

On the Power of Database Update Languages

Michael Lawley

School of Computing and Information Technology

Griffith University

Nathan, Queensland, Australia 4111

Abstract

A database update may be considered a mapping from one database instance to another database instance. A database update language L_1 is considered more powerful than a database update language L_2 if L_1 can express a superset of the mappings expressible in L_2 . We consider several database update languages described in the literature. We present them in a uniform notation with well defined semantics and examine their relative expressive power.

Keywords Database, update, expressive power

1 Introduction

We study the relative expressive power of several database update languages that have been described in the literature. In doing so, we attempt to classify their expressive power in the same way that Chandra⁴ classifies the expressive power of database query languages. A similar classification was done by Abiteboul and Vianu^{1,2} but we deal with a set of languages that is less expressive (and more practical) than theirs.

For any update language with alternation or iteration constructs, the expressive power of the language as a whole will be affected by the power of the language available for expressing the conditions in these constructs. We restrict our study to deterministic update languages, focusing on the

iteration constructs and how their power depends on the language available for expressing their conditions.

An understanding of the expressive power of update languages is essential for the efficient checking of integrity constraints. We intend to explore this application in subsequent work.

The paper is arranged as follows. Section 2 gives a model for database updates and introduces some necessary notation and concepts. Section 3 describes the semantics of several proposed update languages while Section 4 examines the relative power of these languages. Section 5 concludes with a summary of the results from Section 4 and suggestions for further research.

2 Preliminaries

Let \mathcal{U} be a *universal domain* consisting of a countably infinite set of constants. A *database* is a tuple $\langle R_1, \dots, R_n \rangle$, where each R_i is a finite named relation over \mathcal{U}^{k_i} for some $k_i \geq 0$. A *database update* is a mapping from one database, $B = \langle R_1, \dots, R_n \rangle$, to another, $B' = \langle R'_1, \dots, R'_n \rangle$, where each R_i and R'_i have the same arity. We require the mapping to be partially recursive and C -generic for some finite set of constants C . (See Abiteboul and Vianu² for an explanation of the C -genericity condition.)

Some of the update languages we consider contain iteration and alternation constructs which depend on some query. The power of the language used for expressing this query will affect the power of the update language. This is very important since, if we cannot predict the result of a particular query, then we will have difficulty in reasoning about an update involving that query. Figure 1 shows the relative power of the query languages we consider.⁴

Let L be an update language and M a query language. We write L_M to refer to the update language with alternation and iteration conditions expressed in that query language; e.g., L_{SPJ} and L_{FO} refer to L with iteration and alternation conditions specified in the query languages SPJ and FO respectively.

Let L_1 and L_2 be update languages. L_2 is at least as powerful as L_1 if every mapping expressible in L_1 is also expressible in L_2 . We write this as $L_1 \sqsubseteq L_2$. If the set of mappings expressible in L_2 is a strict superset of those expressible in L_1 , then L_2 is more powerful than L_1 and we write

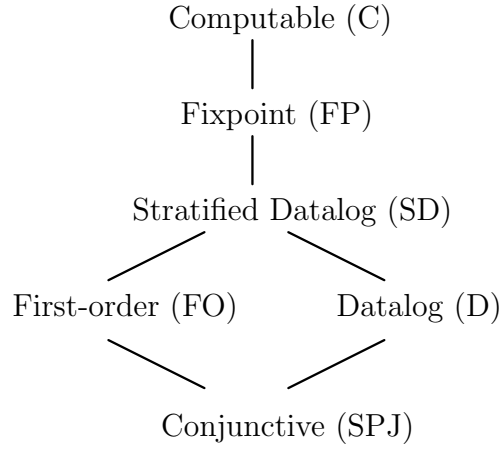


Figure 1: Expressive power of query languages.

$L_1 \sqsubset L_2$. If the set of mappings expressible in L_1 is not fully contained in the set of mappings expressible in L_2 and vice versa then the languages are incomparable and we have $L_1 \not\sqsubseteq L_2$ and $L_2 \not\sqsubseteq L_1$.

3 Some Update Languages

In this section we introduce a selection of update languages from the literature. We believe the selection covers most of the common update language features including the presence/absence of iteration; bounded/unbounded iteration; and general/restricted loop conditions. They are presented in a uniform syntax to make the differences and similarities between them more apparent.

In the following, let \bar{a} and \bar{c} be ground tuples, $F(\bar{x})$ a query in some language with free variables \bar{x} , $R(\bar{x})$ a named relation with unbound variables \bar{x} , and $S[\bar{x}]$ a statement in the update language with free variables \bar{x} .

3.1 LST – An Iteration free Language

In general, these consist of finite sequences of insertions and deletions.^{5,7} They have the standard requirement that a sequence may not both insert and delete the same tuple, and thus the order of operations is unimportant. Figure 2 gives the formal syntax.

$$\begin{aligned} stmt &::= stmt ; stmt \\ &| insert_R(\bar{a}) \\ &| delete_R(\bar{a}) \end{aligned}$$

Figure 2: Syntax of the iteration free language, LST.

3.2 SdetTL – An Iterator Based Language

Abiteboul and Vianu’s language SdetTL² is given in Figure 3.

$$\begin{aligned} stmt &::= stmt ; stmt \\ &| delete_R(\bar{a}) \\ &| insert_R(\bar{a}) \\ &| erase(R) \\ &| while \bar{x} : F(\bar{x}) do stmt \end{aligned}$$

Figure 3: Syntax of the iterator based language, SdetTL.

The *while* construct requires a little explanation. Its semantics are such that, at each iteration, *stmt* is performed in parallel for every \bar{x} satisfying $F(\bar{x})$. The result of each iteration is then the union of each parallel branch. The process is then repeated until $F(\bar{x})$ can no longer be satisfied or until *stmt* no longer has an effect (i.e. a fixpoint has been generated). This is termed *saturation*. It should be noted that these semantics require every branch to delete the same tuple for the deletion of that tuple to be effective, but only one branch need insert a tuple.

An alternative semantics, where *stmt* is performed for only one instantiation of \bar{x} satisfying $F(\bar{x})$ at each iteration, results in non-determinism.

Example 3.1 Consider the update $while \bar{x} : R(\bar{x}) do delete_R(\bar{x})$. With a deterministic semantics, one would expect this update to delete every tuple

in the relation R . However, since each parallel branch deletes a different tuple, the update has no effect on R . This means it is ineffective to put a *delete* in the loop.

On the other hand, with non-deterministic semantics, the update will delete every tuple in R in some, unspecified, order. \square

Non-deterministic languages such as detTL and TL (also proposed by Abiteboul and Vianu²) are outside the scope of this paper.

3.3 WL – A Set Based Language

Wallace’s proposed language¹⁰ is given in Figure 4.

$$\begin{array}{l}
 stmt ::= stmt ; stmt \\
 \quad | \quad insert_R(\bar{a}) \\
 \quad | \quad delete_R(\bar{a}) \\
 \quad | \quad replace_R(\bar{a}, \bar{c}) \\
 \quad | \quad if \ F \ then \ stmt \\
 \quad | \quad foreach \ \bar{x} : F(\bar{x}) \ do \ stmt
 \end{array}$$

Figure 4: Syntax of the set based language, WL.

The statement *foreach* $\bar{x} : F(\bar{x})$ *do* *stmt* is executed by evaluating $F(\bar{x})$ to produce a set of bindings for \bar{x} and then executing *stmt* in the following manner. If *stmt* is an atomic statement ($insert_R(\bar{x})$, $delete_R(\bar{x})$ or $replace_R(\bar{x}, \bar{y},)$), then execute it for each \bar{x} binding in parallel. If *stmt* is a sequence of statements, $stmt_1; \dots; stmt_n$, then execute $stmt_1$ for each \bar{x} binding in parallel, then execute $stmt_2$ for each \bar{x} binding in parallel, and so on. Otherwise, *stmt* is *foreach* $\bar{y} : F'(\bar{x}, \bar{y})$ *do* *stmt'*, where \bar{x} and \bar{y} are disjoint. Evaluate $F'(\bar{x}, \bar{y})$ for each \bar{x} binding to produce a new set of bindings for \bar{x} and \bar{y} . Then execute *stmt'* for the new set of bindings as described above.

Note that the condition of an *if* statement has no free variables, and consequently is a special case of the *foreach* statement. The *while* statement is shorthand for a *delete* followed by an *insert*.

Example 3.2 If the query, $F(x, y)$, produces the set of bindings $\{\langle a, b \rangle, \langle b, a \rangle\}$, then the update:

$$foreach \ x, y : F(x, y) \ do \ (insert_R(x) ; delete_R(y))$$

is equivalent to

$$\text{insert}_R(a) ; \text{insert}_R(b) ; \text{delete}_R(b) ; \text{delete}_R(a)$$

and not

$$\text{insert}_R(a) ; \text{delete}_R(b) ; \text{insert}_R(b) ; \text{delete}_R(a)$$

□

The following result will allow us to ignore nested *foreach* statements when comparing WL with other languages.

Lemma 3.1 *Any nested foreach construct can be flattened to a sequence of simple foreach statements containing only primitive updates.*

Proof sketch: By rewriting *ifs* as *foreach*'s and repeatedly applying the following transformations, we produce an equivalent flat statement.

1. $\text{foreach } \bar{x} : F_1(\bar{x}) \text{ do foreach } \bar{y} : F_2(\bar{x}, \bar{y}) \text{ do } S[\bar{x}, \bar{y}]$
 \Downarrow
 $\text{foreach } \bar{x}, \bar{y} : F_1(\bar{x}) \wedge F_2(\bar{x}, \bar{y}) \text{ do } S[\bar{x}, \bar{y}]$
2. Let R_T be a new empty relation with the appropriate arity.
 $\text{foreach } \bar{x} : F(\bar{x}) \text{ do } (S_1[\bar{x}] ; \dots ; S_n[\bar{x}])$
 \Downarrow
 $\text{foreach } \bar{x} : F(\bar{x}) \text{ do insert}_{R_T}(\bar{x}) ;$
 $\text{foreach } \bar{x} : R_T(\bar{x}) \text{ do } S_1[\bar{x}] ;$
 \vdots
 $\text{foreach } \bar{x} : R_T(\bar{x}) \text{ do } S_n[\bar{x}]$

These transformation will always preserve the original semantics since, within the body of a *foreach*, each operation in a sequence is performed for the entire set of loop tuples before the next one in the sequence. Hence, the effect of inserting, deleting or replacing a tuple multiple times is the same as doing it just once.

3.4 SS – A Non-Orthogonal Language

Figure 5 gives the non-orthogonal language proposed by Stemple and Sheard.⁹

$$\begin{array}{l}
stmt ::= stmt ; stmt \\
| insert_R(\bar{a}) \\
| delete_R(\bar{a}) \\
| foreach \bar{x} : F(\bar{x}) do delete_R(\bar{x}) \\
| foreach \bar{x} : R_1(\bar{x}) do insert_{R_2}(f(\bar{x})) \\
| foreach \bar{x} : R(\bar{x}) \wedge F(\bar{x}) do replace_R(\bar{x}, f(\bar{x})) \\
| foreach \bar{x} : R(\bar{x}) \wedge F_1(\bar{x}) do \\
| \quad if F_2 then delete_R(\bar{x}) else replace_R(\bar{x}, f(\bar{x})) \\
| if F then stmt else stmt
\end{array}$$

Figure 5: Syntax of the non-orthogonal language, SS.

Here, $f(\bar{x})$ is an arbitrary function which takes the input tuple \bar{x} and returns a tuple of appropriate type for the relation.

It can be seen that this is almost a subset of Wallace’s language. The only difference is the *if-then-else* statement which we deal with in Section 4.

3.5 $T - A$ “Query Based” Update Language

Abiteboul and Vianu² describe updates and queries as being instances of “database transformations” in the sense that they map database instances to other database instances. This idea is also presented by Qian.⁸ We now introduce the “T” family of update languages, allowing us to treat query languages as update languages.

For each relation R_i to be altered by an update, let $F_i(\bar{x})$ be a safe query such that the updated relation R'_i is described by the set $\{\bar{x} \mid F_i(\bar{x})\}$. So, for the update language T_{SPJ} , all the queries $F_i(\bar{x})$ are formulae in SPJ, and similarly for T_{FO} , T_D , etc.

We justify the restriction to safe queries by observing that updates can only operate on existing tuples, and thus are inherently range restricted.

Example 3.3 In T_{SPJ} , to insert the intersection of the relations R_2 and R_3 into R_1 we write

$$R_1 \leftarrow \{x, y \mid R_2(x, y) \wedge R_3(x, y)\}$$

□

Example 3.4 In T_{FO} , to delete all paths starting at m from the path relation R we write

$$R \leftarrow \{x, y \mid R(x, y) \wedge \neg(x = m)\}$$

This is neither a valid T_{SPJ} nor T_D update because it contains negation. \square

The purpose here is to get a feel for the power of the various update language constructs when combined with conditions expressed in each of the above query languages. Note that the relative power of T_{SPJ} , T_{FO} , T_D , etc. is the same as that of the query languages illustrated in Figure 1.

4 Relative Power

We first consider results that are independent of the query language used. We then examine the effect of the query language on the power of the update language.

4.1 General Results

Lemma 4.1 *LST is strictly less powerful than SS ($LST \sqsubset SS$).*

Proof sketch: LST does not contain any form of looping construct so an update which is to operate on every tuple in a relation cannot be independent of the current state of that relation. For example, LST cannot delete every tuple from a relation.

Lemma 4.2 *SS is strictly less powerful than WL ($SS \sqsubset WL$).*

Proof sketch: First, any SS update can be rewritten in WL. We need only consider *if-then-else* which can be rewritten as follows

$$\begin{aligned} & \text{if } F \text{ then } stmt_1 \text{ else } stmt_2 \\ & \quad \Downarrow \\ & \text{insert}_{R_T}(\text{else}) \\ & \text{if } F \text{ then (delete}_{R_T}(\text{else}) ; stmt_1) \\ & \text{if } R_T(\text{else}) \text{ then (delete}_{R_T}(\text{else}) ; stmt_2) \end{aligned}$$

where R_T is a new, initially empty, unary relation. Nested *if*'s will also work since the relation R_T is always empty before $stmt_i$ is performed and is left empty no matter which branch is taken.

For WL, we note that, when nested inside a *foreach*, the condition is evaluated for each variable binding in parallel, followed by all the *then* statements in parallel, followed by all the *else* statements in parallel since, as translated above, *if-then-else* consists of two sequential *if* statements. This is a natural semantics for WL, since the alternative of having the *then* and *else* branches performed in parallel introduces the possibility of non-determinism. Because of the constrained syntax of SS updates, the issue of the order of updates does not arise.

Second, there is no way to generate a “product” in SS. This is because the iterative constructs of SS may only range over a single relation and there is no way to nest these constructs.

Lemma 4.3 *WL is less powerful than SdetTL when the query language contains negation ($WL \sqsubseteq SdetTL$) and they are incomparable otherwise. The exact relationship between WL and SdetTL depends on the query language and is given below.*

Proof sketch: Any WL update can be rewritten as an SdetTL update. To do this we must be able to express *foreach* as a *while*. The mapping

$$\begin{array}{l} \text{foreach } \bar{x} : F(\bar{x}) \text{ do } stmt \\ \downarrow \\ \text{insert}_{R_T}(a) \\ \text{while } \bar{x} : F(\bar{x}) \wedge R_T(a) \text{ do } (\text{delete}_{R_T}(a) ; stmt) \end{array}$$

is incorrect since *stmt* may include a delete and, due to the semantics of the *while* requiring the delete to be performed in every parallel branch, this effectively nullifies the deletion. However, deletes can be implemented as follows:

- replace anything of the form $\text{delete}_R(\bar{a})$ by $\text{insert}_{R'}(\bar{a})$,
- replace all the corresponding references $R(\bar{x})$ by $(R(\bar{x}) \wedge \neg R'(\bar{x}))$.
- At the end of the update, copy all tuples in R and not in R' into another temporary relation R'' , erase R , and copy R'' into R .

Example 4.1 The following update

$$\text{foreach } \bar{x} : F(\bar{x}) \text{ do delete}_R(\bar{x})$$

becomes

$$\begin{aligned} & \text{insert}_{R_T}(a) \\ & \text{while } \bar{x} : F(\bar{x}) \wedge R_T(a) \text{ do } (\text{delete}_{R_T}(a) ; \text{insert}_{R'}(\bar{x})) \\ & \text{while } \bar{x} : R(\bar{x}) \wedge \neg R'(\bar{x}) \text{ do insert}_{R''}(\bar{x}) \\ & \text{erase}(R) \\ & \text{while } \bar{x} : R''(\bar{x}) \wedge \neg R(\bar{x}) \text{ do insert}_R(\bar{x}) \end{aligned}$$

Also, any reference $R(\bar{x})$ in $F(\bar{x})$ must be mapped to $R(\bar{x}) \wedge \neg R'(\bar{x})$. \square

In summary, these results give us $\text{LST} \sqsubset \text{SS} \sqsubset \text{WL} \sqsubseteq \text{SdetTL}$, the last inclusion only holding when the query language contains negation.

4.2 Query language specific results

We now examine how the power of the query language affects the relative expressive power of WL and SdetTL. We choose to omit LST and SS from this discussion since both WL and SdetTL are more general.

Lemma 4.4 T_{SPJ} is strictly less powerful than WL_{SPJ} ($T_{\text{SPJ}} \sqsubset WL_{\text{SPJ}}$).

Proof sketch: The *foreach* of WL allows us to express any T update as follows:

$$\begin{aligned} R & \leftarrow \{\bar{x} \mid F(\bar{x})\} \\ & \downarrow \\ & \text{foreach } \bar{x} : F(\bar{x}) \text{ do insert}_R(\bar{x}) \end{aligned}$$

However, T_{SPJ} cannot express the WL_{SPJ} update $\text{insert}_R(a)$, since this requires an increase in the cardinality of the relation R , which is not possible with only *select*, *project*, and *join*.

Lemma 4.5 T_{SPJ} is strictly less powerful than $SdetTL_{\text{SPJ}}$ ($T_{\text{SPJ}} \sqsubset SdetTL_{\text{SPJ}}$).

Proof sketch: Any T update can be rewritten as an SdetTL update as given above but with *while* implementing *foreach*. T_{SPJ} cannot express transitive closure³ while $SdetTL_{\text{SPJ}}$ can, as follows:

while $x, y, z : R(x, y) \wedge R(y, z)$ *do* $\text{insert}_R(x, z)$

Lemma 4.6 *SdetTL_{SPJ} is incomparable with WL_{SPJ}.*

Proof sketch: First, SdetTL_{SPJ} can express updates that WL_{SPJ} cannot. As outlined above, any WL update can be flattened into a (finite) sequence of unnested *foreach* statements. Since the transitive closure of a relation cannot be expressed as an SPJ query, any WL_{SPJ} update can only extend the relation by a fixed number of paths. SdetTL_{SPJ} as shown above, can express the transitive closure of a relation.

Conversely, WL_{SPJ} can express updates that SdetTL_{SPJ} cannot. Consider the following WL_{SPJ} update:

foreach $\bar{x} : R_1(\bar{x})$ *do* $\text{delete}_{R_2}(\bar{x})$

Because of the union-parallel semantics of *while*, the only way to reduce the size of a relation is to erase it then add tuples back. The tuples to be added back are exactly $R_2 \wedge \neg R_1$ which is not expressible with an SPJ query. Hence, this update is not expressible in SdetTL_{SPJ}.

Lemma 4.7 *WL_{SPJ} and T_{FO} are equivalent.*

Proof sketch: To convert from T_{FO} to WL_{SPJ}, convert the FO formula to existential normal form (writing each \forall as $\neg\exists\neg$). Negation can then be implemented with *delete*.

Example 4.2 The following update

$$R \leftarrow \{z \mid R_1(z) \wedge \forall x (R_2(x, z) \rightarrow R_3(x))\}$$

is written as

$$R \leftarrow \{z \mid R_1(z) \wedge \neg\exists x (R_2(x, z) \wedge \neg R_3(x))\}$$

which translates to

foreach $x, z : R_2(x, z)$ *do* $\text{insert}_{R_T}(x, z)$
foreach $x, z : R_3(x) \wedge R_T(x, z)$ *do* $\text{delete}_{R_T}(x, z)$
foreach $x : R(x)$ *do* $\text{delete}_R(x)$
foreach $z : R_1(z)$ *do* $\text{insert}_R(z)$
foreach $x, z : R_T(x, z)$ *do* $\text{delete}_R(z)$

□

To go the other way, from WL_{SPJ} to T_{FO} , flatten the WL_{SPJ} update as described in Lemma 3.1; each *foreach* can then be rewritten as an FO statement as follows.

$$\begin{aligned}
& \text{foreach } \bar{x} : F(\bar{x}) \text{ do } \text{insert}_R(\bar{x}) \\
& \quad \Downarrow \\
& R \leftarrow \{\bar{x} \mid R(\bar{x}) \vee F(\bar{x})\} \\
& \text{foreach } \bar{x} : F(\bar{x}) \text{ do } \text{delete}_R(\bar{x}) \\
& \quad \Downarrow \\
& R \leftarrow \{\bar{x} \mid R(\bar{x}) \wedge \neg F(\bar{x})\} \\
& \text{foreach } \bar{x}, \bar{y} : F(\bar{x}, \bar{y}) \text{ do } \text{replace}_R(\bar{x}, \bar{y}) \\
& \quad \Downarrow \\
& R \leftarrow \{\bar{x} \mid (R(\bar{x}) \wedge \neg F(\bar{x}, \bar{y})) \vee (R(\bar{y}) \wedge F(\bar{y}, \bar{x}))\}
\end{aligned}$$

Lemma 4.8 T_D is strictly less powerful than $SdetTL_{SPJ}$ ($T_D \sqsubset SdetTL_{SPJ}$).

Proof sketch: T_D cannot express the $SdetTL_{SPJ}$ update $\text{delete}_R(\bar{a})$. A T_D update involving a rule of the form

$$R_1(\bar{x}) \leftarrow R_2(\bar{y}) \wedge \cdots \wedge R_1(\bar{z})$$

can be rewritten in $SdetTL_{SPJ}$ as

$$\text{while } \bar{x} : R_2(\bar{y}) \wedge \cdots \wedge R_1(\bar{z}) \text{ do } \text{insert}_{R_1}(\bar{x})$$

where \bar{x} , \bar{y} , and \bar{z} are not necessarily disjoint.

Lemma 4.9 $SdetTL_{SPJ}$ is strictly less powerful than T_{SD} ($SdetTL_{SPJ} \sqsubset T_{SD}$).

Proof sketch: $SdetTL_{SPJ}$ cannot delete sets of tuples due to the semantics of the *while* statement (see Example 3.1). T_{SD} can, however, effectively delete tuples using negation. $SdetTL_{SPJ}$ updates can be rewritten in T_{SD} in a similar manner to Lemma 4.8.

Lemma 4.10 $SdetTL_D$ is equivalent to $SdetTL_{SPJ}$.

Proof sketch: The difference between D queries and SPJ queries is due to the presence of recursion, which is already (effectively) present in $SdetTL_{SPJ}$ in the form of the *while* statement.

Lemma 4.11 WL_{FO} , WL_{SPJ} and T_{FO} are all equivalent in power.

Proof sketch: The difference between SPJ queries and FO queries is due to the presence of disjunction and negation. These can already be simulated by sequence and deletion in WL_{SPJ} , so they add no power to the language WL_{FO} . The translation of WL_{FO} updates to T_{FO} updates and vice versa is trivial.

Lemma 4.12 WL_D , WL_{SD} , and T_{SD} are all equivalent in power.

Proof sketch: The negation present in WL_{SD} can be implemented by deletion in WL_D . The translation of WL_{SD} updates to T_{SD} updates and vice versa is trivial.

Lemma 4.13 $SdetTL_{FO}$ and T_{FP} are equivalent in power.

Proof sketch: $SdetTL_{FO}$ has negation, so can effectively perform deletes. The *while* statement provides a fixpoint operator. This means it can, for example, express *win*,⁶ which is the standard example of a query which can be written in FP and not SD.

Lemma 4.14 WL_{FP} , $SdetTL_{SD}$, $SdetTL_{FP}$, and T_{FP} are all equivalent in power.

Proof sketch: The *while* statement of $SdetTL_{SD}$ can be used to implement the FP queries of WL_{FP} , $SdetTL_{FP}$, and T_{FP} . Similarly, FP queries in WL_{FP} and T_{FP} can be used to implement the *while* statement of $SdetTL_{SD}$ and $SdetTL_{FP}$.

Lemma 4.15 WL_C , $SdetTL_C$, and T_C are all equivalent in power.

This is self-evident.

4.3 Summary of Results

The major observations to be made here are the relationships between the update language constructs (of WL and SdetTL) and the query language constructs (of T). We note that deletions can effectively be performed by inserting the difference of two sets of tuples using negation.

Figure 6 summarises the results presented in this section.




Figure 6: Expressive power of update languages.

5 Conclusion and Future Work

We have shown how the relative power of several update languages can be characterised and have provided a comparison with update languages closely related to a well known set of query languages. These results give a good indication of the difficulty in reasoning about the effects of updates in various languages and provide a basis for examining the relationship between integrity constraints and weakest preconditions.

On a more practical level, we still need to investigate the effect of introducing aggregate functions and arithmetic to the update languages. These capabilities are necessary in any actual implementation of a database update language.

Acknowledgement

Many thanks to Rodney Topor for valuable input on the topic of this paper and to Gill Dobbie for reading and commenting on earlier drafts.

References

- [1] S. Abiteboul and V. Vianu. Procedural and declarative database update languages. In *Proc. Seventh ACM Symposium on Principles of Database Systems*, pages 240–250, Austin, Texas, Mar. 1988.
- [2] S. Abiteboul and V. Vianu. Procedural languages for database queries and updates. *Journal of Computer and System Sciences*, 41(1):181–229, 1990.
- [3] A. V. Aho and J. D. Ullman. Universality of data retrieval languages. In *Conference Record of the Sixth ACM Symposium on Principles of Programming Languages*, pages 110–120, San Antonio, Texas, Jan. 1979.
- [4] A. K. Chandra. Theory of database queries (extended abstract). In *Proc. Seventh ACM Symposium on Principles of Database Systems*, pages 1–9, Austin, Texas, Mar. 1988.
- [5] L. J. Henschen, W. W. McCune, and S. A. Naqvi. Compiling constraint-checking programs from first-order formulas. In H. Gallaire, J. Minker, and J.-M. Nicolas, editors, *Advances in Data Base Theory*, pages 145–169. Plenum Press, New York, 1984.
- [6] P. Kolaitis. The expressive power of stratified logic programs. manuscript, Department of Computer Science, Stanford University, 1987. To appear in *Information and Computation*.
- [7] J. W. Lloyd, E. A. Sonenberg, and R. W. Topor. Integrity constraint checking in stratified databases. *Journal of Logic Programming*, 4(4):331–343, Dec. 1987.
- [8] X. Qian. *The Deductive Synthesis of Database Transactions*. PhD thesis, Stanford University, Nov. 1989.
- [9] T. Sheard and D. Stemple. Automatic verification of database transaction safety. *ACM Transactions on Database Systems*, 14(3):322–368, Sept. 1989.
- [10] M. Wallace. Compiling integrity checking into update procedures. In *Proc. Twelfth International Joint Conference on Artificial Intelligence*, pages 903–908, Aug. 1991.