to annotated relations, provided the annotations come from a positive commutative monoid that has the inner consistency property. Furthermore, the results reported here shed light on the role of the join of two standard relations. Specifically, our results reveal that the only relevant property of the join of two standard relations is that it is a witness to the consistency of the two relations, provided that these two relations are consistent. For the more abstract setting of annotated relations, this property of the standard join is captured by the notion of a consistency witness function, a notion which we systematically utilize in this work.

2012 ACM Subject Classification Information systems → Relational database model; Theory of computation → Database theory

Keywords and phrases annotated relations, gamma-acyclicity, consistency witness functions

Digital Object Identifier 10.4230/LIPIcs.ICDT.2026.16

Funding *Albert Atserias*: Atserias was partially supported by grant PID2022-138506NB-C22 (PROOFS BEYOND) and the Severo Ochoa and María de Maeztu Program for Centers and Units of Excellence in R&D (CEX2020-001084-M) of the AEI, and the CERCA and ICREA Academia Programmes of the Generalitat de Catalunya.

1 Introduction

Annotated databases are databases in which each fact in a relation is annotated with a value from some algebraic structure. Starting with the influential work on database provenance [12, 13], there has been an extensive investigation of several different aspects of annotated databases, including the study of conjunctive query containment for annotated databases [11, 15] and the evaluation of Datalog programs on annotated databases [14]. In these investigations, the annotations are values from some fixed semiring $\mathbb{K} = (K, +, \times, 0, 1)$. Thus, standard relational databases are annotated databases in which the annotations are 1 (true) and 0 (false), while bag databases are annotated databases in which the annotations are non-negative integers denoting the multiplicities. This framework, which is often referred to as *semiring semantics*, has spanned first-order logic [9] and least fixed-point logic [6].

29th International Conference on Database Theory (ICDT 2026).
Editors: Balder ten Cate and Maurice Funk; Article No. 16; pp. 16:1–16:18
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

During the early days of relational database theory it was realized that “acyclic” database schemas possess a number of desirable semantic properties. In fact, three different notions of “acyclicity” were identified and extensively investigated during the 1980s, namely, acyclicity (also known as α-acyclicity), β-acyclicity, and γ-acyclicity. On undirected graphs (equivalently, on database schemas consisting of binary relation symbols only) these notions coincide with the notion of an acyclic graph, but they form a strict hierarchy on hypergraphs (equivalently, on arbitrary database schemas) with β-acyclicity being a stricter notion than acyclicity, and γ-acyclicity being a stricter notion than β-acyclicity.

The study of acyclic schemas was initiated by Yannakakis, who focused on the evaluation of acyclic joins [16]. After this, Fagin, Beeri, Maier, and Yannakakis [4] showed that acyclic schemas are precisely the ones possessing the *local-to-global consistency property*, that is, every collection of pairwise consistent relations $R_1, \dots, R_m$ over such schemas is globally consistent (i.e., there is a relation T whose projection on the attributes of R_i is equal to R_i, for $1 \leq i \leq m$). Fagin et al. [4] also characterized acyclicity in terms of the existence of *monotone sequential join expressions*, i.e., expressions of the form $(((\cdots(R_1 \bowtie R_2) \bowtie \cdots) \bowtie R_{m-1}) \bowtie R_m)$ with the property that if the relations $R_1, \dots, R_m$ are pairwise consistent, then every intermediate sequential join expression $((\cdots(R_1 \bowtie R_2) \bowtie \cdots) \bowtie R_{i-1})$ produces a relation that is consistent with the relation R_i. Results about acyclicity yield results about β-acyclicity, since a schema is β-acyclic if and only if every sub-schema of it is acyclic. Fagin [8] studied γ-acyclicity and showed that a schema is γ-acyclic if and only if every *connected* sequential join expression is monotone. Intuitively, this means that every sequential join expression is monotone, provided no join between relations with disjoint sets of attributes is allowed.

Atserias and Kolaitis [2] studied the interplay between local consistency and global consistency for annotated relations. Since the definition of consistency of annotated relations uses only the projection operation on relations and since projection is defined using only addition $+$, they considered $\mathbb{K}$-relations where the annotations come from a monoid $\mathbb{K} = (K, +, 0)$. They identified a condition on monoids, called the *inner consistency property*, and showed that a positive monoid $\mathbb{K} = (K, +, 0)$ has the inner consistency property if and only if every acyclic schema H has the local-to-global consistency property for $\mathbb{K}$-relations (i.e., every pairwise consistent collection of $\mathbb{K}$-relations over H is globally consistent). It was not clear, however, whether the results about acyclic schemas and sequential join expressions in [4] can be extended to annotated relations, since, as shown in [1], the analog of the standard join for bags need not be a witness to the consistency of two consistent bags. In a subsequent paper, Atserias and Kolaitis [3] introduced the notion of a *consistency witness function* on a positive monoid $\mathbb{K}$, which is a function W that, given two $\mathbb{K}$-relations R and S, returns a $\mathbb{K}$-relation $W(R, S)$ that is a consistency witness for R and S, provided that R and S are consistent $\mathbb{K}$-relations. They also introduced the notion of a *monotone sequential* c*-join expression*, which is analogous to that of a monotone sequential join expression with some arbitrary consistency witness function in place of the standard join. Using these notions, it was shown in [3] that the characterization of acyclicity in terms of monotone sequential join expressions in [4] extends to characterizations of acyclicity in terms of monotone sequential c-join expressions on monoids having the inner consistency property; furthermore, the inner consistency property itself can be characterized in such terms.

Here, we investigate γ-acyclic schemas and establish that the desirable semantic properties of γ-acyclic schemas extend to annotated relations. The two main results are as follows:

1. If $\mathbb{K}$ is a positive commutative monoid and H is a schema such that every connected sequential c-join-expression over H is monotone on $\mathbb{K}$ w.r.t. *some* consistency witness function on $\mathbb{K}$, then H is γ-acyclic.

2. If $\mathbb{K}$ is a positive commutative monoid that has the inner consistency property and H is a schema which is γ-acyclic, then every connected sequential c-join-expression over H is monotone on $\mathbb{K}$ w.r.t. *every* consistency witness function on $\mathbb{K}$.

As a byproduct of these two main results, we obtain a characterization of the inner consistency property in terms of γ-acyclicity and connected sequential c-join expressions. Furthermore, our work sheds light on the role of the join of two standard relations. Specifically, our results reveal that, in the study of the various notions of acyclicity in [8], the only relevant property of the join of two standard relations is that it is a witness to the consistency of the two relations, provided these two relations are consistent. In the setting of annotated relations, this property of the standard join is captured by the notion of a consistency witness function.

The rest of the paper is organized as follows. Section 2 contains the definitions of the basic notions, while Section 3 contains the definition of a consistency witness function and related notions. To make the paper as self-contained as possible, the earlier results about acyclic schemas are summarized in Section 4. Section 5 discusses β-acyclic schemas. Section 6 contains the main results about γ-acyclic schemas and annotated relations.

2 Basic Notions

Monoids. A *commutative monoid* is a structure $\mathbb{K} = (K, +, 0)$, where $+$ is a binary operation on the universe K of $\mathbb{K}$ that is associative, commutative, and has 0 as its neutral element, i.e., $p + 0 = p = 0 + p$ holds for all $p \in K$. A commutative monoid $\mathbb{K} = (K, +, 0)$ is *positive* if for all elements $p, q \in K$ with $p + q = 0$, we have that $p = 0$ and $q = 0$. From now on, we assume that all commutative monoids considered have at least two elements in their universe.

The following are examples of positive commutative monoids.

- The *Boolean monoid* $\mathbb{B} = (\{0, 1\}, \vee, 0)$ with disjunction $\vee$ as its operation and 0 (false) as its neutral element.
- The *bag monoid* $\mathbb{N} = (Z^{\geq 0}, +, 0)$, where $Z^{\geq 0}$ is the set of non-negative integers and $+$ is the standard addition operation. Note that the structure $\mathbb{Z} = (Z, +, 0)$, where Z is the set of integers, is a commutative monoid, but not a positive one.
- A *numerical semigroup* is a submonoid $\mathbb{K} = (K, +, 0)$ of the bag monoid $\mathbb{N} = (Z^{\geq 0}, +, 0)$, such that K is a cofinite set, i.e., the complement $Z^{\geq 0} \setminus K$ is finite. A concrete example of a numerical semigroup is $\mathbb{K} = (\langle 3, 5\rangle, +, 0)$, where $\langle 3, 5\rangle$ is the set of all non-negative integers of the form $3m + 5n$ with $m \geq 0$ and $n \geq 0$, i.e., $\langle 3, 5\rangle = \{0, 3, 5, 6, 8, 9, 10, \ldots\}$.
- The *power set monoid* $\mathbb{P}(A) = (\mathcal{P}(A), \cup, \emptyset)$, where if A is a set, then $\mathcal{P}(A)$ is its powerset, and $\cup$ is the union operation on sets.
- The structure $\mathbb{V} = ([0, 1], \max, 0)$, where $[0, 1]$ is the interval of all real numbers between 0 and 1, and max is the standard maximum operation.
- The structure $\mathbb{R}^{\geq 0} = ([0, \infty), +, 0)$, where $[0, \infty)$ is the set of all non-negative real numbers and $+$ is the standard addition operation.
- The structure $\mathbb{T} = ((-\infty, \infty], \min, \infty)$, where $(-\infty, \infty]$ is the set of all real numbers together with ∞, and min is the standard minimum operation.

$\mathbb{K}$-relations and marginals of $\mathbb{K}$-relations. An *attribute* A is a symbol with an associated set $\mathrm{Dom}(A)$ as its *domain*. If X is a finite set of attributes, then $\mathrm{Tup}(X)$ is the set of X-*tuples*, i.e., the set of functions that take each attribute $A \in X$ to an element of its domain $\mathrm{Dom}(A)$. $\mathrm{Tup}(\emptyset)$ is non-empty as it contains the *empty tuple*, i.e., the function with empty domain. If $Y \subseteq X$ and t is an X-tuple, then the *projection of t on Y*, denoted by $t[Y]$, is the unique Y-tuple that agrees with t on Y. In particular, $t[\emptyset]$ is the empty tuple.

Let $\mathbb{K} = (K, +, 0)$ be a positive commutative monoid and let X be a finite set of attributes.

- A $\mathbb{K}$-*relation over* X is a function $R : \text{Tup}(X) \to K$ that assigns a value $R(t)$ in K to every X-tuple t in $\text{Tup}(X)$. We will often write $R(X)$ to indicate that R is a $\mathbb{K}$-relation over X, and we will refer to X as the set of attributes of R. If X is the empty set of attributes, then a $\mathbb{K}$-relation over X is simply a function that assigns a single value from K to the empty tuple. Note that the $\mathbb{B}$-relations are the standard relations, while the $\mathbb{N}$-relations are the *bags* or *multisets*, i.e., each tuple has a non-negative integer associated with it that denotes the *multiplicity* of the tuple.
- The *support* $\text{Supp}(R)$ of a $\mathbb{K}$-relation $R(X)$ is the set of X-tuples t that are assigned non-zero value, i.e., $\text{Supp}(R) := \{t \in \text{Tup}(X) : R(t) \neq 0\}$. We will often write R' to denote $\text{Supp}(R)$. Note that R' is a standard relation over X. A $\mathbb{K}$-relation is *finitely supported* if its support is a finite set. In this paper, all $\mathbb{K}$-relations considered will be finitely supported, and we omit the term; thus, from now on, a $\mathbb{K}$-relation is a finitely supported $\mathbb{K}$-relation. When R' is empty, we say that R is the empty $\mathbb{K}$-relation over X.
- If $Y \subseteq X$, then the *marginal* $R[Y]$ *of* R *on* Y is the $\mathbb{K}$-relation over Y such that for every Y-tuple t, we have that

$$R[Y](t) := \sum_{\substack{r \in R': \\ r[Y]=t}} R(r). \tag{1}$$

The value $R[Y](t)$ is the *marginal of* R *over* t. For notational simplicity, we will often write $R(t)$ for the marginal of R over t, instead of $R[Y](t)$. It will be clear from the context (e.g., from the arity of the tuple t) if $R(t)$ is indeed the marginal of R over t (in which case t must be a Y-tuple) or $R(t)$ is the actual value of R on t as a mapping from $\text{Tup}(X)$ to K (in which case t must be an X-tuple). Note that if R is a standard relation (i.e., R is a $\mathbb{B}$-relation), then the marginal $R[Y]$ is the projection of R on Y.

The proof of the next useful proposition follows easily from the definitions.

▶ **Proposition 1.** *Let $\mathbb{K}$ be a positive commutative monoid and let $R(X)$ be a $\mathbb{K}$-relation. Then the following hold:*

1. *For all $Y \subseteq X$, we have $R'[Y] = R[Y]'$.*
2. *For all $Z \subseteq Y \subseteq X$, we have $R[Y][Z] = R[Z]$.*

If X and Y are sets of attributes, then we write XY as shorthand for the union $X \cup Y$.

Schemas and hypergraphs.

- A *schema* is a sequence $X_1, \ldots, X_m$ of non-empty sets of attributes.
- A *hypergraph* is a pair $H = (V, F)$, where V is a finite non-empty set and F is a set of non-empty subsets of V. We call V the set of the *nodes* of H and we call F the set of the *hyperedges* of H.

A schema $X_1, \ldots, X_m$ can be identified with the hypergraph $H = (\bigcup_{i=1}^{m} X_i, \{X_1, \ldots, X_m\})$, i.e., the nodes of H are the attributes and the hyperedges of H are the members $X_1, \ldots, X_m$ of the schema. We will use the terms *schema* and *hypergraph* interchangeably.

- A *collection of* $\mathbb{K}$-*relations* over a schema $X_1, \ldots, X_m$ is a sequence $R_1(X_1), \ldots, R_m(X_m)$ such that each $R_i(X_i)$ is a $\mathbb{K}$-relation over X_i.

3 Consistent Relations and Consistency Witness Functions

Let $\mathbb{K} = (K, +, 0)$ be a positive commutative monoid.

- Two $\mathbb{K}$-relations $R(X)$ and $S(Y)$ are *consistent* if there is a $\mathbb{K}$-relation $T(XY)$ such that $T[X] = R$ and $T[Y] = S$. Such a $\mathbb{K}$-relation T is a *consistency witness* for R and S.

- A collection $R_1(X_1), \dots, R_m(X_m)$ of $\mathbb{K}$-relations over a schema $X_1, \dots, X_m$ is *globally consistent* if there is a $\mathbb{K}$-relation $T(X_1 \dots X_m)$ such that $T[X_i] = R_i$, for i with $1 \leq i \leq m$. Such a $\mathbb{K}$-relation T is a *consistency witness* for $R_1, \dots, R_m$.

Note that if $R_1(X_1), \dots, R_m(X_m)$ is a globally consistent collection of $\mathbb{K}$-relations, then the relations $R_1(X_1), \dots, R_m(X_m)$ are pairwise consistent. Indeed, if T is a consistency witness for $R_1(X_1), \dots, R_m(X_m)$, then for all i and j with $1 \leq i, j \leq m$, we have that the $\mathbb{K}$-relation $T[X_iX_j]$ is a consistency witness for R_i and R_j, because

$$R_i = T[X_i] = T[X_iX_j][X_i] \quad \text{and} \quad R_j = T[X_j] = T[X_iX_j][X_j],$$

where, in each case, the first equality follows from the definition of global consistency and the second equality follows from Proposition 1.

It is well known that the converse fails even for standard relations, i.e., there are standard relations that are pairwise consistent but not globally consistent. The main result by Beeri et al. [4] characterizes the schemas for which the pairwise consistency of a collection of standard relations implies that they are globally consistent. Quite recently, this result was extended to $\mathbb{K}$-relations over positive monoids that satisfy a condition called the *inner consistency property* [2]. We will discuss this extension in the next section. For now, we consider the following notion, which was introduced and studied in [3].

- A *consistency witness function on* $\mathbb{K}$ is a function W that takes as arguments two $\mathbb{K}$-relations $R(X)$ and $S(Y)$, and returns as value a $\mathbb{K}$-relation $W(R, S)$ over XY such that if R and S are consistent $\mathbb{K}$-relations, then $W(R, S)$ is a consistency witness for R and S.

▶ **Example 2.** We now give examples of consistency witness functions for several different positive monoids.

1. If $\mathbb{B} = (\{0, 1\}, \vee, 0)$ is the Boolean monoid, then an example of a consistency witness function on $\mathbb{B}$ is the standard join operation $\bowtie$ on standard relations.
2. If $\mathbb{P}(A) = (\mathcal{P}(A), \cup, \emptyset)$ is a powerset monoid, then an example of a consistency witness function on $\mathbb{P}(A)$ is the function W such that for every XY-tuple t, we have that
$$W(R, S)(t) = R(t[X]) \cap S(t[Y]).$$
3. If $\mathbb{V} = ([0, 1], \max, 0)$, then an example of a consistency witness function on $\mathbb{V}$ is the function W such that for every XY-tuple t, we have that
$$W(R, S)(t) = \min\{R(t[X]), S(t[Y])\}.$$
4. If $\mathbb{R}^{\geq 0} = ([0, \infty), +, 0)$, then an example of a consistency witness function on $\mathbb{R}^{\geq 0}$ is the function W such that for every XY-tuple t, we have that
$$W(R, S)(t) = \begin{cases} (R(t[X]) \times S(t[Y]))/R(t[X \cap Y]), & \text{if } R(t[X \cap Y]) \neq 0 \\ 0, & \text{if } R(t[X \cap Y]) = 0. \end{cases}$$
5. If $\mathbb{T} = ((-\infty, \infty], \min, \infty)$, then an example of a consistency witness function on $\mathbb{T}$ is the function W such that for every XY-tuple t, we have that
$$W(R, S)(t) = \begin{cases} R(t[X]) + S(t[Y]) - R(t[X \cap Y]), & \text{if } R(t[X \cap Y]) \neq \infty \\ \infty, & \text{if } R(t[X \cap Y]) = \infty. \end{cases}$$
(in the previous two expressions, $\times$, $/$, $+$, and $-$ are the standard multiplication, division, addition, and subtraction operations on the real numbers).

As shown in [2], the witness functions in the first three preceding examples are instances of the *standard join* $\bowtie_{\mathbb{K},\mathrm{S}}$ on a monoid $\mathbb{K}$ that can be expanded to a bounded distributive lattice. For instance, the powerset monoid $\mathbb{P}(A) = (\mathcal{P}(A), \cup, \emptyset)$ can be expanded to the bounded distributive lattice $(\mathcal{P}(A), \cup, \cap, \emptyset, A)$. Furthermore, the witness functions in the last two preceding examples are instances of the *Vorob'ev join* $\bowtie_{\mathbb{K},\mathrm{V}}$ on a monoid $\mathbb{K}$ that can be expanded to a semifield. For instance, the monoid $\mathbb{T} = ((-\infty, \infty], \min, \infty)$ can be expanded to the *tropical* semifield $((-\infty, \infty], \min, +, -, \infty, 0)$.

Note that the bag monoid $\mathbb{N} - (N, +, 0)$ was not included in the preceding examples. As pointed out in [1], the bag-join operation of bags (the analog of the standard join for bags) is not a consistency witness function on the bag monoid $\mathbb{N}$. Furthermore, the analog of the Vorob'ev join is not well-defined on $\mathbb{N}$. A consistency witness function for $\mathbb{N}$ can be obtained via an iterative algorithm, called the *northwest corner method* (see [2] for details).

A *join expression* is an expression involving standard relations and applications of the join operation $\bowtie$ on standard relations [4, 8]. In [3], the notion of a c-*join expression* and its variants were introduced as a generalization of the notion of join expression to arbitrary consistency witness functions and $\mathbb{K}$-relations. The precise definitions are as follows.

Let $X_1, \ldots, X_m$ be a schema and let $\bowtie_{\mathrm{c}}$ be a binary function symbol, which will be interpreted by some consistency witness function.

- The collection of c-*join expressions over* $X_1, \ldots, X_m$ is the smallest collection of strings that contains each X_i and has the property that if E_1 and E_2 are in the collection, then also the string $(E_1 \bowtie_{\mathrm{c}} E_2)$ is in the collection.
- The collection of *sequential* c-*join expressions over* $X_1, \ldots, X_m$ is the smallest collection of strings that contains each X_i and has the property that if E is in the collection and X is one of the X_i's, then also the string $(E \bowtie_{\mathrm{c}} X)$ is in the collection.

Note that a (sequential) c-join expression over $X_1, \ldots, X_m$ need not contain every set X_i.

Clearly, the string $((X_1 \bowtie_{\mathrm{c}} X_2) \bowtie_{\mathrm{c}} X_3)$ is a sequential c-join-expression, while the string $((X_1 \bowtie_{\mathrm{c}} X_2) \bowtie_{\mathrm{c}} (X_3 \bowtie_{\mathrm{c}} X_4))$ is a c-join expression, but not a sequential one.

Semantics to c-join expressions are assigned in a straightforward way as follows.

Let $X_1, \ldots, X_m$ be a schema and let E be a c-join-expression over $X_1, \ldots, X_m$. If W is a consistency witness function on $\mathbb{K}$ and $R_1(X_1), \ldots, R_m(X_m)$ is a collection of $\mathbb{K}$-relations, we write $E(W, R_1, \ldots, R_m)$ to denote the $\mathbb{K}$-relation over $X_1 \cdots X_m$ obtained by evaluating E when $\bowtie_{\mathrm{c}}$ is interpreted by W and each X_i is interpreted by R_i for $i = 1, \ldots, m$.

The next notion yields a sufficient condition for a c-join expression to give rise to global consistency witnesses.

- Let E be a c-join expression over a schema $X_1, \ldots, X_m$, let W be a consistency witness function on $\mathbb{K}$, and let $R_1(X_1), \ldots, R_m(X_m)$ be a collection of $\mathbb{K}$-relations. We say that E is *monotone with respect to* W *and* $R_1(X_1), \ldots, R_m(X_m)$ if for every sub-expression $E_1 \bowtie_{\mathrm{c}} E_2$ of E, we have that the $\mathbb{K}$-relations $E_1(W, R_1, \ldots, R_m)$ and $E_2(W, R_1, \ldots, R_m)$ are consistent.

▶ **Example 3.** Let H^* be the schema $\{A, B\}, \{A, C\}, \{A, B, C\}$ and consider the sequential c-join-expressions E_1 and E_2 over H^*, where

$$E_1 = ((\{A, B\} \bowtie_{\mathrm{c}} \{A, B, C\}) \bowtie_{\mathrm{c}} \{A, C\}) \quad \text{and} \quad E_2 = ((\{A, B\} \bowtie_{\mathrm{c}} \{A, C\}) \bowtie_{\mathrm{c}} \{A, B, C\}).$$

Consider also the monoid $\mathbb{V} = ([0, 1], \max, 0)$ and the consistency witness function W on $\mathbb{V}$ such that $W(R, S)(t) = \min\{R(t[X]), S(t[Y])\}$, for every XY-tuple t (see Example 2). Let f, t be two values in some domain with $f \neq t$, let c, d be two real numbers such that $0 < c < d < 1$, and let $R_1(A, B), R_2(A, C), R_3(A, B, C)$ be the following $\mathbb{V}$-relations:

- $R_1(f,f) = c$, $R_1(f,t) = d$, and $R_1(x,y) = 0$, for all other pairs.
- $R_2(f,f) = c$, $R_2(f,t) = d$, and $R_2(x,y) = 0$, for all other pairs.
- $R_3(f,f,f) = c$, $R_3(f,t,t) = d$, and $R_3(x,y,z) = 0$, for all other triples.

Let us evaluate E_1 and E_2 on R_1, R_2, R_3 using the consistency witness function W. It is easy to check that $E_1(W, R_1, R_2, R_3) = W(W(R_1, R_3), R_2) = W(R_3, R_2) = R_3$ and that $E_2(W, R_1, R_2, R_3) = W(W(R_1, R_2), R_3) = W(R_4, R_3) = R_3$, where $R_4(A,B,C)$ is the $\mathbb{V}$-relation with $R_4(f,f,f) = c$, $R_4(f,f,t) = c$, $R_4(f,t,f) = c$, $R_4(f,t,t) = d$, and $R_4(x,y,z) = 0$, for all other triples.

While both E_1 and E_2 evaluate to the same $\mathbb{V}$-relation, there is an important difference in the intermediate steps of their evaluations. Specifically, in the case of E_1, the intermediate relation $W(R_1, R_2) = R_3$ is consistent with the last relation R_2; in contrast, in the case of E_2, the intermediate relation $W(R_1, R_2) = R_4$ is *not* consistent with the last relation R_3 (since R_3 and R_4 are different $\mathbb{V}$-relations on the same set of attributes). In particular, this example shows that the sequential c-join-expression E_2 is *not* monotone with respect to W.

The next proposition from [3] is proved in a straightforward way by induction on the construction of c-join expressions and by using Proposition 1.

▶ **Proposition 4.** *Let E be a* c*-join expression over $X_1, \dots, X_m$, let W be a consistency witness function on $\mathbb{K}$, and let $R_1(X_1), \dots, R_m(X_m)$ be $\mathbb{K}$-relations. If E is monotone with respect to W and $R_1, \dots, R_m$, and every X_i occurs in E, then $E(W, R_1, \dots, R_m)$ is a global consistency witness for the $\mathbb{K}$-relations $R_1, \dots, R_m$.*

4 Acyclic Hypergraphs

As mentioned in Section 2, if $R_1, \dots, R_m$ are standard relations that are globally consistent, then they are pairwise consistent, but the converse does not always hold. For example, consider the *triangle* schema $\{A,B\}, \{B,C\}, \{C,A\}$ and the standard relations $R_1(A,B) = \{(0,0),(1,1)\}$, $R_2(B,C) = \{(0,1),(1,0)\}$, $R_3(C,A) = \{(0,0),(1,1)\}$. It is easy to check that these standard relations are pairwise consistent; however, they are not globally consistent since $((R_1 \bowtie R_2) \bowtie R_3) = \emptyset$. Beeri et al. [4] characterized the schemas for which every collection of pairwise consistent standard relations is globally consistent by showing that these are precisely the *acyclic* (also known as α-*acyclic*) schemas. To give the precise definition of an acyclic schema, we need to first introduce some basic notions about hypergraphs.

Let $H = (V, F)$ be a hypergraph in which each vertex belong to at least one hyperedge.

- A *path in H* is a sequence $X_1, \dots, X_k$ of hyperedges of H such that $X_i \cap X_{i+1} \neq \emptyset$, for every i with $1 \leq i < k$. In this case, we say that there is a *path from X_1 to X_k*.
- We say that a set G of hyperedges of H is *connected* if for every two distinct hyperedges X and X' of G, there is a path from X to X'.
- A *connected component* of H is a maximal connected set of hyperedges of H
- We say that H is *connected* if the set F of the hyperedges of H is connected (in other words, H has a single connected component); otherwise, we say that H is *disconnected*.
- H is *reduced* if no hyperedge of H is properly contained in some other hyperedge of H.
- The *reduction* of H is the hypergraph (V, F'), where F' consists of the hyperedges in F that are not properly contained in some other hyperedge in F.
- If $U \subseteq V$, then the *induced hypergraph* $H[U]$ is the hypergraph $(U, \{X \cap U : X \in F\} \setminus \{\emptyset\})$. Also, the *restriction of H on U*, denoted $H \upharpoonright U$, is the reduction of the hypergraph $H[U]$.
- Let H be a reduced hypergraph and let X, X' be two distinct hyperedges. We say that $Y = X \cap X'$ is an *articulation set* of H if the number of connected components of $H \upharpoonright (V \setminus Y)$ is bigger than the number of connected components of H.

- Let H be a reduced hypergraph. We say that H is *acyclic* if the following condition holds: for every set $U \subseteq V$, if $H \restriction U$ is connected and has at least two hyperedges, then it has an articulation set; otherwise, we say that H is *cyclic*.
- We say that H is *acyclic* if its reduction is acyclic; otherwise, it is *cyclic*. Acyclic hypergraphs are also known as *α-acyclic* hypergraphs, while cyclic hypergraphs are also known as α-cyclic hypergraphs.

Admittedly, the notion of an acyclic hypergraph appears to be difficult to grasp when encountered for the first time. Intuitively, it generalizes to hypergraphs the property that a graph is acyclic if and only if every connected component of it with at least two edges has an articulation point. It is well known that there is a polynomial-time algorithm for testing if a hypergraph is acyclic; this algorithm is due to Graham [10] and, independently, to Yu and Ozsoyoglu [17], and it is known as the GYO algorithm (see also [8]).

There are several different structural conditions that are equivalent to acyclicity. We discuss two of these notions next.

Let $H = (V, F)$ be a hypergraph.

- The *Gaifman graph* $G(H)$ of H is the undirected graph with nodes the attributes of H and such that there is an edge between two attributes of H if and only if both these attributes belong to one of the hyperedges of H.
- H is *conformal* if every clique of $G(H)$ is contained in one of the hyperedges of H.
- H is *chordal* if every cycle of $G(H)$ of length at least 4 has a *chord*, i.e., there is an edge of $G(H)$ that is not an edge of the cycle.
- H has the *running intersection property* if there is an ordering $Y_1, \ldots, Y_m$ of the hyperedges of H such that for every $i \leq m$, there is a $j < i$ such that $(Y_1 \cup \cdots \cup Y_{i-1}) \cap Y_i \subseteq Y_j$.

The proof of the next result can be found in [4].

▶ **Proposition 5.** *For every hypergraph H, the following statements are equivalent:*
1. *H is acyclic.*
2. *H is conformal and chordal.*
3. *H has the running intersection property.*

The triangle schema $\{A, B\}, \{B, C\}, \{C, A\}$ is cyclic because it is not conformal (but it is chordal); the *4-cycle* schema $\{A, B\}, \{B, C\}, \{C, D\}, \{D, A\}$ is cyclic because it is not chordal (but it is conformal). For every $n \geq 2$, the *n-path* schema P_n with hyperedges $\{A_1, A_2\}, \{A_2, A_3\}, \ldots, \{A_n, A_{n+1}\}$ is acyclic because it is both conformal and chordal. One can also reason about these schemas using the running intersection property. Finally, consider the schema $\{A, B, C\}, \{C, D, E\}, \{E, F, A\}, \{A, C, E\}$. It has the running intersection property via the ordering $\{A, B, C\}, \{A, C, E\}, \{C, D, E\}, \{E, F, A\}$, hence it is acyclic.

Beeri et al. [4] showed that acyclicity can be characterized in terms of useful semantic properties, where by "semantic" we mean a property of the hypergraph whose definition involves also standard relations. Specifically, consider the following two properties.

Let H be a hypergraph with $X_1, \ldots, X_m$ as its hyperedges.

- H has the *local-to-global consistency property for standard relations* if for every collection $R_1(X_1), \ldots, R_m(X_m)$ of pairwise consistent standard relations over H, we have that this collection is globally consistent.
- H *admits a monotone join expression* if there is a c-join expression E such that
  - i E is monotone with respect to the standard join operation $\bowtie$ and every collection of pairwise consistent relations $R_1(X_1), \ldots, R_m(X_m)$;
  - ii Every hyperedge X_i of H occurs in E.

As mentioned in Section 2, if a collection of relations is globally consistent, then it is pairwise consistent. Thus, if a schema has the local-to-global consistency property for standard relations, then global consistency coincides with pairwise consistency for relations over that schema. Furthermore, if a schema admits a monotone join expression, then, in view of Proposition 4, this join expression can be used to construct witnesses to global consistency of collections of pairwise consistent relations.

With the notions of local-to-global consistency property and monotone join expression at hand, the main result in Beeri et al. [4] can be stated as follows.

▶ **Theorem 6** ([4]). *For every hypergraph H, the following statements are equivalent:*
1. *H is acyclic.*
2. *H has the local-to-global consistency property for standard relations.*
3. *H admits a monotone sequential join expression.*

Thus, the triangle schema and the 4-cycle schema do not have the local-to-global consistency property for standard relations and do not admit a monotone join expression, but each n-path schema does, and so does the schema $\{A, B, C\}, \{C, D, E\}, \{E, F, A\}, \{A, C, E\}$.

In [2], the following question was investigated: does Theorem 6 extend from standard relations to $\mathbb{K}$-relations, where $\mathbb{K}$ is an arbitrary positive commutative monoid?

The first realization in [2] was that acyclicity is a necessary, but not always sufficient, condition for the local-to-global consistency property to hold for $\mathbb{K}$-relations. More formally, we say that a hypergraph $H = \{X_1, \dots, X_m\}$ has the *local-to-global consistency property for* $\mathbb{K}$*-relations* if every collection $R_1(X_1), \dots, R_m(X_m)$ of pairwise consistent $\mathbb{K}$-relations is also globally consistent.

We can now state the precise result about the necessity of acyclicity.

▶ **Theorem 7** ([2]). *The following statements are true:*
1. *For all positive commutative monoids $\mathbb{K}$ and hypergraphs H, if H has the local-to-global consistency property for $\mathbb{K}$-relations, then H is acyclic.*
2. *There are positive commutative monoids $\mathbb{K}$ and hypergraphs H such that H is acyclic and does* not *have the local-to-global consistency property for $\mathbb{K}$-relations. In particular, this holds true for every numerical semigroup $\mathbb{K}$ other than the bag monoid $\mathbb{N} = (N, +, 0)$ and for the 3-path hypergraph P_3.*

Furthermore, in [2], a class of positive commutative monoids $\mathbb{K}$ was identified for which the acyclicity of a hypergraph H is a sufficient condition for H to have the local-to-global consistency property for $\mathbb{K}$-relations.

- Two $\mathbb{K}$-relations $R(X)$ and $S(Y)$ are *inner consistent* if $R[X \cap Y] = S[X \cap Y]$.
- We say that $\mathbb{K}$ has the *inner consistency property* if whenever two $\mathbb{K}$-relations are inner consistent, they are also consistent.

Note that, using Proposition 1, it is easy to verify that if R and S are consistent $\mathbb{K}$-relations, then they are also inner consistent. Consequently, for monoids with the inner consistency property, the notions of consistency and inner consistency coincide. In particular, this holds true for the the Boolean monoid $\mathbb{B} = (\{0, 1\}, \vee, 0)$.

In [1], it was shown that the bag monoid $\mathbb{N} = (N, +, 0)$ has the inner consistency property. In [2], several other monoids were shown to have the inner consistency property, including the monoids $\mathbb{V} = ([0, 1], \max, 0)$, $\mathbb{R}^{\geq 0} = ([0, \infty), +, 0)$, $\mathbb{T} = ((-\infty, \infty], \min, \infty)$, and the power set monoids $\mathbb{P}(A) = (\mathcal{P}(A), \cup, \emptyset)$, for every set A. In contrast, no numerical semigroup other than the bag monoid has the inner consistency property.

To state the extension of Theorem 6 to $\mathbb{K}$-relations, we need some additional notions, which were introduced in [3].

Let $\mathbb{K}$ be a positive commutative monoid, let $X_1, \ldots, X_m$ be a schema, and let E be a c-join expression over $X_1, \ldots, X_m$.

- First, recall from Section 3, that E is monotone with respect to a consistency witness function W on $\mathbb{K}$ and a collection $R_1(X_1), \ldots, R_m(X_m)$ of $\mathbb{K}$-relations if for every sub-expression $E_1 \bowtie_c E_2$ of E, we have that the $\mathbb{K}$-relations $E_1(W, R_1, \ldots, R_m)$ and $E_2(W, R_1, \ldots, R_m)$ are consistent.
- We say that E is *monotone on* $\mathbb{K}$ if there is a consistency witness function W on $\mathbb{K}$ such that E is monotone with respect to W and every collection $R_1(X_1), \ldots, R_m(X_m)$ of pairwise consistent $\mathbb{K}$-relations.
- We say that E is *strongly monotone on* $\mathbb{K}$ if for every consistency witness function W on $\mathbb{K}$, we have that E is monotone with respect to W and every collection $R_1(X_1), \ldots, R_m(X_m)$ of pairwise consistent $\mathbb{K}$-relations.

Finally, we define what it means for a schema to admit a monotone and a strongly monotone c-join expression.

- A schema $X_1, \ldots, X_m$ *admits a monotone* c*-join expression on* $\mathbb{K}$ if there is a c-join-expression E over $X_1, \ldots, X_m$ such that E is monotone on $\mathbb{K}$ and every hyperedge X_i occurs in E.
- A schema $X_1, \ldots, X_m$ *admits a strongly monotone* c*-join expression on* $\mathbb{K}$ if there is a c-join-expression E over $X_1, \ldots, X_m$ such that E is strongly monotone on $\mathbb{K}$ and every hyperedge X_i occurs in E.

We can now state one of the main results from [2, 3].

▶ **Theorem 8** ([2, 3]). *Let $\mathbb{K}$ be a positive commutative monoid that has the inner consistency property. For every hypergraph H, the following statements are equivalent:*

1. *H is acyclic.*
2. *H has the local-to-global consistency property for $\mathbb{K}$-relations.*
3. *H admits a monotone sequential* c*-join-expression on $\mathbb{K}$.*
4. *H admits a strongly monotone sequential* c*-join expression on $\mathbb{K}$.*

Theorem 8 yields Theorem 6 of Beeri et al. [4] by taking $\mathbb{K}$ to be the Boolean monoid $\mathbb{B}$. In fact, Theorem 8 yields something stronger: the standard join $\bowtie$ can be replaced in Theorem 6 by an arbitrary consistency witness function for standard relations. In effect, this means that the *only* property of the standard join needed in Theorem 6 is that the standard join is a consistency witness function for standard relations.

As shown in [2], the inner consistency property actually characterizes the positive commutative monoids $\mathbb{K}$ for which every acyclic hypergraph has the local-to-global consistency property for $\mathbb{K}$-relations.

▶ **Theorem 9** ([2, 3]). *Let $\mathbb{K}$ be a positive commutative monoid. Then the following statements are equivalent:*

1. *$\mathbb{K}$ has the inner consistency property.*
2. *Every acyclic hypergraph has the local-to-global consistency property for $\mathbb{K}$-relations.*
3. *The 3-path hypergraph P_3 has the local-to-global consistency property for $\mathbb{K}$-relations.*
4. *Every acyclic hypergraph admits a monotone sequential* c*-join-expression on $\mathbb{K}$*
5. *Every acyclic hypergraph admits a strongly monotone sequential* c*-join-expression on $\mathbb{K}$.*

5 Beta Acyclic Hypergraphs

Their good structural and semantic properties notwithstanding, acyclic hypergraphs suffer from the drawback that acyclicity is not a hereditary property, that is to say, a sub-hypergraph of an acyclic hypergraph need not be acyclic. For example, the hypergraph $\{A, B, C\}, \{A, B\}, \{B, C\}, \{C, A\}$ is acyclic but it contains as a sub-hypergraph the triangle hypergraph $\{A, B\}$, $\{B, C\}$, $\{C, A\}$, which is cyclic.

Motivated by the preceding considerations, Fagin [8] introduced the following notion.

- A hypergraph H is *β-acyclic* if every sub-hypergraph of H is acyclic; otherwise, H is *β-cyclic*.

For example, the hypergraph H^* with hyperedges $\{A, B, C\}, \{A, B\}, \{A, C\}$ is β-acyclic. There is a polynomial-time algorithm for testing if a hypergraph is β-acyclic (see [8]).

Fagin [8] found several different characterizations of β-acyclicity, including some involving the absence of cycles of certain types. In the sequel, we will use the following one.

- A *weak β-cycle* in a hypergraph H is a sequence $Y_1, A_1, Y_2, A_2, \ldots, Y_k, A_k, Y_{k+1}$ such that $k \geq 3$ and the following properties hold:
  1. $Y_1, \ldots, Y_k$ are distinct hyperedges of H and $Y_{k+1} = Y_1$;
  2. $A_1, \ldots, A_k$ are distinct nodes of H;
  3. For $1 \leq i \leq k$, the node A_i is in $Y_i \cap Y_{i+1}$.
  4. For $1 \leq i \leq k$, the node A_i is not in any hyperedge Y_j other than Y_i and Y_{i+1}.

▶ **Theorem 10** ([8]). *For every hypergraph H, the following statements are equivalent:*
1. *H is β-acyclic.*
2. *H has no weak β-cycles.*

For example, consider the schema $\{A, B, C\}, \{C, D, E\}, \{E, F, A\}, \{A, C, E\}$, which was shown earlier to be acyclic. The sequence $\{A, B, C\}, C, \{C, D, E\}, E, \{E, F, A\}, A, \{A, B, C\}$ is a weak β-cycle, hence this schema is β-cyclic.

Because of the hereditary nature of β-acyclicity, the characterizations of acyclicity in Section 4 give rise to characterizations of β-acyclicity. Thus, the following result holds.

▶ **Corollary 11.** *Let $\mathbb{K}$ be a positive commutative monoid that has the inner consistency property. For every hypergraph H, the following statements are equivalent:*
1. *H is β-acyclic.*
2. *Every sub-hypergraph of H has the local-to-global consistency property for $\mathbb{K}$-relations.*
3. *Every sub-hypergraph of H admits a monotone sequential* c*-join-expression on $\mathbb{K}$.*
4. *Every sub-hypergraph of H admits a strongly monotone sequential* c*-join expression on $\mathbb{K}$.*

6 Gamma-Acyclic Hypergraphs

Fagin [8] introduced and studied γ-acyclic hypergraphs, which form a proper subclass of the class of β-acyclic hypergraphs. As with β-acyclic hypergraphs, there are several equivalent formulations of the notion of a γ-acyclic hypergraph in terms of absence of cycles of certain types, including the following one.

- A *weak γ-cycle* in a hypergraph H is a sequence $Y_1, A_1, Y_2, A_2, \ldots, Y_k, A_k, Y_{k+1}$ such that $k \geq 3$ and the following properties hold:
  1. $Y_1, \ldots, Y_k$ are distinct hyperedges of H and $Y_{k+1} = Y_1$;
  2. $A_1, \ldots, A_k$ are distinct nodes of H;
  3. For $1 \leq i \leq k$, the node A_i is in $Y_i \cap Y_{i+1}$;
  4. For $i = 1, 2$, the node A_i is not in any hyperedge Y_j other than Y_i and Y_{i+1}.
- A hypergraph H is *γ-acyclic* if H has no weak γ-cycle; otherwise, H is *γ-cyclic*.

Clearly, every sub-hypergraph of a γ-acyclic hypergraph is γ-acyclic as well. Observe that the only difference between a weak β-cycle and a weak γ-cycle is in the fourth condition of the definitions of these notions: the requirement that the node A_i belongs only to the hyperedges Y_i and Y_{i+1} holds for every $i \leq k$ in the case of a weak β-cycle, while it holds for $i = 1, 2$ in the case of a weak γ-cycle. In particular, every weak β-cycle is also a weak γ-cycle; consequently, every γ-acyclic hypergraph is also β-acyclic. The converse, however, is not true. To see this, consider the hypergraph H^* with hyperedges $\{A, B, C\}, \{A, B\}, \{A, C\}$, which, as pointed out in Section 5, is β-acyclic. Clearly, the sequence $\{A, B\}, B, \{A, B, C\}, C, \{A, C\}, A, \{A, B\}$ is a weak γ-cycle, hence this hypergraph is γ-cyclic. It is also easy to see that for every $n \geq 2$, the n-path hypergraph P_n with hyperedges $\{A_1, A_2\}, \ldots, \{A_n, A_{n+1}\}$ is γ-acyclic. Note that there is a polynomial-time algorithm, due to D' Atri and Moscarini [7], for testing whether or not a hypergraph is γ-acyclic (see also [8, Section 9.4]).

There are several different structural characterizations of γ-acyclic hypergraphs, including one which is due to Brault-Baron [5] and which involves the hypergraph H^* above. To describe this characterization, let us recall the following basic notion from Section 4.

- Let $H = (V, F)$ be a hypergraph and let U be a subset of the set V of the nodes of H. The *induced hypergraph* $H[U]$ is the hypergraph with hyperedges $\{X \cap U : X \in E\} \setminus \{\emptyset\}$.

▶ **Proposition 12** ([5]). *For every hypergraph H, the following statements are equivalent:*

1. *H is γ-acyclic.*
2. *H is β-acyclic and there do not exist three nodes A, B, C of H such that the hypergraph H^* with hyperedges $\{A, B, C\}, \{A, B\}, \{A, C\}$ is a sub-hypergraph of $H[\{A, B, C\}]$.*

In fact, Brault-Baron [5] defines a γ-acyclic hypergraph to be a hypergraph that satisfies the second condition in Proposition 12; he then shows that this condition is equivalent to the D' Atri and Moscarini algorithm [7] producing the empty hypergraph, hence this condition is equivalent to Fagin's [8] definition of γ-acyclicity.

Fagin [8] established that γ-acyclic hypergraphs have certain desirable semantic properties; the main such property involves the notion of a *connected* join expression.

- A join expression E is *connected* if for each of its sub-expressions $(E_1 \bowtie E_2)$, there is an attribute that appears in both E_1 and E_2.

In particular, if E is a sequential join expression $(((\cdots(Y_1 \bowtie Y_2) \bowtie \cdots) \bowtie Y_{k-1}) \bowtie Y_k)$, then E is connected if and only if for every with $1 \leq i \leq k$, we have that $(Y_1 \cup \cdots \cup Y_{i-1}) \cap Y_i \neq \emptyset$. For example, the sequential join expression $((\{A_1, A_2\} \bowtie \{A_2, A_3\}) \bowtie \{A_3, A_4\})$ is connected, while the sequential join expression $(\{A_1, A_2\} \bowtie \{A_3, A_4\})$ is not connected.

We can now state the main semantic characterization of γ-acyclicity, obtained in [8].

▶ **Theorem 13** ([8]). *For every hypergraph H, the following statements are equivalent:*

1. *H is γ-acyclic.*
2. *Every connected sequential join expression over H is monotone.*

The notion of a connected join expression extends to the notion of a connected c-join expression in a straightforward way.

- A c-join expression E is *connected* if for each of its sub-expressions $(E_1 \bowtie_c E_2)$, there is an attribute that appears in both E_1 and E_2.

In particular, if E is a sequential c-join expression $(((\cdots(Y_1 \bowtie_c Y_2) \bowtie_c \cdots) \bowtie_c Y_{k-1}) \bowtie_c Y_k)$, then E is connected if for every with $1 \leq i \leq k$, we have that $(Y_1 \cup \cdots \cup Y_{i-1}) \cap Y_i \neq \emptyset$.

It is now natural to ask: does Theorem 13 extend and how does it extend to γ-acyclic hypergraphs, c-join expressions, and $\mathbb{K}$-relations, where $\mathbb{K}$ is a positive commutative monoid?

Let $\mathbb{K}$ be a positive commutative monoid and let E be a c-join expression. Recall that E is monotone on $\mathbb{K}$ if there is a consistency witness function W on $\mathbb{K}$ such that E is monotone with respect to W and every collection $R_1(X_1), \ldots, R_m(X_m)$ of pairwise consistent $\mathbb{K}$-relations. Recall further that E is strongly monotone on $\mathbb{K}$ if E is monotone with respect to every consistency witness function W on $\mathbb{K}$ and every collection $R_1(X_1), \ldots, R_m(X_m)$ of pairwise consistent $\mathbb{K}$-relations.

In what follows in this section, we will establish the following results:

1. If H is a hypergraph such that every connected sequential c-join-expression over H is monotone on $\mathbb{K}$, then H is γ-acyclic.
2. If $\mathbb{K}$ has the inner consistency property and H is γ-acyclic, then every connected sequential c-join-expression over H is strongly monotone on $\mathbb{K}$,

▶ **Theorem 14.** *For all positive commutative monoids $\mathbb{K}$ and hypergraphs H, if H is such that every connected sequential* c*-join-expression over H is monotone on $\mathbb{K}$, then H is γ-acyclic.*

Proof. We will prove the contrapositive, that is, if the hypergraph H is not γ-acyclic, then there is a connected sequential c-join expression E over H such that E is not monotone on $\mathbb{K}$. So, assume that H is not γ-acyclic. We distinguish two cases, namely, the case in which H is not β-acyclic and the case in which H is β-acyclic.

Case 1. Assume that H is not β-acyclic. By the definition of β-acyclicity, there is a sub-hypergraph H' of H that is cyclic. Moreover, we may assume that H' is connected, since if every connected component of H' were acyclic, then it is easy to see that H' would be acyclic as well. Let $X_1, \ldots, X_k$ be a list of the hyperedges of H'. Since H' is cyclic, the first part of Theorem 7 implies that there are $\mathbb{K}$-relations $R_1(X_1), \ldots, R_k(X_k)$ that are pairwise consistent, but not globally consistent. Since H' is connected, there is a sequence $Y_1, \ldots, Y_t$ of not necessarily distinct sets of attributes such that the following hold: (a) each Y_j is one of the hyperedges of H', i.e., $Y_j = X_{i_j}$, where $i_j \in \{1, \ldots, k\}$; (b) each hyperedge X_i of H' appears in the sequence $Y_1, \ldots, Y_t$; and (c) for every j with $1 \leq j < t$, we have that $Y_j \cap Y_{j+1} \neq \emptyset$. Let E be the sequential c-join expression $(((\cdots(Y_1 \bowtie_{\mathrm{c}} Y_2) \bowtie_{\mathrm{c}} \cdots) \bowtie_{\mathrm{c}} Y_{t-1}) \bowtie_{\mathrm{c}} Y_t)$. Then E is a connected c-join expression because $Y_j \cap Y_{j+1} \neq \emptyset$ holds, for every j with $1 \leq j < t-1$. We now claim that there is no consistency witness function W on $\mathbb{K}$ such that E is monotone with respect to W. To see this, let W be a consistency witness function on $\mathbb{K}$ and consider the pairwise consistent $\mathbb{K}$-relations $R_1(X_1), \ldots, R_k(X_k)$. If E were monotone with respect to W, then, by Proposition 4, the $\mathbb{K}$-relation $E(W, R_{i_1}, \ldots, R_{i_t})$ is a global consistency witness for the relations $R_{i_1}, \ldots, R_{i_t}$, hence it is a global consistency witness for the relations $R_1, \ldots, R_k$ since every hyperedge X_i appears in the sequence $Y_1, \ldots, Y_t$. This, however, is a contradiction since the relations $R_1, \ldots, R_k$ are not globally consistent.

Case 2. Assume that H is β-acyclic. Since H is not γ-acyclic, Proposition 12 implies that there are three attributes A, B, C and three hyperedges Y_1, Y_2, Y_3 of H such that $Y_1 \cap \{A, B, C\} = \{A, B\}$, $Y_2 \cap \{A, B, C\} = \{A, C\}$, and $Y_3 \cap \{A, B, C\} = \{A, B, C\}$. Let D_1, D_2, D_3 be the remaining sets of attributes in Y_1, Y_2, Y_3, respectively; thus, $Y_1 = \{A, B\} \cup D_1$, $Y_2 = \{A, C\} \cup D_2$, $Y_3 = \{A, B, C\} \cup D_3$. Let E be the sequential c-join expression $((Y_1 \bowtie_{\mathrm{c}} Y_2) \bowtie_{\mathrm{c}} Y_3))$, which is connected since $\{A, B, C\} \subseteq (Y_1 \cup Y_2) \cap Y_3$.

We now claim that there is no consistency witness function W on $\mathbb{K}$ such that E is monotone with respect to W. We define four $\mathbb{K}$-relations in pairs $R_1(A, B, D_1), R_2(A, C, D_2)$, and $S_1(A, B, C, D_1, D_2), S_2(A, B, C, D_1, D_2)$. To define them, let $\mathbf{f}_1$ and $\mathbf{f}_2$ be the tuples $(f, \ldots, f)$ of lengths equal to the cardinalities of D_1 and D_2, and let $\mathbf{f}$ be the tuple $(f, \ldots, f)$ of length equal to the cardinality of $D_1 \cup D_2$. Let a be a non-zero element $\mathbb{K}$ and let

- $R_1(f, f, \mathbf{f}_1) = a$, $R_1(f, t, \mathbf{f}_1) = a$, and $R_1(x, y, \mathbf{z}_1) = 0$, for all other values,
- $R_2(f, f, \mathbf{f}_2) = a$, $R_2(f, t, \mathbf{f}_2) = a$, and $R_2(x, y, \mathbf{z}_2) = 0$, for all other values,

- $S_1(f,f,f,\mathbf{f}) = a$, $S_1(f,t,t,\mathbf{f}) = a$, and $S_1(x,y,z,\mathbf{w}) = 0$, for all other values,
- $S_2(f,f,t,\mathbf{f}) = a$, $S_2(f,t,f,\mathbf{f}) = a$, and $S_2(x,y,z,\mathbf{w}) = 0$, for all other values.

It is easy to verify that $S_1(A,B,C,D_1,D_2)$ and $S_2(A,B,C,D_1,D_2)$ are two different consistency witnesses for the relations $R_1(A,B,D_1)$ and $R_2(A,C,D_2)$. It is easy to verify that $S_1(A,B,C,D_1,D_2)$ and $S_2(A,B,C,D_1,D_2)$ are two different consistency witnesses for the relations $R_1(A,B,D_1)$ and $R_2(A,C,D_2)$.

Now, let W be an arbitrary witness function on $\mathbb{K}$. Since the relations R_1 and R_2 are consistent, we have that the $\mathbb{K}$-relation $W(R_1,R_2)$ is a consistency witness function for R_1 and R_2. We distinguish the following two sub-cases.

Sub-case 1. $W(R_1,R_2)[ABC] = S_1[ABC]$. In this case, let $R_3(A,B,C,D_3)$ be the $\mathbb{K}$-relation such that $R_3(f,f,t,\mathbf{f}_3) = a$, $R_3(f,t,f,\mathbf{f}_3) = a$, and $R_3(x,y,z,\mathbf{w}) = 0$, for all other values, where $\mathbf{f}_3$ is a tuple of f's of length equal to the cardinality of the set D_3. Observe that $R_3[ABC] = S_2[ABC]$. Since S_2 is a consistency witness for R_1 and R_2, we have that the relations R_1, R_2, R_3 are pairwise consistent. However, the $\mathbb{K}$-relations $W(R_1,R_2) = S_1$ and R_3 are not consistent, since $S_1[ABC] \neq S_2[ABC]$ and $S_2[ABC] = R_3[ABC]$.

Sub-case 2. $W(R_1,R_2)[ABC] \neq S_1[ABC]$. In this case, let $R_3(A,B,C,D_3)$ be the $\mathbb{K}$-relation such that $R_3(f,f,f,\mathbf{f}_3) = a$, $R_3(f,t,t,\mathbf{f}_3) = a$, and $R_3(x,y,z,\mathbf{w}) = 0$, for all other values, where $\mathbf{f}_3$ is a tuple of f's of length equal to the cardinality of the set D_3. Observe that $R_3[ABC] = S_1[ABC]$. Since S_1 is a consistency witness for R_1 and R_2, we have that the relations R_1, R_2, R_3 are pairwise consistent. However, the $\mathbb{K}$-relations $W(R_1,R_2)$ and R_3 are not consistent, since $W(R_1,R_2)[ABC] \neq S_1[ABC] = R_3[ABC]$. ◀

Since the hypergraph H^* with hyperedges $\{A,B,C\}, \{A,B\}, \{A,C\}$ is γ-cyclic, Theorem 14 implies that there is a connected sequential c-join expression over H^* that is not monotone on $\mathbb{K}$. Actually, the proof of Theorem 14 implies that $((\{A,B\} \bowtie_c \{A,C\}) \bowtie_c \{A,B,C\})$ is such an expression, and Example 3 is an illustration of this fact.

The preceding Theorem 14 asserts that γ-acyclicity is a necessary condition for a hypergraph H to have the property that every connected sequential c-join expression over H is monotone on $\mathbb{K}$, where $\mathbb{K}$ is an arbitrary positive commutative monoid. The second main result in this section asserts that if $\mathbb{K}$ has the inner consistency property, then γ-acyclicity is a sufficient condition for a hypergraph to have the property that every connected sequential c-join expression over $\mathbb{K}$ is strongly monotone on $\mathbb{K}$.

▶ **Theorem 15.** *Let $\mathbb{K}$ be a positive commutative monoid that has the inner consistency property. If H is a γ-acyclic hypergraph, then every connected sequential* c*-join expression over H is strongly monotone on $\mathbb{K}$.*

Proof. We will establish the contrapositive, i.e., if H does not have the property that every connected sequential c-join expression over H is strongly monotone on $\mathbb{K}$, then H is γ-cyclic.

Assume that H lacks the above property. Then there exist a connected sequential c-join expression $E = (((\cdots(X_1 \bowtie_c X_2) \bowtie_c \cdots) \bowtie_c X_{m-1}) \bowtie_c X_m)$ over H, a consistency witness function W on $\mathbb{K}$, and a collection $R_1(X_1),\ldots,R_m(X_m)$ of pairwise consistent $\mathbb{K}$-relations such that E is not monotone w.r.t. to W and $R_1(X_1),\ldots,R_m(X_m)$. In turn, this means that there is some index $j < m$ such that the $\mathbb{K}$-relation $E_j(W,R_1,\ldots,R_j)$ is not consistent with the $\mathbb{K}$-relation R_{j+1}, where E_j is the sequential c-join expression $(((\cdots(X_1 \bowtie_c X_2) \bowtie_c \cdots) \bowtie_c X_{j-1}) \bowtie_c X_j)$. Let j be the smallest index with this property; thus, if $i < j$, then the $\mathbb{K}$-relation $E_i(W,R_1,\ldots,R_i)$ is consistent with the $\mathbb{K}$-relation R_{i+1}.

Let $Y = (X_1 \cup \cdots \cup X_j) \cap X_{j+1}$. Note that $Y \neq \emptyset$, since E is a connected sequential c-join expression. Note also that $X_1 \cup \cdots \cup X_j$ is the set of attributes of $E_j(W, R_1, \ldots, R_j)$, while X_{j+1} is the set of attributes of R_{j+1}. By hypothesis, $\mathbb{K}$ has the inner consistency property, which means that if two $\mathbb{K}$-relations $Q_1(X_1)$ and $Q_2(X_2)$ are inner consistent (i.e., $Q_1[X_1 \cap X_2] = Q_2[X_1 \cap X_2]$), then they are consistent. Thus, since the $\mathbb{K}$-relations $E_j(W, R_1, \ldots, R_j)$ and R_{j+1} are not consistent, it follows that $E_j(W, R_1, \ldots, R_j)$ and R_{j+1} are not inner consistent, hence $E_j(W, R_1, \ldots, R_j)[Y] \neq R_{j+1}[Y]$.

▷ Claim 1. For every $k \leq j$, we have that $Y \not\subseteq X_k$.

To establish the claim, assume that $Y \subseteq X_k$, for some $k \leq j$. By the pairwise consistency of $R_1, \ldots, R_m$, we have that R_k and R_{j+1} are consistent $\mathbb{K}$-relations, hence

$$R_k[X_k \cap X_{j+1}] = R_{j+1}[X_k \cap X_{j+1}].$$

Since $Y \subseteq X_k \cap X_{j+1}$, we have that $R_k[Y] = R_{j+1}[Y]$. Proposition 4 and the minimality assumption about j imply that the $\mathbb{K}$-relation $E_j(W, R_1, \ldots, R_j)$ is a global consistency witness for $R_1, \ldots, R_j$. Since $k \leq j$, we have that $E_j(W, R_1, \ldots, R_j)[X_k] = R_k$. Furthermore, since $Y \subseteq X_k$, we have that $E_j(W, R_1, \ldots, R_j)[Y] = R_k[Y]$, hence $E_j(W, R_1, \ldots, R_j)[Y] = R_{j+1}[Y]$; this contradicts the earlier finding that $E_j(W, R_1, \ldots, R_j)[Y] \neq R_{j+1}[Y]$, hence Claim 1 has been established.

So, we now know that for every $k \leq j$, we have that $Y \not\subseteq X_k$. Choose an index $k \leq j$ so that the cardinality $|X_k \cap Y|$ of the set $X_k \cap Y$ is the largest of the cardinalities $|X_i \cap Y|$ of the sets $X_i \cap Y$, $1 \leq i \leq j$. Pick a node A_1 such that $A_1 \in Y \setminus X_k$. Since $Y \subseteq X_1 \cup \cdots \cup X_j$, there is some $i \leq j$ such that $A_1 \in X_i$. Furthermore, since $E = (((\cdots(X_1 \bowtie_c X_2) \bowtie_c \cdots) \bowtie_c X_{m-1}) \bowtie_c X_m)$ is a connected sequential c-join expression over H, it is easy to see that $\{X_1, \ldots, X_j\}$ is a connected set of hyperedges of H. Let p be the length of the shortest path within $\{X_1, \ldots, X_j\}$ from X_k to a hyperedge containing A_1. This means that there is a sequence $S_1, \ldots, S_p$ of sets with the following properties:

1. Each set S_i is one of the hyperedges $X_1, \ldots, X_j$.
2. $S_1 = X_k$; $A_1 \in S_p$; $S_i \cap S_{i+1} \neq \emptyset$, $1 \leq i \leq p$.
3. p is the smallest number for which a sequence $S_1, \ldots, S_p$ with properties (1), (2) exists.

The minimality of p implies that $S_1, \ldots, S_p$ are distinct hyperedges of H. By the maximality of k, we have that $|S_p \cap Y| \leq |X_k \cap Y|$. Since $A_1 \in S_p \cap Y$ and $A_1 \notin X_k \cap Y$, there must exist a node $A_2 \in (X_k \cap Y) \setminus (S_p \cap Y)$. Let n be the largest number such that $A_2 \in S_n$ and $1 \leq n < p$. This number n exists because $S_1 = X_k$. By property (4) of the sequence $S_1 \ldots, S_p$, there are nodes B_i, $n \leq i < j$ such that $B_i \in S_i \cap S_{i+1}$. Consider now the sequence

$$(S_p, A_1, X_{j+1}, A_2, S_n, B_n, S_{n+1}, B_{n+1}, \ldots, B_{p-1}, S_p).$$

▷ Claim 2. $(S_p, A_1, X_{j+1}, A_2, S_n, B_n, S_{n+1}, B_{n+1}, \ldots, B_{p-1}, S_p)$ is a weak γ-cycle in H.

We have to verify that this sequence satisfies the conditions defining a weak γ-cycle, which were spelled out in the beginning of this section.

1. The hyperedges $S_p, X_{j+1}, S_n, S_{n+1}, \ldots, S_{p-1}$ are distinct, since the hyperedges $S_1, \ldots, S_p$ are distinct and also are among the hyperedges $X_1, \ldots, X_j$.
2. The nodes $A_1, A_2, B_n, B_{n+1}, \ldots, B_{p-1}$ are distinct for the following reasons: first, $A_1 \neq A_2$ because $A_1 \notin X_k \cap Y$ and $A_2 \in X_k \cap Y$; second, A_1 is different from $B_n, B_{n+1}, \ldots, B_{p-1}$, since $A_1 \notin S_i$ for $n \leq i < p$; A_2 is different from $B_n, B_{n+1}, \ldots, B_{p-1}$ by the choice of n as the largest number such that $A_2 \in S_n$ and the fact that $B_i \in S_i \cap S_{i+1}$; third, the B_i's are distinct, else the path would have been shorter.

3. First, $A_1 \in S_p \cap X_{j+1}$ by the choice of S_p and the fact that $A_1 \in Y \subseteq X_{j+1}$; second, $A_2 \in X_{j+1} \cap S_n$ by the choice of A_2; third, $B_i \in S_i \cap S_{i+1}$ by the choice of the B_i's.
4. A_1 is not in any of the hyperedges $S_n, S_{n+1}, \ldots, S_{p-1}$ by the properties of the path $S_1, \ldots, S_p$; finally, A_2 is not in any of the hyperedges $S_{n+1}, \ldots, S_p$ by the choice of A_2.

This completes the proof of Claim 2, hence the hypergraph H is γ-cyclic. ◀

The reader familiar with Fagin's paper [8] will undoubtedly notice that the proof of Theorem 15 has a very similar structure to the proof of the result that if a hypergraph H is γ-acyclic, then every connected sequential join expression is monotone (see [8, pages 539-540]). The main difference is that here, instead of the standard join, we use an arbitrary consistency witness function and, thus, obtain a stronger result about annotated relations and arbitrary consistency witness functions. Furthermore, it is of the essence that the annotations come from a monoid $\mathbb{K}$ that has the inner consistency property because, in the first part of the proof before Claim 1, we used in a crucial way that $\mathbb{K}$ has the inner consistency property.

▶ **Corollary 16.** *Let $\mathbb{K}$ be a positive commutative monoid that has the inner consistency property. For every hypergraph H, the following statements are equivalent:*
1. *H is γ-acyclic.*
2. *Every connected sequential* c*-join expression over H is strongly monotone on $\mathbb{K}$.*
3. *Every connected sequential* c*-join-expression over H is monotone on $\mathbb{K}$.*

Proof. The implication (1) $\Longrightarrow$ (2) follows from Theorem 15; the implication (2) $\Longrightarrow$ (3) follows from the definitions; the implication (3) $\Longrightarrow$ (1) follows from Theorem 14. ◀

Finally, we characterize the inner consistency property in terms of γ-acyclicity.

▶ **Theorem 17.** *Let $\mathbb{K}$ be a positive commutative monoid. Then the following statements are equivalent:*
1. *$\mathbb{K}$ has the inner consistency property.*
2. *Every γ-acyclic hypergraph H has the property that every connected sequential* c*-join expression over H is strongly monotone on $\mathbb{K}$.*
3. *Every γ-acyclic hypergraph H has the property that every connected sequential* c*-join-expression over H is monotone on $\mathbb{K}$.*
4. *The 3-path hypergraph P_3 has the property that every connected sequential* c*-join-expression over P_3 is monotone on $\mathbb{K}$.*

Proof. The implication (1) $\Longrightarrow$ (2) follows from Theorem 15. The implication (2) $\Longrightarrow$ (3) follows from the definitions. The implication (3) $\Longrightarrow$ (4) holds because the 3-path hypergraph P_3 is γ-acyclic. Towards (4) $\Longrightarrow$ (1), let P_3 be the 3-path hypergraph with hyperedges $\{A_1, A_2\}, \{A_2, A_3\}, \{A_3, A_4\}$. Assume that P_3 has the property that every connected sequential c-join expression over P_3 is monotone on $\mathbb{K}$. Let E be the connected sequential c-join expression $((\{A_1, A_2\} \bowtie_c \{A_2, A_3\}) \bowtie_c \{A_3, A_4\})$ over H and let W be a consistency witness function on $\mathbb{K}$ such that P_3 is monotone with respect to W and every three pairwise consistent $\mathbb{K}$-relations $R_1(A_1, A_2), R_2(A_2, A_3), R_2(A_3, A_4)$. Then P_3 has the local-to-global consistency property property for $\mathbb{K}$-relations, because if $R_1(A_1, A_2), R_2(A_2, A_3), R_2(A_3, A_4)$ are three pairwise consistent $\mathbb{K}$-relations, then, by Proposition 4, the $\mathbb{K}$-relation $E(W, R_1, R_2, R_3)$ is a global consistency witness for the $\mathbb{K}$-relations $R_1(A_1, A_2), R_2(A_2, A_3), R_2(A_3, A_4)$. Thus, by Theorem 9, the monoid $\mathbb{K}$ has the inner consistency property. ◀

7 Concluding Remarks

In this paper, we showed that the main desirable semantic properties of γ-acyclic hypergraphs and standard relations extend to desirable semantic properties of γ-acyclic hypergraphs and annotated relations, as long as the annotations come from a positive commutative monoid with the inner consistency property. It can also be shown that other desirable semantic properties of γ-acyclicity, such as join dependencies and lossless joins (see Fagin [8]), have suitable extensions to γ-acyclic hypergraphs and annotated relations; the notions of a consistency witness function and a c-join expression are used in these results.

Going back to acyclic hypergraphs, Beeri et al. [4] characterized acyclic hypergraphs in terms of semijoin programs and full reducers, which make it possible to evaluate queries efficiently in a distributed setting. It remains an open problem to determine whether or not there is an abstract notion of a semijoin program that can be used to extend these characterizations to acyclic schemas and annotated relations.

References

1 Albert Atserias and Phokion G. Kolaitis. Structure and complexity of bag consistency. In Leonid Libkin, Reinhard Pichler, and Paolo Guagliardo, editors, *PODS'21: Proceedings of the 40th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems, Virtual Event, China, June 20-25, 2021*, pages 247–259. ACM, 2021. doi:10.1145/3452021.3458329.

2 Albert Atserias and Phokion G. Kolaitis. Consistency of relations over monoids. *Journal of the ACM*, 72(3, Article 18):47 pages, 2025. Earlier version in Proc. ACM Manag. Data, 2(2):107, 2024. doi:10.1145/3721855.

3 Albert Atserias and Phokion G. Kolaitis. Consistency witnesses for annotated relations. *SIGMOD Rec.*, 54(2):7–17, 2025. doi:10.1145/3749116.3749118.

4 Catriel Beeri, Ronald Fagin, David Maier, and Mihalis Yannakakis. On the desirability of acyclic database schemes. *J. ACM*, 30(3):479–513, July 1983. doi:10.1145/2402.322389.

5 Johann Brault-Baron. Hypergraph acyclicity revisited. *ACM Comput. Surv.*, 49(3):54:1–54:26, 2016. doi:10.1145/2983573.

6 Katrin M. Dannert, Erich Grädel, Matthias Naaf, and Val Tannen. Semiring provenance for fixed-point logic. In Christel Baier and Jean Goubault-Larrecq, editors, *29th EACSL Annual Conference on Computer Science Logic, CSL 2021, January 25-28, 2021, Ljubljana, Slovenia (Virtual Conference)*, volume 183 of *LIPIcs*, pages 17:1–17:22. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPIcs.CSL.2021.17.

7 Alessandro D'Atri and Marina Moscarini. On the recognition and design of acyclic databases. In Daniel J. Rosenkrantz and Ronald Fagin, editors, *Proceedings of the Third ACM SIGACT-SIGMOD Symposium on Principles of Database Systems, April 2-4, 1984, Waterloo, Ontario, Canada*, pages 1–8. ACM, 1984. doi:10.1145/588011.588013.

8 Ronald Fagin. Degrees of acyclicity for hypergraphs and relational database schemes. *J. ACM*, 30(3):514–550, 1983. doi:10.1145/2402.322390.

9 Erich Grädel and Val Tannen. Semiring provenance for first-order model checking. *CoRR*, abs/1712.01980, 2017. arXiv:1712.01980.

10 Marc H. Graham. On the universal relation. Technical Report, Computer Systems Research Group, University of Toronto, 1980.

11 Todd J. Green. Containment of conjunctive queries on annotated relations. *Theory Comput. Syst.*, 49(2):429–459, 2011. doi:10.1007/s00224-011-9327-6.

12 Todd J. Green, Gregory Karvounarakis, and Val Tannen. Provenance semirings. In Leonid Libkin, editor, *Proceedings of the Twenty-Sixth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, June 11-13, 2007, Beijing, China*, pages 31–40. ACM, 2007. doi:10.1145/1265530.1265535.

13 Grigoris Karvounarakis and Todd J. Green. Semiring-annotated data: queries and provenance? *SIGMOD Rec.*, 41(3):5–14, 2012. doi:10.1145/2380776.2380778.

14 Mahmoud Abo Khamis, Hung Q. Ngo, Reinhard Pichler, Dan Suciu, and Yisu Remy Wang. Convergence of datalog over (pre-) semirings. *J. ACM*, 71(2):8:1–8:55, 2024. doi:10.1145/3643027.

15 Egor V. Kostylev, Juan L. Reutter, and András Z. Salamon. Classification of annotation semirings over containment of conjunctive queries. *ACM Trans. Database Syst.*, 39(1):1:1–1:39, 2014. doi:10.1145/2556524.

16 Mihalis Yannakakis. Algorithms for acyclic database schemes. In *Very Large Data Bases, 7th International Conference, September 9-11, 1981, Cannes, France, Proceedings*, pages 82–94. IEEE Computer Society, 1981.

17 Clement Tak Yu and Meral Z Ozsoyoglu. An algorithm for tree-query membership of a distributed query. In *COMPSAC 79. Proceedings. Computer Software and The IEEE Computer Society's Third International Applications Conference, 1979.*, pages 306–312. IEEE, 1979.