ON THE SATISFIABILITY OF DEPENDENCY CONSTRAINTS IN ENTITY-RELATIONSHIP SCHEMATA

Maurizio Lenzerini
Dipartimento di Informatica e Sistemistica
Università degli Studi di Roma "La Sapienza"
via Buonarroti 12, 00185 Roma, Italy

Paolo Nobili
IASI-CNR
Viale Manzoni 30,
00185 Roma, Italy

Abstract

The problem of satisfiability of a specific class of integrity constraints in data bases, namely the dependency constraints, is analysed. An Entity-Relationship model is used for expressing data schemata. In this model suitable types of dependency constraints, called cardinality ratio constraints, allow one to impose restrictions on the mappings between entities and relationships. We show that, as far as such a class of constraints is concerned, the usual notion of satisfiability is not sufficiently meaningful. For this reason we introduce the notion of strong satisfiability, ensuring that no entity or relationship is compelled to be empty in all of the legal instances of the schema. We propose to model the cardinality ratio constraints of a schema by means of a suitable linear inequality system and we show that a schema is strongly satisfiable if and only if there exist solutions for the associated system. Furthermore, we describe a method for discovering which are the sets of constraints that prevent a schema from being strongly satisfiable.

1. Introduction

Semantic integrity constraints in data bases are used to specify the rules which data have to satisfy in order to reflect the properties of the represented objects in the modeled real world.

Great attention has generally been devoted to a particular class of integrity constraints, the so-called dependency constraints, that are used to specify restrictions on the mappings between the data classes of a schema. They represent a very important and commonly occurring class of constraints ([3], [11]): functional and numerical dependencies (see [6] and [12]) in the relational model, as well as many types of existence constraints expressible in semantic data models (for example [1], [2] and [7]), are meaningful examples of such a kind of constraints.

Dependency constraints have mainly been studied from the perspective of data design, where the goal is to obtain a "good" schema with respect to the efficiency of data base operations. In such a context, the major issues that have been addressed are related to the implication of data dependencies, i.e. the problem of finding sound and complete inference systems for a given class of dependency constraints.

In this paper we deal with one important property of dependency constraints, namely their satisfiability. We remind the reader that the set of integrity constraints of a schema is said to be satisfiable if some instance of the schema (i.e. data base state) exists which satisfies them (in this case the schema itself is said to be satisfiable). It is interesting to observe that, although the concept of satisfiability can be very helpful in verifying the correctness of the data base design process, it has not been deeply addressed in the literature. We shall see in section 3 that, when dependency constraints are considered, the usual notion of satisfiability is not sufficient for capturing significant properties of a schema. In fact, although several instances of the schema may exist that satisfy a set of dependency constraints, it may happen that in all of such instances, some of the classes of the schema are invariably empty. For this reason we define a new property of a schema, called strong satisfiability, ensuring that each class is non-empty in at least one of its instances.

Our work is carried out in the context of the Entity-Relationship approach to data modeling ([5]). Such an approach has largely influenced many methodological proposals for enhancing the effectiveness and the correctness of information systems design ([4]).

Entity-Relationship based formalisms are now widely adopted in the so-called conceptual phase of data base design, whose goal is to obtain a complete, precise and implementation independent description of the objects to be represented in the data base. We shall refer to a particular data model, called Semantic Entity-Relationship Model (SERM, in the following), defined in [8], in which a specific kind of constraints, namely the cardinality ratio constraints, is provided for expressing dependencies between entities and relationships. In section 2 we describe the main characteristics of SERM that are useful for the subsequent sections.

The goal of our work is twofold:

- Providing necessary and sufficient conditions for a SERM schema to be strongly satisfiable. This aspect is dealt with in section 3.
- Presenting a method for discovering possible unsatisfiable sets of cardinality ratio constraints. Such a method is described in section 4.

2. The Data Model

In this section we briefly describe the characteristics of the SERM data model that are useful for the subsequent sections. We assume that the reader is familiar with the concepts and the terminology of the Entity-Relationship model.

An *entity* ("entity type" in [5]) denotes a set of individuals, called its instances, representing real world objects with common properties.

Relationships among entities are used to model logical associations among real world objects. A *relationship* ("relationship type" in [5]) denotes a set of individuals, called its instances: each element of such a set represents a logical association among a different combination of instances of the entities that are connected to the relationship. In the following, we use the term *class* to refer to an entity or a relationship. Since a relationship can be connected to the same entity more than once, the concept of *role* is introduced to distinguish different connections of the same entity with a relationship. More precisely, a role is a name which univocally determines the connection between an entity and a relationship.

For the purpose of this paper, a *SERM schema* consists of a set of entities, a set of relationships, a set of roles, and a set of cardinality ratio constraints, which are defined later in this section.

Using the common conventions of representing Entity-Relationship schemas by means of diagrams, we show an example of schema in fig. 1.

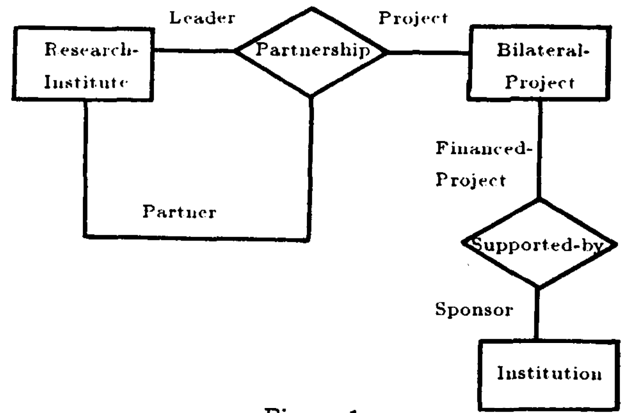

Figure 1.

Notice that in the diagram, roles are associated to the edges connecting the corresponding entities and relationships. The example concerns partnerships and financial supports for bilateral research projects. Each research institute can partecipate to bilateral projects either as project leader or as partner. For this reason, the entity "Research-Institute" is connected to the relationship "Partnership" through two different roles, "Leader" and "Partner".

A *subschema* of a SERM schema S is a schema constituted by a subset of the entities and relationships of S and satisfying the condition that every relationship is connected to the same set of entities and through the same roles as in S. For example, the classes "Research-Institute", "Partnership", and "Bilateral-Project", together with the roles "Leader", "Partner" and "Project" constitute a subschema of the schema shown in fig. 1.

In SERM, the concepts of attribute of entities and relationships and subset relationship between classes are also considered; however, they are not dealt with in the present work.

An *instance of a SERM schema* S is a finite collection of instances of the entities and the relationships of S, satisfying a set of rules (inherent constraints) to be described later. Each instance of a relationship R is linked to a combination of instances of the entities that are connected to R in the schema. The roles are used also at the instance level for identifying the links between relationship and entity instances. For example, an instance of the relationship "Partnership" of the schema shown in fig. 1 can be linked respectively to the instance a_1 of "Research-Institute" through the role "Leader", to the instance a_2 of "Research-Institute" through the role "Partner", and to the instance b_1 of "Bilateral-Project" through the role "Project".

We assume that no limit exists for the number of possible instances of entities and relationships. Furthermore, the instance of a schema in which all the classes have an empty set of instances is called *empty*.

In fig. 2 we give a representation of an instance of the schema shown in fig. 1.

Instances of "Research-Institute" : $\{a_1, a_2, a_3\}$
Instances of "Bilateral-Project" : $\{b_1, b_2\}$
Instances of "Institution" : $\{c_1, c_2\}$
Instances of "Partnership" : $\{p_1, p_2\}$
Instances of "Supported-by" : $\{s_1, s_2\}$
where:
$p_1 = \{<a_1.\text{Leader}>, <a_2.\text{Partner}>, <b_1.\text{Project}>\}$
$p_2 = \{<a_3.\text{Leader}>, <a_1.\text{Partner}>, <b_2.\text{Project}>\}$
$s_1 = \{ <b_1.\text{Financed-Project}>, <c_1.\text{Sponsor}> \}$
$s_2 = \{ <b_2.\text{Financed-Project}>, <c_2.\text{Sponsor}> \}$

Figure 2.

We write:

$$r = \{ <e_1,U_1>, <e_2,U_2>, \ldots, <e_m,U_m> \}$$

to denote the relationship instance r connected to the entity instances $e_1,e_2,\ldots,e_m$, respectively through roles $U_1,U_2,\ldots,U_m$. The pair $<e_i,U_i>$ is said to be the *component* of r corresponding to the role U_i.

Every instance I of a SERM schema must satisfy the following set of rules, called *inherent constraints*:

1. for each relationship R, for each instance $r = \{<e_1,U_1>,\ldots,<e_m,U_m>\}$ of R, for each i $(1 \leq i \leq m)$, e_i is an instance of the entity connected to R through role U_i.
2. for each relationship R, each instance of R has exactly one component corresponding to each role of R;
3. for each relationship R, different instances of R have different sets of components.

It is easy to verify that the above conditions are satisfied by the instance shown in fig. 2.

In SERM, a specific class of integrity constraints, namely the cardinality ratio constraints, are used to express dependency constraints between classes. Informally, a *cardinality ratio constraint* prescribes the minimum and the maximum number of instances of one relationship in which every instance of a connected entity must be involved for a given role. We shall write cardinality ratio constraints in the form:

$$E(U) \rightarrow^{(x,y)} R$$

where:

- E is an entity, R a relationship, and U a role; in particular, E is connected to R by means of role U;
- x is a non-negative integer, called the minimum cardinality of R with respect to E in the role U;
- y is either a positive integer or ∞, called the maximum cardinality of R with respect to E in the role U;
- $y \geq x$.

In order to characterize precisely the meaning of cardinality ratio constraints, we now state the conditions under which an instance of a schema satisfies a given constraint.

Definition 1. *Let $E(U) \rightarrow^{(x,y)} R$ be a cardinality ratio constraint of a schema S. Such a constraint is satisfied by an instance I of S if for every instance e of E in I, the number M of R instances connected to e by means of role U verifies:*

$$x \leq M \leq y$$

Notice that the value 0 for the minimum cardinality and the value ∞ for the maximum cardinality do not actually represent real constraints: however, for the purpose of this paper, they can be treated like any other values.

In the following, we shall write MIN(R,E,U) and MAX(R,E,U) to denote respectively the minimum and the maximum cardinality of the relationship R with respect to the entity E in the role U. If no cardinality ratio constraint $E(U) \rightarrow^{(x,y)} R$ appears in the schema, we assume that MIN(R,E,U)=0 and MAX(R,E,U)=∞. The set of cardinality ratio constraints of a schema S will be denoted by Γ^S. Every instance of S in which all the cardinality ratio constraints in Γ^S are satisfied will be called *legal*.

Referring again to the example of fig. 1, let's assume that the following cardinality ratio constraints are defined in the schema:

Research-Institute(Leader) $\rightarrow^{(0,2)}$ Partnership
Bilateral-Project(Project) $\rightarrow^{(1,1)}$ Partnership
Bilateral-Project(Financed-Project) $\rightarrow^{(1,1)}$ Supported-by
Institution(Sponsor) $\rightarrow^{(1,\infty)}$ Supported-by

The first constraint represents the rule which restricts research institutes to be project leaders of at most two bilateral projects. The second one imposes that each bilateral project is associated with exactly one project leader and one partner. By the third constraint, bilateral projects are supported by one and only one institution. Finally, the fourth constraint represents the fact that only the institutions supporting some project are meaningful for the application. It is easy to verify that the instance of fig. 2 satisfies all of the above constraints and, therefore, is legal.

3. Strong Satisfiability of SERM Schemata

As we said in the introduction, the usual notion of satisfiability is not sufficient for capturing interesting properties of a set of cardinality ratio constraints. In fact, since at least the empty instance of a schema S satisfies all the constraints in Γ^S, it follows that every SERM schema S is satisfiable with respect to Γ^S. On the other hand, it may happen that the cardinality ratio constraints of a SERM schema interact in such a way that no legal instance of the schema other than the empty one exists. Consider, for example, the schema shown in fig. 3 (in the diagram, the minimum and the maximum cardinalities of relationships are associated with the corresponding roles).

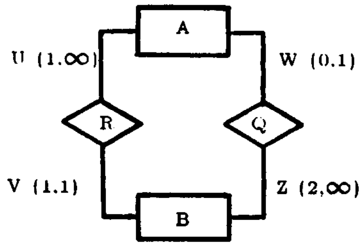

Figure 3.

It is easy to verify that only the empty instance is legal for such a schema: in fact, the constraints defined on relationship R impose that entity A cannot have more instances than entity B, whereas the constraints defined on relationship Q impose that the number of instances of A is at least two times the number of instances of B.

In the general case, the cardinality ratio constraints compel only some of the classes of the schema to be invariably empty. When this happens, we say that such classes cannot be *populated* in the legal instances of the schema. Since a class has to be considered meaningful only if the corresponding instances can be represented in the data base, we look for a new property, ensuring that all of the classes of the schema can be populated. We call such a property *strong satisfiability*.

Definition 2. *A SERM schema S is strongly satisfiable if for each class C of S, there exists at least one legal instance of S in which the set of instances of C is not empty.*

When a schema S is not strongly satisfiable, we shall say that both S and the set of cardinality ratio constraints Γ^S are unsatisfiable.

In the following, every instance of a schema in which no class is empty will be called *fully populated*. The next theorem shows that we can check for the strong satisfiability of a schema by looking for the existence of fully populated legal instances.

Theorem 1. *A SERM schema S is strongly satisfiable if and only if there is at least one fully populated legal instance of S.*

PROOF: *If-part.* It is evident that the existence of a fully populated legal instance of S implies that S is strongly satisfiable.

Only-if-part. If S is strongly satisfiable, then, for every class C_i of S, there exists a legal instance J_i of S in which C_i has a non-empty set of instances. Consider the instance J of S obtained from all the J_i's by means of the following rules: (1) the set of instances of the generic class C in J is the union of the instances of C in all the J_i's; (2) instances of classes coming from different J_i's are considered different. Clearly, J is fully populated. Furthermore, since all the J_i's are legal, J is legal too. QED

In order to characterize the strong satisfiability of SERM schemata we propose to model the cardinality ratio constraints of a schema by means of an associated linear inequality system. This system is defined in such a way that the existence of a fully populated legal instance of the schema (and, therefore, by theorem 1, the strong satisfiability of the schema) is reflected in the existence of some solutions for the corresponding system.

Definition 3. *Given a SERM schema S, we associate with it an inequality system Ψ^S whose unkowns and inequalities are defined as follows.*

Unknowns of Ψ^S:

- *an unknown $\hat{E}$ for each entity E of S;*
- *an unknown $\hat{R}$ for each relationship R of S.*

Inequalities of Ψ^S:

- *an inequality of the form:*

$$\hat{R} \geq x \cdot \hat{E}$$

for each cardinality ratio constraint $E(U) \rightarrow^{(x,y)} R$ in S, with $x \neq 0$;

- *an inequality of the form:*

$$\hat{R} \leq y \cdot \hat{E}$$

for each cardinality ratio constraint $E(U) \rightarrow^{(x,y)} R$ in S, with $y \neq \infty$;

- *an inequality*

$$\hat{E} > 0 \quad (\hat{R} > 0)$$

for each entity E (relationship R) of S.

Notice that by the above definition, Ψ^S is homogeneous (i.e. all of its constant terms are equal to zero) and has integer coefficients.

A preliminary result can now be proved regarding a sufficient condition for the strong satisfiability of a simple SERM schema, constituted by one relationship R connected to a collection of entities, each entity being possibly connected through several roles.

Lemma 2. *Let S be a schema constituted by one relationship R connected to m different entities $E_1,\ldots,E_m$. Assume that each E_i is connected to R through roles $U_{i1},U_{i2},\ldots,U_{iq_i}$, with $q_i \geq 1$, for each i $(1 \leq i \leq m)$. Given m+1 non-negative integer numbers $\rho,\sigma_1,\ldots,\sigma_m$ satisfying Ψ^S when substituted for the unknowns respectively associated with $R,E_1,\ldots,E_m$, a legal instance* *of S exists with σ_1 instances of E_1, σ_2 instances of $E_2, \ldots, \sigma_m$ instances of E_m and ρ instances of R if and only if the following condition holds:*

$$\rho \leq \prod_{i=1}^{m} \sigma_i^{q_i} \qquad (1)$$

PROOF: *If-part.* Suppose condition (1) holds. We now exhibit a legal instance I of S with σ_i instances of entity E_i for each i $(1 \leq i \leq m)$ and ρ instances of R. The way in which the instances of R are connected with the instances of $E_1, \ldots, E_m$ in I, is specified by means of the following procedure:

- number arbitrarily from 0 to $\rho - 1$ the instances of R, and, for each i $(1 \leq i \leq m)$, number arbitrarily from 0 to $\sigma_i - 1$ the instances of E_i;
- for i := 1 to m and for j := 1 to q_i do
  1. for k := 0 to $\rho - 1$ assign the component <e_{ih},U_{ij}> (e_{ih} is the h-th instance of E_i) to r_k (the k-th instance of R), where h = k mod σ_i;
  2. renumber the instances of R in such a way that instances having the same set of associated components are contiguous.

Notice that the inherent constraint 1 is trivially satisfied by I. Therefore, in order to show that I is legal, it remains to prove that the following conditions are satisfied by I:

α) every instance of R has exactly one component corresponding to each role U_{ij};
β) different instances of R have different sets of components;
γ) every cardinality ratio constraint of Γ^S is satisfied.

With regard to condition α, notice that the above procedure assigns exactly one component to each instance of R, for each role of R (step 1).

With regard to condition β, after the first execution of step 2 of the procedure, the largest group of instances of R linked to the same instance of E_1 by means of role U_{11} has at most:

$$\rho_1 = \lceil r/\sigma_1 \rceil \leq \sigma_1^{q_1-1} \cdot \prod_{i=2}^{m} \sigma_i^{q_i}$$

elements. After the second execution of step 2, the largest group of instances of R having the same set of component has at most:

$$\rho_2 = \lceil \rho_1/\sigma_1 \rceil \leq \sigma_1^{q_1-2} \cdot \prod_{i=2}^{m} \sigma_i^{q_i}$$

elements. Analogously, after the q_1-th execution, we have:

$$\rho_{q_1} = \lceil \rho_{q_1-1}/\sigma_1 \rceil \leq \prod_{i=2}^{m} \sigma_i^{q_i}$$

and, after the n-th execution (with $n = \sum_{i=1}^{m} q_i$), the largest group of instances of R having the same set of components has at most:

$$\rho_n = 1$$

elements. Therefore, condition β is satisfied by I.

With regard to condition γ, note that, since $\rho, \sigma_1, \ldots, \sigma_m$ constitute a solution for Ψ^S, we have:
MIN(R,E_i,U_{ij}) $\leq \lfloor \rho/\sigma_i \rfloor \leq \lceil \rho/\sigma_i \rceil \leq$ MAX(R,E_i,U_{ij}) for each i $(1 \leq i \leq m)$ and for each j $(1 \leq j \leq q_i)$. Now it is easy to see that for each entity E_i $(1 \leq i \leq m)$, for each role U_{ij} $(1 \leq j \leq q_i)$, and for each instance e_k of E_i $(0 \leq k \leq \sigma_i - 1)$, the number of instances of R that are associated with the component <e_k,U_{ij}> by the above procedure, is either $\lfloor \rho/\sigma_i \rfloor$ or $\lceil \rho/\sigma_i \rceil$. Therefore, condition γ is satisfied by I.

Only-if-part. Suppose that condition (1) doesn't hold, i.e. suppose that:

$$\rho > \prod_{i=1}^{m} \sigma_i^{q_i} \qquad (2)$$

Consider any instance J of S with σ_i instances of E_i $(1 \leq i \leq m)$ and ρ instances of R. Notice that the number of different sets of components that can be associated with an instance of R in J is $\prod_{i=1}^{m} \sigma_i^{q_i}$. If condition (2) holds, then there exist at least two instances of R in J having the same set of associated components and, hence, the above condition β is not satisfied in J. Therefore, J is not legal. QED

We are now ready to relate the strong satisfiability of a SERM schema S to the existence of solutions for the associated system Ψ^S. The following theorem gives a necessary and sufficient condition for the strong satisfiability of SERM schemata. Its proof makes use of the following lemma:

Lemma 3. *If a linear homogeneous inequality system H with rational coefficients admits a positive solution, then it also admits an integer positive solution.*

PROOF: Let $v_1, v_2, \ldots, v_n$ be the (possibly irrational) values assigned to the unknowns $x_1, x_2, \ldots, x_n$ of H by a positive solution X_0. By adding the following set of inequalities to H

$$x_i \geq b_i \qquad (1 \leq i \leq n)$$

where, for each i, b_i is any positive rational number less than or equal to v_i, we obtain a new inequality system H'. It is evident that the solution space of H' is a polyhedron that is not empty (in fact it includes at least X_0) and is contained in the solution space of H. Moreover, each vertex of such a polyhedron corresponds to a rational positive solution of H. An integer positive solution for H can be easily obtained by multiplying any of such rational solutions to a suitable number. QED

Theorem 4. *A SERM schema S is strongly satisfiable if and only if there exist solutions for the associated system* Ψ^S.

PROOF: *If-part.* From lemma 3 it follows that Ψ^S admits integer positive solutions. Let X be one of them. For each class C of S, let X[C] denote the value assigned by X to the unknown corresponding to the entity or relationship C. We can assume that for every relationship R of S, the following condition holds:

$$X[R] \leq \prod_{i=1}^{m} X[E_i]^{q_i} \qquad (3)$$

where $E_1, \ldots, E_m$ are the entities connected to R and q_i is the number of times entity E_i is connected to R. In fact, if some relationships existed for which condition (3) was not satisfied we could multiply X to a suitably large constant, obtaining another solution of Ψ^S such that condition (3) is satisfied for every relationship of S. Now, since condition (3) is satisfied for all the relationships of S, it follows from lemma 2 that for every subschema T of S, consisting of one relationship R and the connected entities $E_1, \ldots, E_m$, a legal instance I^T exists with X[R] instances of R, $X[E_1]$ instances of $E_1, \ldots, X[E_m]$ instances of E_m. Therefore, a legal instance I of S with X[C] instances for each class C can be easily obtained by merging the various I^T's and unifying the sets of instances of each entity belonging to more than one I_T. Since I is fully populated, we can conclude that S is strongly satisfiable.

Only-if-part. Suppose a legal fully populated instance I of S exists and consider one relationship R and one entity E connected to it by means of role U. Let σ and ρ be the number of instances of E and R in I. Since every cardinality ratio constraint is satisfied by I, the following condition holds:

$$\text{MIN(R,E,U)} \cdot \sigma \leq \rho \leq \text{MAX(R,E,U)} \cdot \sigma$$

In other words, the positive integers σ and ρ satisfy Ψ^S when substituted for the unknowns $\hat{E}$ and $\hat{R}$. By generalization, it is easy to see that we can construct a positive solution of Ψ^S by assigning to each unknown $\hat{C}$ the value corresponding to the number of instances of C in I. QED

The above result ensures that the problem of verifying the strong satisfiability of a SERM schema S can be solved in polynomial time with respect to the number of classes of S. In effect, by theorem 4, such a problem can be reduced to the one of testing the polyhedral cone defined by Ψ^S (i.e. its solution space) for non-emptyness, which can be done in polynomial time (see, for example, [10], pp. 170-185).

4. Analysis of Unsatisfiable Schemata

As already noticed in section 3, a SERM schema S which is not strongly satisfiable includes one or more sets of cardinality ratio constraints (i.e. subsets of Γ^S) that are unsatisfiable. As a result, some classes of S are compelled to be invariably empty in all of the legal instances of the schema. The goal of this section is to provide a method for discovering each unsatisfiable set of cardinality ratio constraints of a schema. Such a method can be very helpful for identifying those dependency constraints whose specification is erroneous.

In what follows we make use of graph concepts to represent SERM schemata. In particular, we show that, for a schema S which is not strongly satisfiable, information about the sets of unsatisfiable cardinality constraints can be obtained from a suitable analysis on a particular graph associated with S.

Definition 4. *Given a SERM schema S, the associated graph* G^S *is a directed multigraph <N,A> labeled in the arcs, where:*

- *the set of nodes N is in one-to-one correspondence with the set of classes of S;*
- *the set of arcs A is determined by the following rules: for each connection in S between an entity E and a relationship R through role U, two arcs* e_1 *and* e_2 *are in A;* e_1 *is directed from the node corresponding to E to the node corresponding to R and is labelled with MAX(R,E,U);* e_2 *is directed from the node corresponding to R to the node corresponding to E and is labelled either with 1/MIN(R,E,U), if MIN(R,E,U)≠0, or with* ∞, *if MIN(R,E,U)=0.*

Every graph associated with a SERM schema will be called an ER-graph. The label of an arc e will be denoted by $LABEL$(e). Moreover, if π is a path (or a cycle) of an ER-graph G^S, then the *weight* of π (denoted by $WEIGHT(\pi)$) is defined as follows:

$$WEIGHT(\pi) = \prod_{e \in \pi} LABEL(e).$$

If γ is a cycle, and $WEIGHT(\gamma)<1$, then we say that γ is *critical.*

An *assignment* ϕ for an ER-graph G=<N,A> is a mapping

$$\phi\colon \text{N} \longrightarrow \text{R}^+$$

associating positive rational numbers with its nodes.

An assignment is said to be *correct* if for each arc $e=<n_1,n_2>$ in A, the following condition holds:

$$\frac{\phi(n_2)}{\phi(n_1)} \leq LABEL(e).$$

An ER-graph is said to be *inconsistent* if no correct assignment exists for it, *consistent* otherwise.

We observe that, if π is a path from n_a to n_b (not necessarily distinct) of an ER-graph G, then $WEIGHT(\pi)$ represents an upper bound for $\frac{\phi(n_b)}{\phi(n_a)}$, for each correct assignment ϕ of G, i.e.:

$$\frac{\phi(n_b)}{\phi(n_a)} \leq WEIGHT(\pi). \qquad (4)$$

Notice that the assignment obtained by multiplying a correct assignment by a rational number is also correct. Therefore, whenever a correct assignment exists for a graph, an integer assignment (i.e. an assignment associating integer numbers to the nodes) also exists which is correct.

By the above definitions and by theorem 4, one can easily verify that the problem of checking a schema for strong satisfiability is isomorphic to the one of finding a correct assignment for the associated graph.

The goal of this section is to show that there exists a strict correspondence between sets of unsatisfiable cardinality ratio constraints in a schema and critical cycles in the associated graph. In particular, it will be shown by theorem 6 that a critical cycle is an inconsistent ER-graph and, on the converse, any inconsistent ER-graph contains a critical cycle. The following lemma introduces such a theorem.

Lemma 5. *Let G be a consistent ER-graph and n_a and n_b two of its nodes. Let Φ be the collection of all the correct assignments for G, and Π the collection of all paths from n_b to n_a in G. Then, it holds that*

$$\max_{\phi\in\Phi} \frac{\phi(n_a)}{\phi(n_b)} = \min_{\pi\in\Pi} WEIGHT(\pi)$$

where, if Π is empty, the right hand side is to be interpreted as ∞.

PROOF: It is obvious that $\max_{\phi\in\Phi} \frac{\phi(n_a)}{\phi(n_b)}$ is less than or equal to $\min_{\pi\in\Pi} WEIGHT(\pi)$, so that we can consider only the case where the left hand side is finite. Let ϕ' be a correct assignment such that $\frac{\phi'(n_a)}{\phi'(n_b)}$ is maximum, and suppose that $\frac{\phi'(n_a)}{\phi'(n_b)} < \min_{\pi\in\Pi} WEIGHT(\pi)$. This means that every path π_i in Π contains an arc e_i (say from p to q) such that $\frac{\phi'(q)}{\phi'(p)} < LABEL(e_i)$. It is easy to verify that the set of all such e_i's constitutes a cut (N_a,N_b) of G, with $n_a \in N_a$ and $n_b \in N_b$. Hence, by multiplying every $\phi'(n)$ (for $n \in N_a$) to a suitable positive number, we can obtain a correct assignment ψ such that $\frac{\psi(n_a)}{\psi(n_b)} > \frac{\phi'(n_a)}{\phi'(n_b)}$, contradicting the hypothesis on ϕ'. QED

Theorem 6. *An ER-graph G is inconsistent if and only if it contains a critical cycle.*

PROOF: *If-part.* Assume that G contains a critical cycle γ with weight w, and suppose that ϕ is a correct assignment for G. Then, for any node n in γ, it would hold that

$$\frac{\phi(n)}{\phi(n)} \leq w \qquad (5)$$

obtained by applying (4) to γ. Since (5) cannot be satisfied, we can conclude that no correct assignment exists for G.

Only-if-part. Assume that G is inconsistent. Let G′ be a maximal consistent subgraph of G (notice that such a subgraph always exists, since at least the graph obtained from G by eliminating all of its arcs is consistent). Let $e=<n_a,n_b>$ be any arc of G−G′. Lemma 5 ensures us that there is at least one path from n_b to n_a in G′ (otherwise we could find a correct assignment for $G' \cup \{e\}$, contradicting the hypothesis that G′ is maximally consistent). From lemma 5 again, it follows that there is a correct assignment ϕ for G′ and a path π from n_b to n_a such that $\frac{\phi(n_a)}{\phi(n_b)} = WEIGHT(\pi)$. Consider now the cycle γ constituted by π and e, and suppose that $WEIGHT(\gamma) \geq 1$, i.e. $\frac{1}{WEIGHT(\pi)} \leq LABEL(e)$. It follows that $\frac{\phi(n_b)}{\phi(n_a)} \leq LABEL(e)$, and, therefore, ϕ is correct also for $G' \cup \{e\}$. Since this contradicts the hypothesis that G′ is maximally consistent, it follows that γ is critical. QED

The above theorem shows that critical cycles are responsible for the inconsistency of a graph associated with a SERM schema. Taking into account the correspondence between the strong satisfiability of a SERM schema and the consistency of the associated graph, each critical cycle singles out an unsatisfiable set of cardinality ratio constraints, namely those corresponding to the labels of the component arcs. Notice that discovering critical cycles can be done in polynomial time, for example using a variant of the Floyd-Warshall algorithm for the determination of the shortest path between two nodes in a graph (see [10], pp. 129-133).

5. Conclusions

In this paper we have shown that significant properties of a SERM schema can be recognized by means of suitable computations performed on an associated inequality system and by an analysis on a corresponding graph. In particular, the results reported in section 3 show that strong satisfiability can be checked in polynomial time with respect to the size of the schema by looking for solutions of the associated system; moreover, in section 4 it is shown that, for a schema which is not strongly satisfiable, information about the sets of unsatisfiable cardinality ratio constraints can be obtained by discovering critical cycles in the corresponding graph.

Since realistic schemata could lead to corresponding sizeable systems and graphs, it is worth noting that some simplifications of the schemata can be adopted and should obviously be considered, for the sake of efficiency, when applying the described techniques. However, the analysis of the possibilities of improving the efficiency of the proposed methods is beyond the scope of this paper and is investigated in [9].

In the same paper, we study the problem of checking a schema for strong satisfiability when further capabilities of SERM are taken into account. In particular, we demonstrate that even small enrichments of the expressive power of the data model used in this paper, may result in a dramatical increase of the complexity of such a problem.

Acknowlegdements

We wish to thank Giorgio Ausiello for his useful comments on an earlier version of this paper.

References

[1] Abrial J.R., "Data Semantics", in J.W. Klimbie, and K.L. Koffeman, (eds.), *Data Base Management*, North-Holland Publishing Company, Amsterdam, 1974.

[2] Albano A., Cardelli L., and Orsini R., "Galileo: A Strongly Typed Interactive Conceptual Language, *ACM Transactions on Database Systems*, 10(2), 1985.

[3] Brodie M.L., "On the Development of Data Models", in Brodie M.L., Mylopoulos J., and Schmidt J.W. (eds.), *On Conceptual Modelling*, Springer-Verlag New York Inc., 1984.

[4] Ceri S. (ed.), "*Methodology and Tools for Data Base Design*", North-Holland Publishing Company, Amsterdam, 1983.

[5] Chen P.P.S., "The Entity-Relationship Model: Toward a Unified View of Data", *ACM Transactions on Database Systems*, 1(1), 1976.

[6] Grant J., and Minker J., "Numerical Dependencies", *Proc. of the International Conference on Logical Bases for Data Bases*, Tolouse, 1982.

[7] Hammer M., and McLeod D., "Database Description with SDM: A Semantic Database Model", *ACM Transactions on Database Systems*, 6(3), 1981.

[8] Lenzerini M., "SERM: Semantic Entity-Relationship Model", *Proc. of the Fourth International Conference on Entity-Relationship Approach*, IEEE Computer Society, 1985.

[9] Lenzerini M., and Nobili P., "On the Satisfiability of Numerical Dependencies", Dipartimento di Informatica e Sistemistica, Università di Roma "La Sapienza", Internal Report, 1987, forthcoming.

[10] Papadimitriou C.H., and Steiglitz K., "*Combinatorial Optimization*", Prentice-Hall, Englewood Cliffs, N.J., 1982.

[11] Tsichritzis D., and Lochovsky F., "*Data Models*", Prentice-Hall, Englewood Cliffs, N.J., 1982.

[12] Ullman J., "*Principles of Database Systems*", second edition, Computer Science Press, 1982.