



PB96-150669

NTIS
Information is our business.


CONNECTIONS IN ACYCLIC HYPERGRAPHS

19970415 047

STANFORD UNIV., CA

DTIC QUALITY INSPECTED 2

MAY 81



U.S. DEPARTMENT OF COMMERCE
National Technical Information Service

DISTRIBUTION STATEMENT A

Approved for public release;
Distribution Unlimited

BIBLIOGRAPHIC INFORMATION

PB96-150669

Report Nos: STAN-CS-81-853

Title: Connections in Acyclic Hypergraphs.

Date: May 81

Authors: D. Maier and J. D. Ullman.

Performing Organization: Stanford Univ., CA. Dept. of Computer Science.

Sponsoring Organization: *National Science Foundation, Washington, DC.*Air Force Office of Scientific Research, Bolling AFB, DC.

Contract Nos: AFOSR-80-0212, NSF-IST-80-21358, NSF-IST-79-18264

NTIS Field/Group Codes: 62B (Computer Software)

Price: PC A03/MF A01

Availability: Available from the National Technical Information Service, Springfield, VA. 22161

Number of Pages: 11p

Keywords: *Relational data bases, Graph theory, Canonical forms, Trees(Mathematics), Query Languages, *Acyclic hypergraphs, Query processing, Graham reduction.

Abstract: We demonstrate a sense in which the equivalence between blocks (subgraphs without articulation points) and biconnected components (subgraphs in which there are two edge-disjoint paths between any pair of nodes) that holds in ordinary graph theory can be generalized to hypergraphs. The result has an interpretation for relational databases that the universal relations described by acyclic join dependencies are exactly those for which the connections among attributes are defined uniquely. We also exhibit a relationship between the process of Graham reduction of hypergraphs and the process of tableau reduction that holds only for acyclic hypergraphs.

May 1981

Report. No. STAN-CS-81-853



PB96-150669

Connections in Acyclic Hypergraphs

by

David Maier

Jeffrey D. Ullman


Research sponsored in part by

National Science Foundation
and

Air Force Office of Scientific Research

Department of Computer Science

Stanford University
Stanford, CA 94305



REPRODUCED BY: **NTIS**
U.S. Department of Commerce
National Technical Information Service
Springfield, Virginia 22161

CONNECTIONS IN ACYCLIC HYPERGRAPHS

David Maier†

SUNY, Stony Brook, N. Y.

Jeffrey D. Ullman†

Stanford Univ., Stanford, Calif.

We demonstrate a sense in which the equivalence between blocks (subgraphs without articulation points) and biconnected components (subgraphs in which there are two edge-disjoint paths between any pair of nodes) that holds in ordinary graph theory can be generalized to hypergraphs. The result has an interpretation for relational databases that the universal relations described by acyclic join dependencies are exactly those for which the connections among attributes are defined uniquely. We also exhibit a relationship between the process of Graham reduction [G] of hypergraphs and the process of tableau reduction [ASaU] that holds only for acyclic hypergraphs.

I. Definitions

A hypergraph $H = (\mathcal{N}, \mathcal{E})$ is a finite set of nodes $\mathcal{N}$ and a finite set $\mathcal{E}$ of edges, which are subsets of $\mathcal{N}$. We shall by default assume that a hypergraph is *reduced*, meaning that no edge is a subset of another edge, but nonreduced hypergraphs will be introduced explicitly at times.

A set N of nodes is *connected* if for every n and m in N there is a sequence of edges $E_1, \dots, E_k$, $k \geq 1$, such that n is in E_1 , m is in E_k , and for $1 \leq i < k$, $E_i \cap E_{i+1} \neq \emptyset$. A *component* of a hypergraph is a maximal connected set of nodes. We shall, for convenience, assume that all our hypergraphs are connected.

An *articulation set* is the intersection of two edges, say $X = E \cap F$, such that the removal of set of nodes X from the hypergraph, and therefore, from all edges containing such nodes, increases the number of components of the hypergraph.

If $H = (\mathcal{N}, \mathcal{E})$ is a hypergraph, and $N \subseteq \mathcal{N}$, then the *node-generated set of edges* generated by N is $\mathcal{F} = \{E \cap N \mid E \in \mathcal{E}\}$, with any edge that is a proper subset of another removed from $\mathcal{F}$. We may view a node-generated set of edges as a hypergraph whose set of nodes is the set that generates it. The edges in a node-generated set are, in general, subsets of the edges of the original hypergraph. To emphasize that an edge is a subset of an original edge, we shall use the term *partial edge* to refer to any subset of an edge of a given hypergraph.

A hypergraph is *acyclic* if every node-generated set of edges is either a single edge or has an articulation set; otherwise it is *cyclic*. This notion of acyclicity in hypergraphs is from [FMU], and a number of properties of the class were explored in [BFMY].

We should observe that our definition of "acyclic" is less restrictive than the standard one given by [B], or the definitions used by [G, Z] in connection with relational database schemes. However, we believe our use of the term is justified by its many natural properties, as indicated by [FMU, BFMY] as well as by the present paper.

II. Graham Reduction

The *Graham reduction* of hypergraph $H = (\mathcal{N}, \mathcal{E})$ is defined by applying the following two operations to H until they can be applied no more.

1. (Node removal) If a node n appears in only one edge, delete n from $\mathcal{N}$ and from that edge. Note the result may not be a reduced hypergraph, a problem that can always be remedied by the next operation.
2. (Edge removal) Delete one edge E from $\mathcal{E}$ if there is another edge F for which $E \subseteq F$.

The process of Graham reduction is from [G]. [BFMY] shows that a hypergraph reduces to nothing by this process if and only if it is acyclic. [YO] also discussed a similar algorithm and claimed it identified the "tree queries" of [BG], a class of (ordinary) graphs that are equivalent in a sense to acyclic hypergraphs [BFMY].

† Work partially supported by Air Force grant AFOSR-80-0212, in accordance with NSF agreement IST-80-21358 and NSF grant IST-79-18264.

† Work partially supported by Air Force grant AFOSR-80-0212, in accordance with NSF agreement IST-80-21358.




Fig. 1. An example acyclic hypergraph.

For our purposes, we need a modification of Graham reduction, where there is a set X of sacred nodes that may not be deleted by node removal. Let $GR(H, X)$ denote the result of applying rules (1) and (2) above to hypergraph H , with the node deletion modified to disallow the deletion of any node in X , the application of rules proceeding until such time as no more applications are possible. We shall write $H \rightarrow_N G$ and $H \rightarrow_E G$ when hypergraph H becomes hypergraph G by one application of node or edge removal, respectively. $H \rightarrow G$ means that H becomes G after one application of either node or edge removal, and $H \rightarrow^* G$ means that H becomes G after zero or more applications of these operations.

Lemma 1: $GR(H, X)$ is defined uniquely, independent of the order in which node and edge removals are applied.

Proof: Intuitively, if a rule is applicable to a node or edge, applying a rule elsewhere does not make the first inapplicable, so reduction proceeds to do everything that is ever possible, independent of the order of reductions chosen. More precisely, we can use the theory of "finite Church-Rosser systems" [ASeU] to construct a formal proof. By that paper, all we have to do is show that all reduction sequences are finite (obvious, since each step removes a node or edge), and that if $H \rightarrow H_1$ and $H \rightarrow H_2$, then there is a G such that $H_1 \rightarrow^* G$ and $H_2 \rightarrow^* G$.

We shall not go through the various cases, which are all easy, but shall give a sample. Let $H \rightarrow_N H_1$ by removal of node n , and $H \rightarrow_E H_2$ by deleting edge $E \subseteq F$. Then let G be H with node n and edge E deleted. Now n could not be in E , as node removal could eliminate it from H . Thus $H_2 \rightarrow^* G$, in fact, in this case, the sequence of zero or more applications of rules is exactly one application of (1). Also, in H_1 , E must still be a subset of F , so $H_1 \rightarrow^* G$ by one application of edge removal. $\square$

Example 1: Figure 1 shows an interesting example of an acyclic hypergraph H . Conventionally, we show edges by regions surrounding their member nodes. Suppose $X = \{A, D\}$. Then we can compute $GR(H, X)$ as follows. First, we may remove nonsacred nodes F and B , because they appear in only one edge. We cannot, of course, delete D , even though it appears only once, because it is sacred. Now, edge $\{A, E, F\}$ has become $\{A, E\}$, which is a subset of edge $\{A, C, E\}$, so we may delete the former by edge removal. Similarly, we may delete edge $\{A, B, C\}$, which has become $\{A, C\}$. The resulting hypergraph consists of edges $\{A, C, E\}$ and $\{C, D, E\}$; it cannot be further reduced. $\square$

III. Tableau Reduction

Another way in which hypergraphs can be transformed to simpler hypergraphs is by the process of *tableau reduction* [ASaU]. In our context, a *tableau* is a table whose columns correspond to the nodes of the hypergraph, in a fixed order, and whose rows consist of a *summary*, which for our purposes indicates the set of "sacred" nodes, and other rows (which we exclusively, from here on, call *rows*), that correspond to the edges. The rows consist of symbols, one for each column. For each column (node), there is a *special symbol*, and the special symbol for a column appears in exactly those rows whose edge contains the node. Further, certain of the special symbols appear in the summary entry for their column; these symbols are also designated *distinguished*. Everywhere else, a symbol that appears nowhere else is placed. All symbols that are not distinguished (including special symbols, perhaps) are *nondistinguished*.

Example 2: The tableau for Fig. 1 is shown in Fig. 2. The rows correspond to edges $\{A, B, C\}$, $\{C, D, E\}$, $\{A, E, F\}$, and $\{A, C, E\}$, in that order. The summary is placed at the top, between horizontal lines. In this tableau we adopt the useful convention that blanks represent symbols that appear nowhere else. We show the special symbols explicitly, whether or not they appear more than once; these symbols are $a, b, \dots, f$. We have assumed that nodes A and D are sacred, so symbols a and d are distinguished, as well as special. $\square$

a		d	
a	b	c	
		c	d
a			e
a		c	e

Fig. 2. Tableau for hypergraph of Fig. 1.

a		d	
		c	d
a		c	e

Fig. 3. Reduced tableau from Fig. 2.

We reduce tableaux by applying a mapping h from the rows to a subset of the rows, called the *target* subset, subject to the constraints that

1. If a row r is in the target subset, then $h(r) = r$.
2. If a symbol appears in two or more rows (in our tableaux, that symbol would be special and would appear in the same column in each row), say rows r_1 and r_2 , then rows $h(r_1)$ and $h(r_2)$ agree on that column.
3. If a row r has a distinguished symbol in a column, then $h(r)$ has the same symbol in that column.

Call such a mapping h a *row mapping*. By (2) it makes sense to treat h as mapping symbols to symbols as well as rows to rows. That is, we can define $h(a)$ to be the symbol appearing in the same column as a in all rows $h(r)$ such that row r has symbol a . The reader should consult [ASaU] for a proof that tableaux and their row mappings form a finite Church-Rosser system, and that therefore there is a unique (up to renaming of symbols) minimal subset of rows in any tableau such that there are no row mappings for those rows other than the identity, and such that there exists a row mapping from the full set of rows to this subset.

Tableaux as defined here are not as general as in [ASaU]. In particular, we can make an important observation about our tableaux, directly from the edge removal rule and the fact that only special symbols ever appear in more than one row.

Lemma 2: If a is a special symbol appearing in row r , h is a row mapping, and $h(r)$ does not have a in a 's column, then all rows with a are mapped by h to row $h(r)$. $\square$

Let h be a row mapping on the tableau for hypergraph H . We shall use $h(E)$ in place of $h(r)$ if r is the row of the tableau that corresponds to edge E . Let X be a set of sacred nodes for H . Then define $h(H)$ to be the set of partial edges formed by taking those edges E of H such that $h(E) = E$, i.e., the edges in the target of h , and deleting from those edges the nodes not in X that appear in only one of those edges. Note the resulting hypergraph may not be reduced.

Define $TR(H, X)$, for hypergraph H and set of sacred nodes X to be the set of partial edges formed as follows.

1. Construct the tableau for H with the special symbols for the nodes of X made distinguished.
2. Reduce that tableau by finding, according to [ASaU], a minimal set of rows that have only identity row mappings. By [ASaU], this set is unique up to renaming of symbols, although that renaming may at times exchange special and nonspecial symbols (an example is Fig. 1 with A and C sacred).
3. Let h be a row mapping from the rows of H 's tableau to its reduced version. Then $TR(H, X) = h(H)$. Note that $TR(H, X)$ will always be a reduced hypergraph, or we could have eliminated an additional row in the tableau minimization.

Example 3: Continuing with Example 2, a minimal set of rows is the second and fourth, shown in Fig. 3. The desired row mapping h sends rows 1, 3, and 4 to 4, and 2 to 2. Thus, for example, b is mapped to the nondistinguished symbol appearing in column 2 of row 4. Note that we left that entry blank, since the particular symbol appearing there is irrelevant, but a symbol does "exist" there.

Figure 3 represents the minimal set of rows, since its first row cannot map to its second (distinguished symbol d would be mapped to a nondistinguished symbol), nor can the second be mapped to the first (the

same problem with a prevents it). The resulting set of partial edges is $\{C, D, E\}$ for the first row and $\{A, C, E\}$ for the second. Note that c and e appear twice, and a and d are distinguished, so all special symbols do yield nodes in the partial edges, but a nondistinguished special symbol appearing only once would not cause its node to appear in a partial edge of $TR(H, X)$. $\square$

An important fact about Graham and tableau reductions is stated in the following lemma and theorem.

Lemma 3: Let H be a hypergraph, and X a set of sacred nodes. Then there is a row mapping from the tableau of H with the special symbols for X made distinguished to the subset of rows that correspond to the edges that remain after some Graham reduction of H with X sacred.[†]

Proof: The proof is an induction on the number of steps in the Graham reduction. The basis, zero steps, is trivial. Suppose that $H \rightarrow G \rightarrow *I$.

Case 1: $H \rightarrow_N G$, where node A is eliminated by node removal. Then A is not sacred, and the column for A in the tableau of H has no two symbols the same. By the inductive hypothesis, there is a row mapping from the tableau of G to a subset of its rows that corresponds to $GR(H, X)$. Since the Graham reductions of G and H are the same, and their tableaux differ only in that H has a column of different symbols (none of which is distinguished) that G does not have, the same row mapping is legal for H .

Case 2: $H \rightarrow_E G$, where edge $E \subseteq F$ is deleted by edge removal. Then the row mapping for the tableau of G can legally be extended to map E to the same row that F is mapped to. This row mapping serves to map H to its Graham reduction. $\square$

Theorem 1: Let H be an acyclic hypergraph with set X of sacred nodes. Then $GR(H, X) = TR(H, X)$.

Proof: With Lemma 3, we know that every Graham reduction can be duplicated by a row mapping. We must show by induction on the number of edges in H that for every row mapping h , there is a sequence of Graham reduction steps that yields $h(H)$. The basis, where H has only one edge, is trivial.

For the induction, assume H has more than one edge. By [BFMY], we may find an articulation set Y for H that breaks H into two components, one of which is a single edge E with Y removed.

Case 1: $h(E) = E$. Then the special symbols for the nodes in Y are mapped to themselves, so we may as well assume the nodes in Y are sacred. Let G be H with E removed. Then h is also a row mapping for G , even if all the nodes in $X - E \cup Y$ are sacred. By the inductive hypothesis, we can Graham-reduce the latter hypergraph to $h(G)$. Since all of Y was taken to be sacred, this reduction did not remove any of those nodes, and therefore the same Graham reduction is valid in H and yields the same result with edge E added. We can use node removal to remove from E those nodes in $E - Y - X$, which yields $h(H)$.

Case 2: $h(E) \neq E$. Then no node in $E - Y$ can be in X , or else a distinguished symbol would be mapped to a nondistinguished symbol in the row for $h(E)$. Thus we can use node removal to delete from H all nodes in $E - Y$. Since Y is an articulation set, there is some other edge F such that $E \cap F = Y$. Thus we can delete what remains of E by edge removal. We now have the hypergraph G , which is H with E removed. By the inductive hypothesis, we can obtain $h(G)$ by a Graham reduction. Since $h(G)$ is also $h(H)$, we are done. $\square$

Note that Theorem 1 is false for cyclic graphs. For example, let us take edges $\{A, B\}$, $\{A, C\}$, $\{B, C\}$, and $\{A, D\}$, with only D sacred. Then the tableau reduction consists only of D , since all edges can be mapped to $\{A, D\}$, yet all four edges remain when Graham reduction is attempted.

Another fact about tableau reduction is brought out in the next lemma.

Lemma 4: For any H and X , $TR(H, X)$ is a node-generated set of edges.

Proof: Let $TR(H, X) = (\mathcal{N}, \mathcal{E})$, and suppose there is an edge E of H such that $E \cap \mathcal{N}$ is not a subset of any edge in $\mathcal{E}$. Let h be the mapping from the edges of H to $\mathcal{E}$ that performs the tableau reduction, and suppose $h(E) = F$. Then surely $F \cap \mathcal{N}$ is in $\mathcal{E}$.

We claim $E \cap \mathcal{N} \subseteq F \cap \mathcal{N}$, for the special symbol for each node in $E - F$ is mapped by h to the symbol in the tableau's row for F . It follows that all edges with a node in $E - F$ are also mapped to F . Hence, the columns for the nodes in $E - F$ have no repeated symbol in the reduced tableau. We may conclude from the construction of $TR(H, X)$ that no symbol in $E - F$ can be in $\mathcal{N}$. But then, $E \cap \mathcal{N} \subseteq F \cap \mathcal{N}$, as we wished to prove, and E does not contradict the node-generated condition for $TR(H, X)$, as was to be proved. $\square$

[†] Note that the set of remaining edges is not necessarily unique, since we could at some time have two identical partial edges and eliminate either one by edge removal. The lemma refers to any set of remaining edges. The set of partial edges in $GR(H, X)$ is, of course, unique.

Corollary 1: If H is acyclic, so is $TR(H, X)$.

Proof: Directly from the definition of "acyclic," any node-generated subset of an acyclic hypergraph must itself be acyclic. Note that we could also have used Theorem 1 to prove this corollary, since it is easy to show that Graham reductions yield node-generated sets of edges. However, Lemma 4 is interesting in its own right, since it applies to all hypergraphs, not just acyclic ones. $\square$

We close this section with three simple lemmas about tableau reductions that we shall find useful later.

Lemma 5: If $X \subseteq Y$ then $TR(H, X) \subseteq TR(H, Y)$.

Proof: Any row mapping legal for hypergraph H when Y is the set of sacred nodes will also be legal when X is the set of sacred nodes. The result follows immediately from the finite Church-Rosser property of row mappings. $\square$

Lemma 6: Let h be a row mapping that reduces hypergraph H to $TR(H, X)$. If node n of H is a member of any edge E such that $h(E)$ is an edge that does not contain n , then n does not appear in $TR(H, X)$.

Proof: By Lemma 2, all edges containing n have their rows mapped to the row for $h(E)$. But then, in the reduced tableau there is no special symbol in the column for n . Thus the construction of the hypergraph $TR(H, X)$ will omit node n . $\square$

Lemma 7: If Y is an articulation set of hypergraph H , and N is a component of H after Y is removed, with $X \cap N = \emptyset$, then $TR(H, X)$ contains no node of N .

Proof: When reducing H with set of sacred nodes X , we can map all the edges of N into any one edge that contains Y , of which there must be at least two, since Y is an articulation set. If this edge itself has no node in N , we are done. If it has nodes of N , then it is now the only such edge. Therefore, the nodes of this edge that are in N do not appear in the hypergraph $TR(H, X)$. $\square$

IV. Technical Lemmas

Before getting to our main results, we must provide some preparation. The first lemma states that certain rings of edges really do imply the hypergraph is cyclic. The reader should be reminded that every ring of edges does not make the hypergraph in question cyclic, and Fig. 1 is a canonical example, where edges $\{A, B, C\}$, $\{C, D, E\}$, and $\{A, E, F\}$ form a "ring." The problem with this ring is that there is a single edge, $\{A, C, E\}$, that contains three intersections of edges.

Lemma 8: Let H be a (not necessarily reduced) hypergraph. If there is a sequence of $k \geq 3$ sets of nodes $N_0, N_1, \dots, N_{k-1}, N_k = N_0$ such that

- a) $N_i \cup N_{i+1}$ is a subset of some edge E_i of H for $0 \leq i < k$, and
- b) there are no i, j , and ℓ such that $0 \leq i < j < \ell < k$, and all of N_i, N_j , and N_ℓ are contained in one edge of H ,

then H is cyclic.

Proof: We shall suppose that H is the smallest counterexample; that is, H is acyclic, satisfies the hypotheses of the lemma, has as few edges as any counterexample, and has as few nodes as any counterexample with the same number of edges. Then surely H has at least one edge, and so is not Graham-reduced. Let $H \rightarrow G$ be the first step in reducing H . If the step is of type (1), surely it is not a node of some N_i that was eliminated, so the same sequence of sets of nodes proves that G is a smaller counterexample.

If the first step is of type (2), and no E_i is eliminated, then again G is a smaller counterexample. If $E_i \subseteq F$, where F is not among the E_j 's, then let F replace E_i , and we still have the conditions that prove G a smaller counterexample. Finally, if $E_i \subseteq E_j$, then E_j is a superset of all of N_i, N_{i+1}, N_j , and N_{j+1} . No matter what distinct values i and j have, there are at least three distinct sets of nodes among the four just mentioned, so condition (b) would not have been satisfied for H as we had supposed. $\square$

Lemma 9: If $H = (N, \mathcal{E})$ is acyclic, and Y is an articulation set for $TR(H, X)$, then Y is also an articulation set for H , and no components disconnected by Y in $TR(H, X)$ are connected in H .

Proof: To begin, we must show that Y is the intersection of two members of $\mathcal{E}$. We know that $Y = E \cap F$, where E and F are edges of $TR(H, X)$. By Theorem 1, E and F are in $GR(H, X)$. Surely $E \subseteq E_1$ and $F \subseteq F_1$ for some E_1 and F_1 in $\mathcal{E}$. But $E \cap F = E_1 \cap F_1$, since no node in $E_1 \cap F_1$ could be eliminated by Graham reduction. In proof, note that no such node can be eliminated by node removal as long as E_1




Fig. 4. Diagram for proof of Lemma 9.




Fig. 5. Acyclic hypergraph with two apparent paths between nodes

and F_1 remain in the Graham reduction. But these edges remain unless one becomes a subset of the other. That evidently never happened, since neither E nor F can be subsets of the other.

Now suppose that there are two components of $TR(H, X)$ with Y removed, say with sets of nodes N_1 and N_2 , that are not separated when Y is removed from H . Then there exists a sequence of edges of H , $E_1, \dots, E_k$, $k \geq 1$, such that

1. $E_1 \cap N_1 \neq \emptyset$,
2. $E_k \cap N_2 \neq \emptyset$,
3. $E_i \cap E_{i+1} \neq \emptyset$ for $1 \leq i < k$, and
4. none of the E_i 's are in the target of h , where h is a row mapping that yields $TR(H, X)$ from H .

The situation is shown in Fig. 4.

Let $h(E_1) = F$. Then F is disjoint from N_2 . For if not, it must be disjoint from N_1 . In proof, we observe that by Lemma 4, $TR(H, X)$ is node-generated. If F were disjoint from neither N_1 nor N_2 , then these sets of nodes would not be disconnected. If F were disjoint from N_1 , Lemma 6 would apply to $E_1 \cap N_1$, and we would conclude that these nodes are not in $TR(H, X)$, a contradiction of (1) above. We conclude that $F \cap N_2 = \emptyset$.

By Lemma 2, we can show by induction on i that $h(E_i) = F$ for all i . When $i = k$ we observe that Lemma 6 applies to the nodes in $E_k \cap N_2$, implying that these nodes are not present in $TR(H, X)$ and contradicting (2) above. Thus the sequence of edges $E_1, \dots, E_k$ cannot exist. $\square$

V. Canonical Connections and Independent Trees

It would be nice if we could develop a theorem that said of cyclic hypergraphs that "there are two paths between any two points," at least as long as the points in question were within one block, i.e., a component with no articulation sets. The trouble is that we must be very careful what we mean by "paths." For example, in the acyclic hypergraph shown in Fig. 5, there appear to be two distinct paths between nodes A and G , since we may eliminate either the second or the third of these edges.

Rather than try to modify the definition of "path," we have chosen to rely on the notion that acyclic graphs (and hypergraphs) really look like trees; there is a unique way to get between points. In cyclic graphs, on the other hand, we may pick a spanning tree, but then there is always some other way to get between at least one pair of nodes. Thus we shall develop the notion of "canonical connections" in hypergraphs, which correspond to the paths along trees. For the example above, all four edges will be in the canonical connection. We shall also develop the notion of an "independent path," which is a path that takes you outside the canonical connection.

Let $H = (N, \mathcal{E})$ be a hypergraph and $X \subseteq N$. The canonical connection for X in H is simply $TR(H, X)$. We denote this canonical connection by $CC_H(X)$, or just $CC(X)$ if H is understood. The canonical connection is intended, at least in the case that H is acyclic, to be the natural set of partial edges



Fig. 6. An independent tree.

with which to link the nodes in X .

We may also think about a tree T that links sets of nodes of H . In T , the nodes are sets of nodes of H , and the edges are pairs of sets of nodes of H that are contained within one edge of H . We shall call the nodes and edges of T *tree nodes* and *tree edges*. We also wish to have in our definition of such trees, a condition of minimality, which we shall phrase by requiring that no three tree nodes be contained within one edge of H .

More formally, let us define a *connecting tree* to be a collection of sets of nodes $\{N_1, \dots, N_k\}$ together with a tree structure on these sets, that is, a set of $k - 1$ pairs (N_i, N_j) such that, if the N_i 's were nodes and these pairs were edges, the resulting (ordinary) graph would be an undirected, unrooted tree. This connecting tree is said to be *for* the collection of sets of nodes at its leaves.

A connecting tree T for $\{N_1, \dots, N_n\}$ is said to be an *independent tree* if some tree node of T (treated as a set of nodes of H) is not wholly contained within the set of nodes of the hypergraph $CC_H(\cup_{i=1}^n N_i)$.

Example 4: Let H be the hypergraph with edges $\{A, B, C\}$, $\{C, D, E\}$, and $\{A, E, F\}$, that is, the hypergraph of Fig. 1 with edge $\{A, C, E\}$ removed. Let $X = \{A, C\}$ be the set of sacred nodes. Then $CC(X)$ consists of the single partial edge $\{A, C\}$, since there is a row mapping of all rows to the row for edge $\{A, B, C\}$, and from that row we need not take node B , because it is not sacred and its special symbol appears only once in the reduced tableau.

However, the collection of sets of nodes $\{\{A\}, \{E\}, \{C\}\}$ forms an independent tree with the simple structure shown in Fig. 6. The edges of H that supply the required tree edges are $\{A, E, F\}$ and $\{C, D, E\}$, of course. The tree is independent because one of its sets, namely $\{E\}$, is not contained within the canonical connection. Note that although the sets of nodes of H in Fig. 6 are all singletons, there is no requirement that they be so.

Note that if our hypergraph were the acyclic one of Fig. 1, then Fig. 6 would not represent an independent tree. The reason is that the edge $\{A, C, E\}$ would contain three of the sets of nodes on the tree. $\square$

A connecting tree in the form of a single path (no node of the tree is in more than two edges of the tree) is called a *connecting path*, and an independent tree that is a connecting path is called an *independent path*. Note that each independent path connects two sets of nodes, each of which is contained in an edge of the hypergraph (perhaps the same edge). Another useful observation about any connecting path $N_1, \dots, N_n$ is that it is possible that $N_1 \subseteq N_2$ or that $N_n \subseteq N_{n-1}$, but no other containment relationships can hold, or else three sets are contained in one edge. The next lemma points out that we can restrict our attention from independent trees to independent paths.

Lemma 10: If there is some collection of sets of nodes $N_1, \dots, N_n$ for which an independent tree T exists, then for the same hypergraph H there is a pair of sets of nodes for which an independent path exists in H .

Proof: The tree T must have a pair of leaves, say N_i and N_j , such that on the path in T between these leaves the tree node N occurs, and N is not contained within $CC(\cup_{\ell=1}^n N_\ell)$. By Lemma 5, $CC(N_i \cup N_j) \subseteq CC(\cup_{\ell=1}^n N_\ell)$. Thus the path in T between N_i and N_j is an independent path, containing the node N that is not contained in $CC(N_i \cup N_j)$. $\square$

VI. The Main Theorem

Theorem 2: A hypergraph H is acyclic if and only if for no pair of sets of nodes of H , say N and M , each contained within an edge of H , is there an independent path.

Proof: (if): We perform an induction on the number of edges in $CC(N \cup M)$.

Basis: One edge. Let the edge in the canonical connection be E , and let $N = N_1, \dots, N_k = M$ be the independent path. Surely $k \geq 3$, since N and M are each contained in E , and some N_i must not be. By the definition of an independent path, no three of the sets of nodes are contained in a single edge. Thus Lemma




Fig. 7. Diagram for if part of Theorem 2.

8 immediately implies that H is cyclic, a contradiction.

Induction: Suppose H has more than one edge. Then $CC(N \cup M)$ must have an articulation set Y . Let $\mathcal{N}_1$ and $\mathcal{N}_2$ be two sets of nodes in $CC(N \cup M)$ that exclude Y and are disconnected when Y is removed. ($\mathcal{N}_1$ and/or $\mathcal{N}_2$ may themselves consist of several connected subsets.) Figure 6 illustrates the configuration, including the positions of certain sets of nodes that we shall deduce shortly.

Assume without loss of generality that N includes a node of $\mathcal{N}_1$. Then surely N includes no node of $\mathcal{N}_2$, or $\mathcal{N}_1$ and $\mathcal{N}_2$ would not be disconnected. Similarly, M cannot contain nodes from both $\mathcal{N}_1$ and $\mathcal{N}_2$. If M were wholly contained within $\mathcal{N}_1$, then by Lemma 7, $\mathcal{N}_2$ would be empty. We conclude that N and M are on different sides of Y , as shown in Fig. 7, although either could, in principle, intersect Y .

Thus, the independent path $N = N_1, \dots, N_k = M$ includes nodes of both $\mathcal{N}_1$ and $\mathcal{N}_2$. There must therefore be some i for which $N_i \subseteq Y$, as if not, $\mathcal{N}_1$ and $\mathcal{N}_2$ would be connected by whatever edge contained N_{p-1} and N_p , where N_p is the first set along the path to contain a node in $\mathcal{N}_2$.

Suppose without loss of generality that a set N_j not contained in $CC(N \cup M)$ is found between N and N_i along the path, as suggested by Fig. 7. Further, let i be the smallest subscript greater than j such that $N_i \subseteq Y$, and let ℓ be the largest subscript less than j for which $N_\ell \subseteq Y$, or $\ell = 1$ if there is no such. By Lemma 6, since the nodes of $N_i \cup N_\ell$ are present in $TR(H, N \cup M)$, the row mapping that effects the reduction must not map the special symbol for any node in $N_i \cup N_\ell$ to another symbol. Thus these nodes may as well have been distinguished. That is, $TR(H, N \cup M) = TR(H, N \cup M \cup N_i \cup N_\ell)$.

Thus, by Lemma 5, $TR(H, N_i \cup N_\ell)$ consists of a subset of the nodes of $TR(N \cup M)$. By our selection of i, j , and ℓ , the portion of the independent path from N to M that lies between N_ℓ and N_i is an independent path for those two sets of nodes, as N_j is surely not contained within $CC(N_i \cup N_\ell)$. As $N_i \cup N_\ell \subseteq \mathcal{N}_1$, $CC(N_i \cup N_\ell)$ has no node of $\mathcal{N}_2$, by Lemma 7. We conclude that the pair of sets of nodes N_i and N_ℓ have a smaller canonical connection than N and M . The inductive hypothesis applies to them, contradicting the existence of the independent path for them and proving the induction.

(only if): Suppose not, and let H be a smallest counterexample, that is, a cyclic hypergraph in which there are no independent paths. By Lemma 7, we may assume H has no articulation sets at all, because if Y were an articulation set, and any of the components of H that result when Y is removed had an independent path, say between N and M , the canonical connection between N and M would be the same in H , and an independent path in a component is easily seen to be independent in H .

Let F and G be two edges of H such that $X = F \cap G$ is maximal, that is, no intersection of edges of H properly includes X . Since H has no articulation sets, it remains connected when X is removed. That is, there is some sequence of two or more sets of nodes, $M_1, \dots, M_k$, such that

1. $M_1 = F - X$,
2. $M_k = G - X$, and
3. For each i , $1 \leq i < k$, there is an edge E_i that contains $M_i \cup M_{i+1}$. Note that E_1 may be F and M_2 a subset of M_1 , and a similar statement holds about E_{k-1} .




Fig. 8. Diagram for only-if part of Theorem 2.

The arrangement is shown in Fig. 8.

Consider the canonical connection between $F - X$ and X , that is, $TR(H, F)$. It is easy to see that a tableau mapping of all rows to F is legal, so this canonical connection consists of exactly the edge F . We claim that the sequence of sets $M_1, \dots, M_k, X$ is an independent path from $M_1 = F - X$ to X . Since it surely has a set $M_k = G - X$ not contained within (or even intersecting) the canonical connection, we have only to prove that no edge of H contains three of the sets.

We shall show by induction on k that whenever there are, in H , edges F and G , with $F \cap G = X$, X maximal, and a sequence of nonempty sets of nodes $M_1, \dots, M_k, X$ such that $M_1 \subseteq F$, $M_k \subseteq G$, and there is an edge E_i containing M_i and M_{i+1} for $1 \leq i < k$, then either no edge of H contains three of the sets, or there is a shorter sequence of sets, perhaps between different F and G , but with all the above qualifications, forming an independent path between the first and last sets in the path.

Basis: $k = 2$. Then any edge containing three of the sets would have to contain M_1 , M_2 , and X . The intersection of this edge with F would be a proper superset of X , violating the maximality of X . Thus, when $k = 2$, the given path is surely an independent path between M_1 and X .

Induction: Suppose $k > 2$, and there is an edge E containing three of the sets. If none of these three sets are X , or if they are X , M_i , and M_j , for $i < j - 1$, then we can surely eliminate one or more sets from the path and invoke the inductive hypothesis to claim that there is a shorter independent path in H .

The only other possibility is that E contains X , M_i , and M_{i+1} for some i . If $i > 1$, let E replace F and let the sequence of sets begin with M_{i+1} . If $i = 1$, let E replace F and begin the sequence at M_2 . Again we may invoke the inductive hypothesis to prove there is an independent path shorter than $M_1, \dots, M_k, X$. $\square$

Corollary 2: A hypergraph is acyclic if and only if it has no independent trees.

Proof: Immediate from Theorem 2 and Lemma 10. $\square$

VII. Conclusions

The result in Theorem 2 is a statement of the way in which the proper definition of "acyclic hypergraph" allows us to generalize the theorem about (nontrivial) ordinary graphs that the absence of articulation points is equivalent to the existence of nonunique paths between some sets of points.

There is also an important interpretation in database theory. In [FMU] there appears the notion that queries over a universal relation are answered by joining all the "objects" (in essence, relations) in the database and applying the query to the join. Objects are sets of "attributes," and we can construct a hypergraph that represents the structure of the universal relation by letting the nodes be attributes and the edges be the objects. [FMU] provides the motivation.

When the query is of a type that can be represented by a tableau, as many are, tableau minimization can then be applied to the resulting query. The effect of this tableau minimization is to transform the query into one on the join of those objects that are in the canonical connection of those attributes that are mentioned in the query explicitly. In this world, the interpretation of Theorem 2 is that universal relations whose objects form an acyclic hypergraph are exactly those for which the set of objects connecting any set of attributes can be uniquely defined.[†]

Put another way, there is, for sets of objects forming acyclic hypergraphs, and only for these, is there

[†] This connection, the canonical connection, is unique in acyclic hypergraphs, as we have shown, but the reader should be aware of cases like Fig. 5, where subsets of the canonical connection can serve to connect the nodes in question.

the assurance that the above interpretation of queries about the universal relation will connect attributes in exactly the way the objects connect them. This observation is a strong argument for using universal relation semantics when the structure of the universal relation can be described adequately by an acyclic set of objects, but it is a warning that the straightforward implementation of queries will not work when the underlying structure is cyclic; then some additional semantics, such as proposed in [MU], must be applied.

Acknowledgements

We have profited from discussions of the matters covered in this paper held with Ron Fagin and Richard Hull.

References

- [ASaU] Aho, A. V., Y. Sagiv, and J. D. Ullman, "Equivalence of relational expressions," *SIAM J. Computing* 8:2 (1979), pp. 218-246.
- [ASeU] Aho, A. V., R. Sethi, and J. D. Ullman, "Code optimization and finite Church-Rosser systems," in *Design and Optimization of Compilers* (R. Rustin, ed.), Prentice Hall, Englewood Cliffs, N. J., 1972.
- [B] Berge, C., *Graphs and Hypergraphs*, North Holland, Amsterdam, 1973.
- [BFMY] Beeri, C., R. Fagin, D. Maier, and M. Yannakakis, "On the desirable properties of acyclic database schemas," manuscript in preparation, IBM, San Jose, Calif.
- [BG] Bernstein, P. A. and N. Goodman, "The theory of semijoins," TR CCA-79-27, Computer Corp. of America, Cambridge, Mass., 1979.
- [FMU] Fagin, R., A. O. Mendelzon, and J. D. Ullman, "A simplified universal relation assumption and its properties," RJ2900, IBM, San Jose, Calif., Nov., 1980.
- [G] Graham, M. H., "On the universal relation," technical report, Univ. of Toronto, Sept., 1979.
- [MU] Maier, D. and J. D. Ullman, "Maximal objects and the semantics of universal relation databases," TR-80-016, Dept. of C. S., SUNY, Stony Brook, N. Y., 1980.
- [YO] Yu, C. T. and M. Z. Ozsoyoglu, "An algorithm for tree-query membership of a distributed query," *Proc. COMPSAC-79*, pp. 306-312, 1979.
- [Z] Zaniolo, C., "Analysis and design of relational schemata for database systems," Ph. D. thesis, UCLA, 1976.



U.S. DEPARTMENT OF COMMERCE
Technology Administration
National Technical Information Service
Springfield, VA 22161 (703) 487-4650
