

Incremental Tabling in Support of Knowledge Representation and Reasoning

Terrance Swift

Coherent Knowledge Systems, Inc. and NOVALincs, Universidade Nova de Lisboa
(e-mail: terranceswift@gmail.com)

submitted 1 January 2003; revised 1 January 2003; accepted 1 January 2003

Abstract

Resolution-based Knowledge Representation and Reasoning (KRR) systems, such as Flora-2, Silk or Ergo, can scale to tens or hundreds of millions of facts, while supporting reasoning that includes Hilog, inheritance, defeasibility theories, and equality theories. These systems handle the termination and complexity issues that arise from the use of these features by a heavy use of tabled resolution. In fact, such systems table by default all rules defined by users, unless they are simple facts.

Performing dynamic updates within such systems is nearly impossible unless the tables themselves can be made to react to changes. Incremental tabling as first implemented in XSB (Saha 2006) partially addressed this problem, but the implementation was limited in scope and not always easy to use. In this paper, we introduce *transparent incremental tabling* which at the semantic level supports updates in the 3-valued well-founded semantics, while guaranteeing full consistency of all tabled queries. Transparent incremental tabling also has significant performance improvements over previous implementations, including lazy recomputation, and control over the dependency structures used to determine how tables are updated.

1 Introduction

Tabled Logic Programming has supported a variety of applications that would be difficult to implement in Prolog alone, including model checking, program analysis, ontology-based deductions and decision making for collaborative agents. Typically such applications are written mainly as Prolog programs, but with a subset of the predicates tabled in order to support termination, reduce complexity, to use well-founded negation or to exploit other features.

However, systems such as Flora-2 (Yang et al. 2013) and its extensions: Silk (cf. silk.semwebcentral.org), Ergo (cf. coherentknowledge.com/publications) and the RAVE system (cf. www.sri.com/about/people/grit-denker) have been recently developed for knowledge representation and reasoning (KRR), and rely on tabled resolution for their computational underpinning. For instance, Flora-2 (Yang et al. 2013), which is based on XSB (Swift and Warren 2012), supports the non-monotonic inheritance of F-logic, prioritized defeasibility with multiple levels of conflicts, rule identifiers, function symbols, logical constraints, and HiLog. Silk and Ergo, both based on Flora-2, support all of the above features plus *omni axioms*, which are contrapositional rules whose bodies and heads are comprised of any formulas that can be supported by the Lloyd-Topor transformation (Lloyd and Topor 1984).

As an example of using these features, given the sentence: *A contractile vacuole is inactive in an isotonic environment* from (Reece et al. 2010), a tool called Linguist (www.haleyai.com) produces a Silk or Ergo formula in a mostly automatic manner (knowledge engineers may have to choose between translations in ambiguous cases), resulting in the axiom:

$$\begin{aligned}
& \text{forall}(\text{?x6})^{\wedge} \text{contractile}(\text{vacuole})(\text{?x6}) \\
& \implies \text{forall}(\text{?x9})^{\wedge} \text{isotonic}(\text{environment})(\text{?x9}) \\
& \implies \text{inactive}(\text{in}(\text{?x9}))(\text{?x6});
\end{aligned}$$

Such an axiom is next translated into several Flora-2 rules about conditions of contractile vacuoles, inactive contractile vacuoles, and isotonic environments. These Flora-2 rules are then transformed to support HiLog, defeasibility and other features, resulting in numerous normal rules executed in XSB. Once a knowledge base has been constructed from axioms such as the one above, queries can be made such as: *If a Paramecium swims from a hypotonic environment to an isotonic environment, will its contractile vacuole become more active?* The translation of queries is similar to that of knowledge, but may include *hypothetical* information, e.g., that *?x* is a Paramecium swimming from a hypotonic environment to an isotonic environment. Knowledge bases themselves are built from a collection of rules and omni axioms usually written by different knowledge engineers using a shared background vocabulary. The limited coordination among knowledge engineers is critical for producing knowledge bases at a low cost.

All of the the KRR-systems mentioned above employ what may be called *pervasive tabling* where a predicate is tabled unless it is explicitly declared non-tabled. Such programs have an operational behavior that is vastly different from (tabled) Prolog. Among other matters, as many of these tables represent background knowledge, it is critical for good system performance to reuse tables between queries. However, because queries may include hypothetical knowledge, and because knowledge bases are created by interactively adding or modifying rules, good performance demands the use of *incremental tabling* (Saha and Ramakrishnan 2005; Saha 2006).

The main idea behind incremental tabling is to maintain an *Incremental Dependency Graph (IDG)*, indicating how tabled goals depend both on dynamic code and on one another. When an update is made to dynamic code, the IDG is traversed, and affected tables are updated if necessary. However, while previous versions of incremental tabling were robust enough to support a commercial application (Ramakrishnan et al. 2007), they were not sufficient to support high-level KRR applications. Most significantly, a programmer had to decide when tables were updated: either an update was forced immediately upon an assert or retract, or the programmer performed “bulk” updates, after which a command propagated the updates to all affected tables. This methodology was complicated and had semantic drawbacks: unless an update was manually invoked, there was no guarantee that tables would be updated and no provision for stronger forms of view consistency. In fact, because of the brittleness caused by the need for low-level control along with other drawbacks, previous versions of incremental tabling, (designated here as *manual incremental tabling*) were suitable only for careful use by tabling experts.

Support for pervasive tabling requires that a tabling engine be redesigned in several ways, including the mechanisms whereby tables are updated. This paper introduces *transparent incremental tabling* to support applications that rely on pervasive tabling such as the KRR-systems described above. The papers major contributions are:

- A description of core changes that allow table updates to be made in a safe and efficient manner: first, tables are updated automatically and efficiently by *lazy recomputation*; second, updates always guarantee view consistency for incremental tables.
- A description of how incremental recomputation is extended to support updates according to the three-valued well-founded semantics.
- Introduction of the notion of *IDG abstraction* to reduce the size of the IDG when necessary.
- Detailed performance analyses of transparent incremental tabling for both small program




Fig. 1. A Program P_{inc} , and schematic Incremental Dependency Graph (IDG) for the query $t_1(X)$

fragments and for KRR-style examples over Extensional Databases (EDBs) up to size $\mathcal{O}(10^7)$. These results indicate that transparent incremental tabling efficiently supports the KRR uses previously mentioned, and may also provide a basis for *reactive KRR*.

Transparent incremental tabling is available in the current version of XSB. In addition to the extensions mentioned above, its implementation is based on a significant rewriting of the previous implementation of manual incremental tabling. Incremental tabling is not yet available in tabling engines other than XSB. However, while transparent incremental tabling adds data structures such as the IDG, it interfaces with a tabling engine mostly through routines for maintaining table space. Accordingly, most of the features described below are relatively portable, as tabling engines have similar table space operations, and sometimes similar data structures.

2 A Review of Manual Incremental Tabling

In this section we describe the previous version of incremental tabling using the main data structures and algorithms of (Saha 2006), which form the starting point for the features of transparent incremental tabling described in later sections. The description is as self-contained as possible, but sometimes uses the terminology of the SLG-WAM (Sagonas and Swift 1998).

Fig. 1 shows an XSB program P_{inc} where predicates are declared to use incremental tabling. In general both tables and dynamic code may be declared with various attributes: not only *incremental* as here, but also *subsumptive*, *trie-indexed*, and so on. Note that *tnot/1* is an XSB operator for tabled negation. Execution of the query $t_1(X)$ creates the *Incremental Dependency Graph (IDG)* schematically shown in Fig. 1. The IDG has a node for each tabled subgoal but not for non-tabled subgoals such as $nt_1(X)$ – though the bindings made by the rules for $nt_1/1$ are implicitly propagated. Leaf nodes in the IDG correspond to predicates such as $p/1$ and $q/1$ that are declared to be both dynamic and incremental. Each downward edge in a IDG represents an element of the *direct dependency* relation; the inverse relation is the *direct affected* relation. Note that paths in the IDG may be cyclic.

At the level of data structures, each node in the IDG is represented via an *IDG node frame* (Fig. 2). For a tabled incremental subgoal t/n , the IDG node frame is created by the *tabletry* instruction, by registering it into the subgoal trie for t/n ¹, and linking it with the *subgoal frame*, which contains information about each tabled subgoal. For dynamic incremental subgoals a new

¹ In XSB, the default data structure for tabled subgoals and their answers is based on tries (Ramakrishnan et al. 1999). While XSB offers basic support for answers that are “hash-consed” (Zhou and Have 2012) and not maintained as tries, our presentation assumes subgoal and answer tries throughout.

SLG-WAM instruction, `try_dynamic_incremental` performs these tasks. Each time a (tabled or dynamic) incremental subgoal S is called, the IDG may be updated. If S is new, an IDG node frame is created; also whether or not S is new, if S has a nearest tabled subgoal S_{par} as an ancestor, edges between S and S_{par} are added if not already present. As answers are derived for S , their count is maintained in the *nbr_of_answers* field of the IDG node frame.

<i>affected_edges</i>	Subgoals that this subgoal directly affects
<i>dependent_edges</i>	Subgoals upon which this subgoal directly depends
<i>subgoal_frame</i>	Pointer back to the subgoal frame
<i>nbr_of_answers</i>	Counts the number of answers rederived
<i>previous_IDG_node</i>	Used to determine if re-evaluation has changed the set of answers
<i>new_answer</i>	set to true if a new answer has been derived
<i>falsecount</i>	determines whether subgoal is valid

Fig. 2. The IDG node frame for incremental tables

At a high level, the use of the IDG is easy to understand. If a fact, say $p(g(2))$, is asserted, the incremental update subsystem must call `traverse_affected_nodes()` (Fig. 3) to traverse the IDG. Separate traversals start from each leaf node with which $p(g(2))$ unifies, and the traversals will increment the *falsecount* field of their IDG node frame (cf. Fig. 2), marking them as *invalid* (i.e., having a *falsecount* greater than 0). As it is unclear whether sensible semantics can be given to updating a subgoal that is incomplete (i.e., that is still being computed), a permission error is thrown if this is attempted. In our running example, assuming that no nodes in the IDG are already invalid, the algorithm will traverse depth-first through all nodes affected by $p(g(X))$ (directly or indirectly). In so doing, the affected non-leaf nodes are added to a global *invalid list* in the same order. In our example, the nodes for $t_5(X)$, $t_4(X)$ and $t_1(X)$ are traversed, and the invalid list represents this sequence.

Several properties of the traversal are worth noting. First, use of the *falsecount* field in `traverse_affected_nodes()` prevents the same node from being traversed multiple times. Also, note that invalidation simply represents *some* change in the underlying data so that retracts are handled in the same manner as asserts, and both positive and negative dependencies are treated in the same way. In fact, since the traversal starts with dependency leaf nodes that unify with a given atom, propagation of a rule update is handled in the same manner as a fact update: `traverse_affected_nodes()` is invoked for leaf nodes that unify with the rule head. In either case, the unification of leaf nodes with a given atom can also prevent unnecessary updates: for instance, if the fact $q(g(2))$ were added, it would not cause any update, since no leaf node of the IDG unifies with this fact.

After the invalidation phase is finished, reevaluation of the affected nodes may be done either immediately, or at a later time through an explicit command. Note that once the invalid list has been set up, the affected tables can be updated in a bottom-up manner simply by removing them in order from the list. Specifically, for each IDG node $IDGN$ removed from the invalid list, `incremental_reeval(IDGN)` called (Fig. 3). If $IDGN.falsecount$ is 0, the subgoal does not need to be recomputed. Otherwise, the answers for T , the table associated with $IDGN$, are marked as deleted, although their space is not reclaimed². A new IDG node $IDGN_{new}$ is created for T , and its *previous_IDG_node* field is set to the old IDG node, $IDGN$ (cf. Fig. 2). The subgoal

² The *answer list* of an answer trie, which allows easy traversal of all answers in the trie, is reclaimed at the completion of each non-incremental table, but retained by incremental tables for traversals during re-evaluation.

```
|  |  |
| --- | --- |
| traverse_affected_nodes(IDG node frame IDGN) | |
| /* IDGN is the IDG node frame for an incrementally tabled predicate */ | |
| If the table associated with IDGN is not completed, throw a permission exception | |
| For each IDGNaff that is directly affected by IDGN | |
| IDGNaff.falsecount++; | |
| If (IDGNaff.falsecount == 1) traverse_affected_nodes(IDGNaff) | |
| Add IDGN to the global invalid list. | |
| incremental_reeval(IDG node frame IDGN) | /* S is the subgoal to be recomputed */ |
| If IDGN.falsecount > 0 | |
| Let ST be the subgoal frame associated with IDGN (i.e., ST = IDGN.subgoal.frame) | |
| For each A in ST.answer_list | |
| Mark A as deleted, but do not adjust answer trie choice points or reclaim space | |
| Create a new IDG node IDGNnew for ST | |
| IDGNnew.new_answer := false; IDGNnew.falsecount = IDGNnew.nbr_of_answers = 0 | |
| Call S and for each new derived answer Aderiv | |
| Increment IDGNnew.nbr_of_answers | |
| If Aderiv was marked as deleted, remove the deletion mark | |
| Else IDGNnew.new_answer = true | |
| After completion of S, for each A in ST.answer_list | |
| If A is still marked as deleted, remove A from ST.answer_list | |
| Reset answer trie choice points and reclaim space for A | |
| If IDGNnew.new_answer = false and IDGNnew.nbr_of_answers = IDGN.nbr_of_answers | |
| propagate_validity(IDG node frame IDGN) | |
| propagate_validity(IDG node frame IDGN) | |
| For each IDGNaff that is directly affected by IDGN | |
| IDGNaff.falsecount- - | |
| if IDGNaff.falsecount == 0 propagate_validity(IDGNaff) | |

```

Fig. 3. Schematic algorithms for manual incremental tabling

for *T* is re-evaluated, and for each answer *A*, *IDGN_{new}.nbr_of_answers* is incremented; in addition if *A* is new, (i.e., the addition of the answer *A* does not undelete a previously obtained answer) *IDGN_{new}.new_answer* is incremented. Clearly, if *IDGN.nbr_of_answers* is not equal to *IDGN_{new}.nbr_of_answers*, the answers for *T* have changed; also if the two numbers are the same but *IDGN_{new}.new_answer* is set, the answers for *T* have changed. Otherwise, the answers for *T* have not changed, and the subgoals *T* affects are traversed to decrement their *falsecount* fields, which may transitively prevent other subgoals from having to be recomputed (cf. *propagate_validity()* in Fig. 3).

3 Ensuring Transparency through Lazy Recomputation and View Consistency

Perhaps the main drawback of manual incremental tabling is the level of control it requires from a programmer. A programmer can specify that an incremental update is to be done immediately after an assert or retract, but this is inefficient when multiple updates are required. Alternatively, a programmer can specify that an assert or retract simply invalidate affected subgoals, but later must make a call to reevaluate subgoals on the invalid list. In either case, if choice points exist to an incrementally tabled subgoal *S* that is completed, the semantics of an update are undefined (and in fact the program may crash). In addition to these issues, manual incremental tabling may

cause unnecessary work as all affected goals are recomputed even if they are never re-queried. We show how these problems are fixed in transparent incremental tabling.

3.1 Lazy Recomputation

In lazy recomputation assert and retract hooks invalidate tables when a change is made to a dynamic incremental predicate. However, an invalid subgoal S is not re-evaluated until it is called, at which time incremental tabled subgoals upon which S depends are also re-evaluated. The algorithm for lazy recomputation is shown in Fig. 4 within a schematic description of the SLG-WAM’s `tabletry` instruction, which is executed upon calling a tabled subgoal³. Specifically, if S is completed and invalid, lazy recomputation is handled within lines 12-17, using a `reeval_ready` field, which transparent incremental tabling adds to each IDG node frame. If the `reeval_ready` field for S is set to `compute_dependencies_first` the IDG nodes upon which S depends are traversed in a depth-first manner by `traverse_dependent_nodes()` and the traversed subgoals are added to the invalid list (Fig. 4). This predicate, analogous to `traverse_affected_nodes()` of Fig. 3, traverses dependency edges rather than affected edges. Once the invalid list is constructed, its subgoals are recomputed by `recompute_dependent_tables()` which iteratively calls the version of `incremental_reeval()` in Fig. 5. By default the `reeval_ready` field is set to `compute_dependencies_first`, but when `traverse_dependent_nodes()` adds a subgoal S' to the invalid list the `reeval_ready` field for S' is set to `compute_directly` so that the next call to S' will not add it again. Later, the `reeval_ready` field is reset to `compute_dependencies_first` in `incremental_reeval()` after its associated goal is re-evaluated (Fig. 5, line 22); or it is reset when the IDG node frame’s `falsecount` is set to 0 by `propagate_validity()` (this change to Fig. 3 is not shown).

The implementation of line 15 of Fig. 4 uses a general interrupt mechanism whereby a given goal G may dynamically interrupt the current execution environment Env so that G is immediately executed and success and failure continuations of G are (a modification of) Env ⁴. In line 17, the interrupt mechanism intersperses a call to `recompute_dependent_tables()`, to traverse the invalid list and recompute subgoals. When `recompute_dependent_tables()` finishes, its continuation will make a fresh call S , which will see a completed and valid table, and will then simply backtrack through answers for S (starting with line 18 of Fig. 4).

3.2 View Consistency

A fundamental principle of databases is to support view consistency: that is, to ensure that answers to a query Q should be those derivable at the time Q was begun, and should not be affected by any updates. Accordingly, the ISO standard for Prolog (ISO working group JTC1/SC22 1995) specifies that an update v to dynamic code should not affect the behavior of choice points that were created before v . Extending view consistency to incremental tables is critical for understandable system behavior, especially when KRR features such as hypothetical reasoning must be supported. Because XSB’s incremental tabling does not allow updates that affect tables that are still being computed (Section 2), supporting view consistency effectively means ensuring consistency for choice points into completed tables. As such choice points correspond to database cursors, we term them *Open Cursor Choice Points*, (*OCCPs*).

³ XSB’s `tabletry` instruction is substantially more complex as it supports call subsumption, subgoal abstraction, multi-threaded tabling and other features.

⁴ In XSB, as in other Prologs, such interrupts are used to handle unification of attributed variables, signaling among Prolog threads, and other tasks.

```

Instruction tabletry /* SLG-WAM instruction for calling a tabled subgoal S */
  Check whether there is a table for a variant of S and make a table for S if not
  If S is incremental create an IDG node frame
    If S has a nearest tabled ancestor  $S_{anc}$  add IDG edges between S and  $S_{anc}$  if not present
5  If there was not a table for S
    Create a subgoal frame for S
    Create a generator choice point to produce answers via program clause resolution
    Else if S is incomplete /* all answers for S may not yet have been derived */
      Create a consumer choice point to perform answer resolution
10  Else, if S is completed /* all answers for S have been derived */
    Set up a consumer choice point to perform answer resolution
    If S is incremental and invalid
      If  $S.IDG\_node.reeval\_ready == compute\_dependencies\_first$ 
         $invalid\_list = traverse\_dependent\_nodes(S.IDG\_node.invalid)$ 
15      Interrupt to call  $recompute\_dependent\_tables$  with continuation S
      Else /*  $S.IDG\_node.reeval\_ready == compute\_directly$  */
         $incremental\_reeval(S.subgoal\_frame)$ 
    Branch to the instruction of the root of the answer trie for S

IDGN_{dep} upon which IDGN directly depends
    If  $(IDGN_{dep}.reeval\_ready == compute\_dependencies\_first)$ 
      Add  $IDGN_{dep}$  to the global invalid list.
      traverse_dependent_nodes( $IDGN_{dep}$ )

```

Fig. 4. Schematic pseudo-code for lazy recomputation

The approach to view consistency adopted by transparent incremental tabling is summarized in this section, with further details provided in Appendix A. A main goal is to avoid overhead when there are no choice points whose “view” needs to be maintained (including those of non-incremental tables). For this purpose, an *occp_num* field is maintained in the subgoal frame of a completed incremental table T to indicate whether there are OCCPs for T (Appendix A.1). *occp_num* is incremented when the subgoal for T is called; and decremented when the last answer for T has been returned to the call, or when a cut or throw removes the call from the choice point stack. Only if $occp_num > 0$ must the OCCP’s view be preserved. Transparent incremental tabling performs this preservation during *incremental_reeval()* by calling *preserve_occp_views()* (Fig. 5, line 4). While *preserve_occp_views()* is fully described in Appendix A.2, its main actions are as follows. The choice point stack is traversed, and for each OCCP CP_T for T , the answer substitutions that have not yet been resolved by CP_T are determined and then copied from T into the heap as a list (making sure that their heap space is frozen so they are not lost upon backtracking). For each answer substitution that corresponds to an answer whose truth value is *undefined*, the copying includes a special marker *undef*. Next, the structure of CP_T is altered, and its instruction is modified to backtrack through the list on the heap rather than through the table. Once *preserve_occp_views()* has executed, *incremental_reeval()* proceeds as it would otherwise do. Later, when the modified version of CP_T is backtracked into, a new instruction, *preservedViewMember* is called to return the answer substitutions for the preserved view (Appendix A.3) using the correct truth value. When the answers in the list have been exhausted, the heap space used for the list is unfrozen if it is safe to do so.

4 Supporting Well-Founded Negation

A necessary extension for incremental tabling to support KRR applications of the type mentioned in the introduction is to support full well-founded negation. KRR applications make use of the *undefined* truth value to represent conflicts in the defeasibility theory used by a program, as well

as to handle infinite models through a type of answer abstraction called *restraint* (Grosz and Swift 2013), and to support debugging of KRR programs (cf. (Naish 2006)).

As mentioned in the previous section, the IDG maintains information about the dependency and affected relations without representing whether these changes are positive or negative. One advantage of this is that manual incremental tabling is correct for stratified negation — here, meaning well-founded negation with two-valued models. However, to support full well-founded negation, the update process must handle tables in which some atoms are *undefined*. To explain how this is done, we overview those aspects of well-founded negation in SLG resolution (Chen and Warren 1996) that are relevant to the incremental update algorithms.

Essentially, a query evaluation by SLG resolution builds up a partial model of those parts of a program P that are relevant to the query. To make this specific, an SLG evaluation $\mathcal{E}$ is modeled as a sequence of states, called *forests*. Let $\mathcal{F}$ be one such forest in $\mathcal{E}$. $\mathcal{F}$ contains a set of tabled subgoals that have been encountered so far in $\mathcal{E}$. Each such tabled subgoal S in $\mathcal{F}$ is associated with a table T_S containing computed answers for S ; T_S may be marked as *completed* in $\mathcal{F}$ if it has been determined that all necessary resolution has been performed to derive answers for S . To support 3-valued interpretations of $\mathcal{F}$, answers are distinguished as *unconditional* answers representing true derivations, and *conditional* answers representing derivations of atoms with truth value *undefined*. Accordingly, let T_S be a completed table for a subgoal S in $\mathcal{F}$, and let S_G be an atom in the ground instantiation of S . S_G is true if it is in the ground instantiation of some unconditional answer in T_S , and is false if it is *not* in the ground instantiation of *any* answer in T_S (conditional or unconditional).

Formally, for a subgoal S , a conditional answer has the form $S\theta:-DL$ where $S\theta$ is termed the *answer substitution*; and DL , the *delay list*, is a list of literals needed to prove $S\theta$ but whose resolution has been delayed because they do not have a well-founded derivation (based on the current state of the evaluation if T_S is not completed). During the course of an evaluation, if a literal L in a delay list becomes *true* or *false*, the SLG SIMPLIFICATION operation respectively removes L from the delay list or indicates that the conditional answer itself is false.

Example 4.1

The goal $p(X)$ to the program

$p(1) \quad p(2):- \text{not } q(2) \quad p(2):- \text{not } q(3) \quad q(X):- \text{not } p(X)$

has an unconditional answer $p(1)$ along with two conditional answer: $p(2):- \text{not } q(2)$, $p(2):- \text{not } q(3)$. Note that the delay lists for answers to $p(2)$ contain only *undefined* literals upon which $p(2)$ directly depends (e.g., $\text{not } q(2)$, $\text{not } q(3)$), but not indirect dependencies such as $\text{not } p(3)$. $\square$

As mentioned above, XSB represents both tabled subgoals and their answers using tries, a representation that is supported by other Prologs such as YAP (Santos Costa et al. 2012) and Ciao (Hermenegildo et al. 2012). In XSB this representation is extended as follows (Sagonas et al. 2000). If an atom A_{cond} is *undefined*, the leaf node of the answer representing A_{cond} points to an *answer information frame* which in turn points to other answers conditional on A_{cond} as well to a *delay trie* representing all delay lists upon which A_{cond} is conditional. In Example 4.1 the delay trie for $p(2)$ would contain the lists $[\text{not } q(2)]$ and $[\text{not } q(3)]$. Whenever an unconditional answer $S\theta$ is derived in a table for subgoal S , the answer information frame and delay trie for conditional answers to $S\theta$ are deallocated if they exist.

To extend incremental recomputation to correctly handle changes involving conditional answers, several previously unconsidered cases must be addressed for a given answer substitution $S\theta$ in a table S . Each case below considers only those answers in the table S .

- *Informational Weakening 1*, There were previously no answers for $S\theta$; after the update there are one or more conditional answers for $S\theta$.
- *Informational Weakening 2*, There was previously an unconditional answer for $S\theta$; after the update there are one or more conditional answers for $S\theta$.
- *No Informational Change*, There were previously one or more conditional answers for $S\theta$; after the update further conditional answers $S\theta$ were added, or some but not all conditional answers for $S\theta$ were deleted.
- *Informational Strengthening 1*, There were previously one or more conditional answers for $S\theta$; after the update $S\theta$ becomes *true*, with an unconditional answer.
- *Informational Strengthening 2*, There were previously one or more conditional answers for $S\theta$; after the update $S\theta$ becomes *false*, with no answers.

The cases above are grouped by their action on the information ordering of truth values, where both *true* and *false* are stronger than *undefined*. From the perspective of table updates, no action need be taken in the case of *No Informational Change*, as the truth value of $S\theta$ is unchanged. To see this, recall that delay lists contain only direct dependencies. Thus any answer A' that is conditional on $S\theta$ will contain $S\theta$ or *not* $S\theta$ in its delay lists so that changes to the delay list of $S\theta$ need not be propagated. Strengthening and weakening of answers are addressed by the extensions to `incremental_reeval()` shown underlined in Fig. 5.

```

incremental_reeval(IDG node frame IDGN)           /* S is the subgoal to be re-computed */
  If IDGN.falsecount > 0
    Let  $S_T$  be the subgoal frame associated with IDGN (i.e.,  $S_T = IDGN.subgoal\_frame$ )
    if  $S_T.occp\_num > 0$  preserve_occp_views( $S_T$ )    /* Ensure view consistency: see Section 3.2 */
5   For each answer  $S\theta$  in  $S_T.answer\_list$ 
      $S\theta.deleted = true$ 
     If  $S\theta$  is unconditional  $S\theta.unconditional = true$  else  $S\theta.unconditional = false$ 
     Create a new IDG node  $IDGN_{new}$  for  $S_T$ 
      $IDGN_{new}.new\_answer := false$ ;  $IDGN_{new}.falsecount = IDGN_{new}.nbr\_of\_answers = 0$ 
10  Call  $S$  and for each new derived answer  $S\theta:DL$ 
     If  $S\theta.deleted == true$ 
        $S\theta.deleted = false$ ;  $IDGN_{new}.nbr\_of\_answers ++$ 
       If  $S\theta.unconditional == false$  but  $S\theta$  is now unconditional
          $S\theta.unconditional = true$ ; invoke simplification
15  Else /*  $S\theta.deleted$  was false */  $IDGN_{new}.new\_answer = true$ 
     After completion of  $S$ , for each  $S\theta$  in  $S_T.answer\_list$ 
       If  $S\theta.deleted == true$ , remove  $S\theta$  from  $S_T.answer\_list$ 
       If  $S\theta.unconditional = false$  invoke simplification
       Adjust trie choice points and reclaim space for  $S\theta$ 
20  Else if  $S\theta.unconditional == true$ , and  $S\theta$  is now conditional
      $IDGN_{new}.new\_answer = true$ 
      $IDGN.reeval\_ready = compute\_dependencies\_first$ 
     If  $IDGN_{new}.new\_answer == false$  and  $IDGN_{new}.nbr\_of\_answers = IDGN.nbr\_of\_answers$ 
       propagate\_validity(IDGN)

```

Fig. 5. Schematic algorithm for updates in transparent incremental tabling

As shown in Fig. 5 setup for the re-derivation of S now also sets a new *unconditional* field of an answer, representing whether the answer was unconditional at the start of the re-derivation (line 7). In the re-derivation, $IDGN_{new}.nbr_of_answers$ is incremented whenever a new answer substitution $S\theta$ is encountered, whether $S\theta$ is conditional or unconditional (Fig. 5 lines 11-12), so that $IDGN_{new}.nbr_of_answers$ will be updated at most once regardless of how many conditional answers exist for $S\theta$. Thus, there are no changes required for *Informational Weakening 1* as the addition of new conditional and unconditional answer substitutions is handled in the

same manner. Also, if more than one conditional answer is derived for $S\theta$, only the first will increment $IDGN_{new}.nbr_of_answers$, in effect handling the case of *No Informational Change*. A similar check of $S\theta.unconditional$ during re-derivation (lines 13-14) handles *Informational Strengthening 1*, the case where $S\theta$ had only conditional answers, but is now unconditional. This case can actually be handled directly by SLG simplification and does not require propagation through the incremental update system. Once S has been rederived, its answer list is traversed as before (cf. Fig. 3). During this traversal, line 18 handles the case of *Informational Strengthening 2* where $S\theta$ had been conditional but is now false and uses simplification; Lines 20-21 handle *Informational Weakening 2* where $S\theta$ had been *false* but now is *undefined*.

Fig. 5 reflects a bilattice of the information ordering and the truth ordering (where $true > undefined > false$). As discussed, SLG simplification propagates changes of an answer's truth value when it is informationally strengthened. Changes to the truth value of $S\theta$ that reflect a strengthening in the truth ordering can be detected during re-derivation (*Information Weakening 1*, *Information Strengthening 1*). Changes that reflect a weakening in the truth ordering must wait until the re-derivation is complete (*Information Weakening 2*, *Information Strengthening 2*).

5 Abstracting the IDG

The IDG is clearly essential to efficiently update incremental tables, but in certain situations constructing the IDG can cause non-trivial overheads in query time and table space. These overheads can be addressed in many cases by *abstracting* the IDG. When a tabled subgoal S is called, rather than creating an edge between S and its nearest tabled ancestor S' (if any), one could abstract S , S' or both. The semantics and implementation of subgoal abstraction was defined in (Riguzzi and Swift 2013), here we appeal to an intuitive notion of *depth abstraction*: given a subgoal S and integer k , subterms of S with depth $k + 1$ are replaced by unique new variables. For instance, in Fig. 1, abstracting $q(f(1))$ at level 1 gives $q(f(X_1))$; abstracting at level 0 gives $q(X_1)$.

Figure 6 illustrates an important case where abstracting the IDG can be critical to good performance for incremental tabling. In the case of left-linear recursion, if no abstraction is used a new node will be created for each call to *edge/2* as shown on the left side of this figure. If a large number of data elements are in fact reachable, the size of the IDG can be very large. If calls to the *edge/2* predicate make use of depth-0 abstraction, the graph may be much smaller as seen on the right side of Fig. 6. Whether abstracting a IDG in this manner is useful or not is application dependent; however, performance results in the next section illustrate cases where abstraction greatly reduces both query time and space.




Fig. 6. A left-linear program and schematic IDGs: Left without IDG abstraction; Right: with IDG abstraction

Abstracting the *edge/2* predicate has subtle differences from abstracting tabled subgoals. In the first place, the *edge/2* predicate of Fig. 6 is not tabled. Furthermore, the actual *edge/2* subgoal itself should not be abstracted to depth 0 since losing the first argument instantiation would prevent

the use of indexing. Rather, only the IDG’s representation of the subgoal should be abstracted. Fortunately, in XSB the code to intern dynamic goals for the IDG shares code used for tabling, so that extending abstraction to handle dynamic incremental predicates is relatively straightforward. In XSB, abstraction of dynamic code for the IDG can be specified via the declaration:

```
:- dynamic edge/2 as incremental, abstract(0).
```

6 Performance Results and Analysis

The performance of manual incremental tabling in XSB has been analyzed previously, most extensively in (Saha 2006). By and large the behavior of manual incremental tabling features are not affected by the rewriting to support transparent incremental tabling. Accordingly, the performance questions addressed here analyze new features, scalability, and the behavior of incremental tabling for KRR-style computations. A summary of performance results is given in this section, with tables and other details provided in Appendix B.

Left-Linear Recursion Recursion is heavily used in KRR-style programs that make use of features such as Hilog or defeasibility. As a first test, queries of the form $reach(\langle free \rangle, \langle free \rangle)$ were made to a left recursive predicate (Fig. 6) with and without IDG abstraction on the *edge/2* predicate (cf. Appendix B.1). In the benchmarks, *edge/2* consists of ground facts representing randomly generated graphs of 50,000 – 5,000,000 edges. As shown in Fig. B 1 if IDG abstraction is not used, creating the IDG adds a CPU time overhead of roughly 50% and a table space overhead of about 300% compared to non-incremental tabling. By using IDG abstraction at depth 0, the table space overhead becomes approximately 30%, and the time overhead 5-10%. Fig. B 2 shows that for a batch updates (0.02%-2% of EDB), the overhead of re-evaluation is negligible, particularly if abstraction is used.

Non-Stratified Linear Left Recursion Similar tests were made using the predicate *ureach/2* (Fig. B 3), constructed to perform transitive closure, but producing answers with truth value *undefined*. The query $ureach(\langle free \rangle, \langle free \rangle)$ was evaluated on a graph of 500,000 edges. Overhead results for the initial query (Fig. B 4) are similar to those for $reach(\langle free \rangle, \langle free \rangle)$ in terms of time; however the space overhead for incremental tabling is proportionally less (around 10-15% with IDG abstraction) as storing the conditional answers used in this test imposes its own space overhead. Fig. B 5 shows the time to perform various inserts that cause new answers to be added to the table for $ureach(\langle free \rangle, \langle free \rangle)$, and that also change the truth value of some known answers from *undefined* to *true* as discussed in Section 4. The figure shows that updating conditional answers imposes essentially no overhead compared to unconditional answers.

Performance Analysis on a Program with KRR Features The program in Fig. B 6 represents a social network in which certain members of a population are at risk, and other members of the population may influence the behavior of the at-risk members. While the program contains stratified negation, its main computational challenge arises from its heavy use of equality between constants and functional terms – a reasoning capability similar in flavor to some description logics. This use of equality over functional terms quickly leads to non-termination and unsafe negative subgoals during query evaluation. As discussed in Appendix B.2, these behaviors are handled using various tabling mechanisms, so that the ability to incrementally maintain tables for queries to this program requires the ability to update three-valued models that arise from answer abstraction (Grosz and Swift 2013), tabled negation and subgoal abstraction.

As a first benchmark of this program, a small EDB of about 10,000 facts was generated, and $good_influence(\langle bound \rangle, \langle free \rangle)$ was queried for 200 randomly chosen values for its first argument. In this case, incremental tabling caused a time overhead of about 240% and a space

overhead of 280% – although further exploration of abstraction would likely reduce these numbers. Next, updates of substantial sizes were performed on the EDB, and the revaluation time was computed (Fig. B 7). While most of these times are near the level of noise, recomputation of several of the predicates timed out. Analysis of these timeouts showed that they arose because the additional facts caused a large number of new tables to be created when the 200 queries were re-evaluated. Fig. B 8 shows the times to assert or retract, plus time to invalidate affected subgoals via `traverse_affected_nodes()`. Except for updates to *parent_of_edb/2* which directly affect equality, invalidation did not take a significant amount of time

Scalability Analysis on a Program with KRR Features As a next step, the computational burden of the equality relation in the previously mentioned program was reduced by specializing it as discussed in Appendix B.2.1. Tests were performed on EDBs from around 100,000 – 10,000,000 facts. As shown in Fig. B 9, the space and time for these computations scales roughly linearly. For the EDB of about 10,000,000 facts, times were obtained for various large batch updates and for query re-evaluation (Figs. B 10 and B 11). Except for updates to *parent_of_edb/2*, re-evaluation time was very low compared to initial query time (even for initial queries using non incremental tabling), illustrating the promise of incremental tabling for large, reactive systems. These benchmarks also demonstrate the scalability of this implementation, even for very large IDGs. In Fig. B 9, the IDG contained over 750 million edges; after the update sequences mentioned above were applied, it contained more than 1 billion edges.

7 Discussion

This paper has introduced transparent incremental tabling, which improves previous versions of incremental tabling in both semantics and efficiency. The semantics of lazy recomputation (Section 3) together with the preservation of view consistency (Section 3.2 and Appendix A) guarantee that incremental tables will *always* reflect the state of the underlying knowledge base at the time they were queried. This view consistency takes tabled logic programming a step closer to deductive databases, and supports hypothetical reasoning in KRR applications. In addition, the ability to update 3-valued computations (Section 4) is necessary when defeasibility is used over the well-founded semantics, as well as for other features such as answer abstraction. In terms of efficiency, lazy recomputation avoids recomputing invalidated queries until they are requested, and IDG abstraction (Section 5) can significantly reduce the amount of time and space required for queries. The efficiency and scalability of the resulting implementation was summarized in Section 6 and discussed in detail in Appendix B. Appendix C provides further information about how to use transparent incremental tabling in practice.

Although the major semantic issues for incremental tabling have been addressed in this paper, KRR-style computations incur a heavy computational burden, and the benchmark programs do show cases where transparent incremental tabling incurs more cost than is desirable. An important goal is to “guarantee” bounds for transparent incremental tabling when used on representative KRR programs. For instance, for constant bounds b_i , initial query time using incremental tabling should never be more than b_i times that of non-incremental tabling; recomputation time should never be significantly more than initial query time; and the space for the IDG should never be more than b_j times the space of the tables themselves. Such bounds may be obtained through a mixture of program analysis (some of which may itself be incremental cf. (Hermenegildo et al. 2000)) and adaptive incremental tabling algorithms. Even today, incremental tabling is starting to be used to prototype applications in stream deductive databases and event monitoring; continued efficiency improvements should make commercial applications in these areas possible.

References

- CHEN, W. AND WARREN, D. S. 1996. Tabled Evaluation with Delaying for General Logic Programs. *Journal of the ACM* 43, 1, 20–74.
- GROSOFF, B. AND SWIFT, T. 2013. Radial restraint: A semantically clean approach to bounded rationality for logic programs. In *Conference of the American Association for Artificial Intelligence*.
- HERMENEGILDO, M., PUEBLA, G., MARRIOTT, K., AND STUCKEY, P. 2000. Incremental Analysis of Constraint Logic Programs. *ACM TOPLAS* 22, 2 (March), 187–223.
- HERMENEGILDO, M. V., BUENO, F., CARRO, M., LÓPEZ-GARCÍA, P., MERA, E., MORALES, F., AND PUEBLA, G. 2012. An overview of Ciao and its design philosophy. *Theory and Practice of Logic Programming* 12, 1-2, 219–252.
- ISO WORKING GROUP JTC1/SC22. 1995. Prolog international standard ISO-IEC 13211-1. Tech. rep., International Standards Organization.
- LINDHOLM, T. AND O’KEEFE, R. A. 1987. Efficient implementation of a defensible semantics for Prolog. In *International Conference on Logic Programming*. 21–40.
- LLOYD, J. AND TOPOR, R. 1984. Making Prolog more expressive. *Journal of Logic Programming* 1, 3, 225–240.
- NAISH, L. 2006. A three-valued semantics for logic programmers. *Theory and Practice of Logic Programming* 6, 5, 509–538.
- RAMAKRISHNAN, C., RAMAKRISHNAN, I., AND WARREN, D. S. 2007. XcelLog: A deductive spreadsheet system. *Knowledge Engineering Review* 22, 3, 269–279.
- RAMAKRISHNAN, I. V., RAO, P., SAGONAS, K., SWIFT, T., AND WARREN, D. S. 1999. Efficient access mechanisms for tabled logic programs. *Journal of Logic Programming* 38, 1, 31–55.
- REECE, J., URRY, L., CAIN, M., WASSERMAN, S., MINORSKY, P., AND JACKSON, R. 2010. *Campbell Biology*. B. Cummings. 9th Edition.
- RIGUZZI, F. AND SWIFT, T. 2013. Well-definedness and efficient inference for probabilistic logic programming under the distribution semantics. *Theory and Practice of Logic Programming* 13, 2, 279–302.
- SAGONAS, K. AND SWIFT, T. 1998. An abstract machine for tabled execution of fixed-order stratified logic programs. *ACM TOPLAS* 20, 3 (May), 586 – 635.
- SAGONAS, K., SWIFT, T., AND WARREN, D. S. 2000. An abstract machine for efficiently computing queries to well-founded models. *Journal of Logic Programming* 45, 1-3, 1–41.
- SAHA, D. 2006. Incremental evaluation of tabled logic programs. Ph.D. thesis, SUNY Stony Brook.
- SAHA, D. AND RAMAKRISHNAN, C. 2005. Incremental and demand-driven points-to analysis using logic programming. In *Principles and Practice of Declarative Programming*. 117–128.
- SANTOS COSTA, V., DAMAS, L., AND ROCHA, R. 2012. The YAP prolog system. *Theory and Practice of Logic Programming* 12, 1-2, 5–34.
- SWIFT, T. AND WARREN, D. 2012. XSB: Extending the power of Prolog using tabling. *Theory and Practice of Logic Programming* 12, 1-2, 157–187.
- YANG, G., KIFER, M., WAN, H., AND ZHAO, C. 2013. *FLORA-2: User’s Manual Version 0.99.3*. <http://flora.sourceforge.net>.
- ZHOU, N. AND HAVE, C. 2012. Efficient tabling of structured data with enhanced hash-consing. *Theory and Practice of Logic Programming* 12, 4-5, 547–563.

Acknowledgements

The research in this paper was partially funded by Vulcan, Inc. and Coherent Knowledge Systems and FCT Project ERRO PTDC/ EIACCO/121823/2010. The author would like to thank Paulo Moura for latex-related help, Fabrizio Riguzzi for making the University of Ferrara server available for benchmarks, and anonymous reviewers for their careful comments. Finally, the author would like to thank Michael Kifer for finding and reporting many, many bugs in transparent incremental tabling.

Appendix A View Consistency and Table Updates

As discussed in Section 3.2 the approach to maintaining view consistency for transparent incremental tabling has three main parts. (1) a count of the OCCPs for an incremental table T is always maintained. (2) when an update affects an incremental table, the view of an OCCP is preserved by copying its unconsumed answers onto the heap and altering the OCCP to use the copied answers. (3) a new instruction returns answers from the preserved views upon backtracking.

More than other aspects of transparent incremental tabling, the details of view consistency support rely on tabling data structures and algorithms used by XSB, some background for which is presented here.

- *Answer Tries* Steps 1 and 2 use the sequence of choice points set up when backtracking through an answer trie, the default data structure used by XSB to represent answers (Ramakrishnan et al. 1999). Answer tries are constructed to support *substitution factoring*, so that they contain only the information used to bind variables in the associated subgoals, i.e., the answer substitution introduced in Section 4. Each trie node contains an SLG-WAM instruction, so that returning an answer substitution directly corresponds to traversing a path from the root of a trie to a leaf, and backtracking through the trie corresponds to traversing the trie in a fixed depth-first order. A choice point is created whenever traversing a new node that has multiple children and is removed when all children have been traversed (through a *trust-style* instruction).
- *Freeze Registers* Steps 2 and 3 make use of the SLG-WAM’s **HF** (heap freeze) register, which is used to protect terms in the heap from being over-written when tabled computations are repeatedly suspended and resumed.

While these data structures are not unique to XSB, other engines that differ from XSB in their representation of answers or in their implementation of suspension and resumption may implement this approach with suitable modifications.

A.1 Maintaining a Count of OCCPs for a Completed Incremental Table

To maintain a count of OCCPs, a field called *occp_num* is added to subgoal frames. In addition, the first choice point created in backtracking through an answer trie, CP_{first} is modified so that it increases the OCCP number when CP_{first} is created, and decreases the number when CP_{first} is removed. Finally, any routines that remove choice points must also be modified to reset the *occp_num*, including code that removes choice points upon executing a cut, and when executing a throw operation.

A.2 Preserving Views and Altering OCCPs

In order to preserve the views of the current OCCPs for a table T , `incremental_reeval()` of Fig. 5 is modified to check whether the `occp_num` in the subgoal frame S_T of T is non-zero. If so, `preserve_occp_views()` is called using S_T (described at a highly schematic level in Fig. A 1). This routine traverses the choice point stack from top downwards until all OCCPs for T have been located⁵. When a choice point CP is encountered whose failure continuation points to (the instruction field of) a node in the answer trie for T , the process begins of copying the answers that have not yet been consumed by CP . First, an associated choice point CP_{root} must be found. The process of backtracking through an answer trie can create a series of trie choice points of which CP will be the last in the segment due to the order of the choice point stack traversal. However this series of choice points will always form a connected segment in the choice point stack, so that finding the first choice point of the series CP_{root} , is relatively simple. Next, using CP and CP_{root} , the unconsumed portion of the answer trie for T is traversed; each time the traversal encounters a leaf, a pointer to the leaf is added to a list, *Unconsumed*⁶. Next, a *preservedList* is constructed on the heap, by traversing the elements of *Unconsumed*. Each element of the *preservedList* contains a binary term $ret_2(Substitution, Condition)$. *Substitution* represents a given answer substitution consisting of *AnsSubstSize* terms, one for each distinct variable in the associated subgoal. It is represented as a term $ret_{AnsSubstSize}(Args)$ where each argument corresponds to an element of the answer substitution. *Condition* is *null* for unconditional answers, and points to a special answer *undef* whose truth value is *undefined*, and whose use is explained below.

Once the *preservedList* is constructed, the choice points between CP and CP_{root} inclusive are coalesced via `coalesce_choice_points()` into a new choice point, $CP_{coalesced}$. This routine is easiest to illustrate by its results (Fig. A 2). The address of $CP_{coalesced}$ is that of CP , but when $CP_{coalesced}$ is backtracked into, it will restore the engine environment as it would be if backtracking into CP_{root} , and when its choices are exhausted, it will backtrack into the choice point prior to CP_{root} . Of course, $CP_{coalesced}$ also contains a *preservedList* field.

In Fig. A 2 the values of $CP_{coalesced}$ come from CP_{root} rather than from CP , with the exception of fields representing heap values. In the stack-oriented backtracking used by Prolog, the *preservedList* can be protected by setting $CP_{coalesced}$ to the value of the **H** register after the construction of *preservedList*. If there is a possibility that tabling will suspend and resume computations, `preserve_occp_views()` needs to freeze the heap space containing these answers so that the heap cells containing them will not be overwritten. If (**HF** reg $\neq$ bottom of heap), then there is an active tabled computation, and the heap freeze register is set to the value of the **H** register after construction of *preservedList*. The previous value of the **HF** will be reset using $CP_{coalesced}.previous_hfreg$ once backtracking through *preservedList* is done.

⁵ Unlike some other Prologs, XSB has a choice point stack separate from the local stack. The traversal of the choice point stack uses the *previous_top* field of choice points; this field was not part of the SLG-WAM design presented in (Sagonas and Swift 1998), but was added to support various forms of garbage collection.

⁶ As mentioned previously, (e.g., Section 2) an answer list is preserved for incremental tables. While this answer list contains a pointer to each leaf of an answer trie, its ordering does not correspond to the traversal needed to obtain the unconsumed answers of an OCCP.

```

preserve_occp_views(subgoal_frame  $S_T$ )
  Traverse the choice point stack from top until  $S_T.occp\_num$  OCCPs have been located
  For each choice point  $CP$  in the choice point stack
    If  $CP.failure\_continuation$  points into the answer trie for  $S_T$ 
      Determine the root choice point,  $CP_{root}$ , for  $CP$ 
      Construct a list of pointers,  $Unconsumed$ , to leaves of unconsumed answers
       $preservedList = copy\_answer\_substitutions\_to\_heap(Unconsumed, CP_{root}.AnsSubstSize)$ 
       $coalesce\_choice\_points(CP, CP_{root}, preservedList)$ 
  If ( $\mathbf{HF}$  reg  $\neq$  bottom of heap)  $\mathbf{HF}reg = \mathbf{H}reg$ 
   $S_T.occp\_num = 0$ 

copy_answer_substitutions_to_heap(List_of_trie_leaves  $Unconsumed$ , int  $AnsSubstSize$ )
  For each  $leaf\_ptr$  in  $Unconsumed$ 
    Create a list element with the following information
    Let  $Ans_{heap}$  be a skeleton with argument  $ret_{AnsSubstSize}$  and  $AnsSubstSize$  free variables
    Instantiate each argument of  $Ans_{heap}$  with an element of the answer substitution
    If  $leaf\_ptr$  corresponds to a conditional answer
      Create a non-trailed term on the heap  $ret_2(Ans_{heap}, undefined\_ptr)$ 
    Else create a non-trailed term on the heap  $ret_2(Ans_{heap}, null)$ 
  Return the head of the List

```

Fig. A 1. Schematic pseudo-code for preserving views and altering OCCPs

$CP_{coalesced}$	preservedViewMember $ereg_{root}$ $ebreg_{root}$ $hreg$ $trreg_{root}$ $dreg_{root}$ $rsreg_{root}$ $previous_cp_{root}$ $previous_top_{root}$ $answerSubstSize$ $M_{root}+1$ $answerSubst$ $[M_{root}]$: $answerSubst$ $[0]_{root}$ $preservedList$ $previous_hfreg$	/* Failure Continuation */ /* Top environment in stack (E reg)*/ /* Environment of top choice point (EB reg) */ /* Top of heap (H reg) */ /* Top of trail (TR reg)*/ /* SLG-WAM delay register */ /* SLG-WAM root subgoal register */ /* Pointer to previous choice point */ /* Pointer to the previous top of CP stack */ /* Number of Variables in Answer Substitution */ /* Previous SLG-WAM heap freeze register */
------------------	---	--

Fig. A 2. Choice point stack after coalescing

A.3 Backtracking through Preserved Views

Fig. A 3 shows the new SLG-WAM instruction that returns an answer through a preserved OCCP view when a coalesced choice point is backtracked into. The instruction reconstructs the SLG-WAM state at the time of its call (except for the heap register which was adjusted to protect the *preservedList*). Each answer substitution cell of the coalesced choice point is dereferenced to a heap or local stack cell, the dereferenced cell is bound to an element of the answer substitution, and the binding trailed. Afterwards, the *preservedList* field is reset to point to the next list element if one is present; otherwise the **HF** register is set to its before the view was preserved, and the **B** register is set to the previous choice point.

```

Instruction preservedViewMember
  undo_bindings(B register)           /* does not affect answers that were copied to heap */
  Restore SLG-WAM program registers
  Set up pointers to access  $ret_2(Substitution, Condition)$ 
  If ( $Condition \neq null$ ) delay_negatively( $Condition$ )
  For each cell,  $answerSubst[i]$ , of  $B.answerSubst$ 
    Bind argument  $i$  of  $Substitution$  to the dereferenced value of  $answerSubst[i]$ 
    and trail the binding
  If  $B.preservedList$  has been consumed
     $B = B.previous\_cp$ 
    HF  $reg = B.previous\_hfreg$ 
  Else make  $B.preservedList$  point to the next list element

```

Fig. A 3. Schematic pseudo-code for backtracking through preserved views

A.4 Discussion of View Consistency in Transparent incremental tabling

Of course, other approaches to view consistency are possible besides the one just presented. Before the above was implemented, answer tries were extended to include timestamps indicating when a given answer was valid (analogous to that of (Lindholm and O’Keefe 1987) for dynamic Prolog code). However, the time and space overhead of this approach was deemed to be too high. The actual implementation of the heap copying approach presented here uses XSB’s general tabling code as much as possible, so that the cost to traverse tries and copy answers is generally very low.

It should be noted that the approach to view consistency is more closely linked to the data structures of the XSB engine than are other features of transparent incremental tabling, as view consistency interfaces with XSB’s heap and stack freezing mechanisms.

Appendix B Performance Results

In the benchmarks that follow, all times are measured in seconds, and all space is measured in bytes unless otherwise specified ⁷.

B.1 Transparent Incremental Tabling and Linear Left Recursion

Recursion is heavily used in KRR-style programs that make use of features such as Hilog or de-feasibility. As a first test, queries of the form $reach(\langle free \rangle, \langle free \rangle)$ were made to a left recursive predicate (Fig. B 3) with and without IDG abstraction on the $edge/2$ predicate. As discussed in Section 5 and shown in Fig. 6, the IDG created for such a query may differ greatly depending on whether abstraction is used. In the benchmarks, $edge/2$ consists of ground facts representing a randomly generated graph $G(N/M)$ where N is the number of possible nodes in the graph, while M is the number of directed edges. Because of the left recursive form of $reach/2$ together with its query form, the IDG nodes for $edge/2$ are associated with subgoals $edge(\langle free \rangle, \langle free \rangle)$ from clause 1 of $reach/2$, and $edge(\langle ground \rangle, \langle free \rangle)$ where argument 1 is instantiated by different values of Z in clause 2 of $reach/2$. Using the re-evaluation strategies described in previous sections, any update to $edge/2$ will cause a re-evaluation of the subgoal $reach(\langle free \rangle, \langle free \rangle)$ so that (in this program fragment) maintaining nodes of the form $edge(\langle ground \rangle, \langle free \rangle)$ provides no benefit, as their dependencies will be captured by $edge(\langle free \rangle, \langle free \rangle)$.

Nodes	No incr. tabling		Incr. tabling		Incr. tabling + abstraction	
	CPU time	Table space	CPU time	Table space	CPU time.	Table space
100,000	0.12	7,663,728	0.21	21,671,136	0.13	10,273,672
1,000,000	2.19	72,121,240	3.43	211,184,888	2.34	92,746,112
10,000,000	40.9	701,364,952	59.7	2,070,845,368	41.2	902,048,352

Fig. B 1. Overhead for transparent incremental tabling on query evaluation of $reach(\langle free \rangle, \langle free \rangle)$ over randomly generated graphs $G(Nodes/\frac{Nodes}{2})$

Nbr of asserts	Incr. tabling		Incr. tabling + abstraction	
	Time to read/assert/invalid.	Query time	Time to read/assert/invalid.	Re-query time
100	0.004	3.53	0.003	2.29
1,000	0.023	3.67	0.022	2.29
10,000	0.19	4.20	0.17	2.38

Fig. B 2. Updates of $edge/2$ for the query $reach(\langle free \rangle, \langle free \rangle)$ over a randomly generated graph $G(1,000,000/500,000)$

As shown in Fig. B 1 if IDG abstraction is not used, creating the IDG adds a CPU time overhead of roughly 50% and a table space overhead of about 300%. By using IDG abstraction at depth 0, the table space overhead becomes approximately 30%, and the time overhead 5-10%. Regardless of whether abstraction is used, Fig. B 1 demonstrates scalability for 2 orders of magnitude; the time scales log-linearly due to the need to maintain indices. Fig. B 2 shows that for

⁷ Except for those reported in Section B.2.1, the benchmarks below were performed on a MacBook Pro, with a dual core 2.53 Ghz Intel i5 chip and 4 Gbytes of RAM. The benchmarks for Section B.2.1 were performed on a server at the University of Ferrara with 3 Intel dual-core 3.47 GHz CPUs and 188 megabytes of RAM running under Fedora Linux. The default 64-bit, single-threaded SVN repository version of XSB was used for all tests. Benchmark programs can be obtained at www.cs.sunysb.edu/~tsswift/interpreters.html.

a batch updates (0.02%-2% of EDB), the overhead of re-evaluation is negligible, particularly if abstraction is used.

B.1.1 Non-Stratified Linear Left Recursion

Similar tests were made using the predicate *ureach/2* (Fig. B 3). The query *ureach*(*free*),(*free*) was evaluated on the $G(1000000, 500000)$ graph of *edge/2* facts, so that all answers to the query had the truth value *undefined*. Overhead results for the initial query (Fig. B 4) are similar to those for *reach*(*free*),(*free*) in terms of time; however the space overhead for incremental tabling is proportionally less as storing conditional answers requires its own space overhead (Sagonas et al. 2000). Fig. B 5 shows the time to add various numbers of *edge_1* facts, which causes new answers to be added to the table for *ureach*(*free*),(*free*), and also changes the truth value of some known answers from *undefined* to *true* as discussed in Section 4. From Fig. B 5 it can be seen that updating conditional answers imposes essentially no overhead compared to updating unconditional answers.

```
:- table ureach/2 as incremental.
:- dynamic edge/2, edge_1/2 as incremental.
ureach(X,Y):- reach(X,Z),edge(Z,Y).
ureach(X,Y):- edge(X,Y),undefined.
ureach(X,Y):- edge_1(X,Y).
```

Fig. B 3. Benchmark program for non-stratified left linear recursion

Nodes	No incr. tabling		Incr. tabling		Incr. tabling + abstraction	
	CPU time	Table space	CPU time	Table space	CPU time.	Table space
100,000	0.14	21,333,304	0.24	35,540,760	0.15	24,143,168
1,000,000	2.30	208,352,144	3.61	347,416,664	2.42	228,977,672

Fig. B 4. Overhead for transparent incremental tabling on query evaluation of the non-stratified program *ureach*(*free*),(*free*) over randomly generated graphs $G(Nodes/\frac{Nodes}{2})$

Nbr of asserts	Incr. tabling		Incr. tabling + abstr.	
	Time to read/assert/invalid.	Query time	Time to read/assert/invalid.	Re-query time
100	0.005	3.78	0.004	2.591
1,000	0.025	3.83	0.25	2.57
10,000	0.21	3.86	0.22	2.58

Fig. B 5. Updates of *edge_1/2* for the query *ureach*(*free*),(*free*) over a randomly generated graph $G(1000000, 500000)$

B.2 Analysis of Transparent Incremental Tabling on a Program with KRR-style Features

The program in Fig. B 6 represents a social network in which certain members of a population are at risk, and other members of the population may influence the behavior of the at-risk members. Although the program is simplified and idealized in its content, computationally it requires the use of some sophisticated reasoning features. While the program contains stratified negation, its main computational challenge arises from its use of equality, which provides a reasoning capability similar in flavor to some description logics. The predicate *equals/2* allows terms using the

function symbol *parent_of/1* (formed from the EDB predicate *parent_of_edb/2*) to be considered as equal to constants representing individuals.

```

good_influence(P1,P2):- influences(P1,P2),
    sk_not(high_risk(P1)),sk_not(possible_risk(P1)),
    (high_risk(P2) ; possible_risk(P2)).

:- table high_risk_association/2 as incremental.
high_risk_association(Per1,Per2):- high_risk_contact(Per1,Per2),has_disease(Per2).
high_risk_association(Per1,Per2):- high_risk_association(Per1,Per3),high_risk_contact(Per3,Per2).

high_risk_contact(Per1,Per2):- may_share_needle(Per1,Per2).
high_risk_contact(Per1,Per2):- may_have_unprotected_sex(Per1,Per2).

:- table high_risk/1 as incremental.
high_risk(Per):- high_risk_association(Per,_),!.

:- table possible_risk_association/2 as incremental, answer_abstract(3).
possible_risk_association(Per1,Per2):- might_be_sexual_partner(Per1,Per2),
    high_risk_contact(Per2,_).
possible_risk_association(Per1,Per2):- possible_risk_association(Per1,Per3),
    might_be_sexual_partner(Per3,Per2).

:- table possible_risk/1 as incremental.
possible_risk(Per):- possible_risk_association(Per,_),!.

influences(Per1,Per2):- loves(Per2,Per1).
influences(Per1,Per2):- works_for(Per2,Per1).
influences(Per1,Per2):- attends_church(Per2,Church),pastor(Church,Per1).
influences(Per1,Per2):- lives_at(Per1,Loc),lives_at(Per2,Loc).

may_share_needle(Per1,Per2):- obtained_needle(Per1,Needle,_Loc1), returned_needle(Per2,Needle,_Loc2),Per1 Per2.
may_share_needle(Per1,Per2):- share_needle_report(Per1,Per2,_Per3).

might_be_sexual_partner(Per1,Per2):- loves(Per1,Per2),sk_not(related(Per1,Per2)).
might_be_sexual_partner(Per1,Per2):- sexual_partner_report(Per1,Per2,_Per3).

:- table related/2 as incremental.
related(Per1,Per2):- equals(Per1,parent_of(Per2)).
related(Per1,Per2):- equals(Per1,parent_of(parent_of(Per2))).

:- table loves/2 as incremental.
loves(X,Y):- loves(Y,X).
loves(X,Y):- friend(X,Y).
loves(X,Y):- equals(parent_of(X),Y).
loves(X,Y):- grandparent_of(X,Y).

:- table equals/2 as incremental, subgoal_abstract(3).
equals(X,Y):- equals(Y,X).
equals(parent_of(X),parent_of(X)).
equals(parent_of(X),Y):- parent_of_edb(X,Y).
equals(parent_of(parent_of(X)),Y):- parent_of(X,Z),equals(parent_of(Z),Y).

father_of(X,Y):- equals(parent_of(X),Y),male(Y).
mother_of(X,Y):- equals(parent_of(X),Y),female(Y).
grandparent_of(X,Y):- equals(parent_of(parent_of(X)),Y).

:- dynamic friend/2, returned_needle/3, obtained_needle/3, share_needle_report/3, sexual_partner_report/3 as incremental.
:- dynamic has_disease/1, works_for/2, may_have_unprotected_sex/2, pastor/2, parent_of_edb/2, lives_at/2,attends_church/2
    as incremental,abstract(0).

```

Fig. B 6. A social network example showing KRR features

The EDB for this program consists of 12 different dynamic predicates as seen at the bottom of the program ⁸. The use of the *parent_of/1* function within *equals/2* quickly leads to non-termination and unsafe negative subgoals during query evaluation. Unsafe negative subgoals are soundly addressed by XSB’s *sk_not/1* which skolemizes non-ground variables in an atomic subgoal for the purpose of calling a negative subgoal. Non-termination is addressed in two ways. The use of subgoal abstraction in *equals/2* ensures that there will be only a finite number of tabled queries to this predicate, and in general ensures termination for programs with finite models (Riguzzi and Swift 2013). However, the predicate *possible_risk_association/2*, produces an infinite number of answers for the benchmark data set. The use of answer abstraction (or restraint) for this predicate, ensures sound (but not complete) terminating query evaluation (Groszof and Swift 2013) ⁹.

Thus, the ability to incrementally maintain tables for queries to this program requires the ability to update three-valued models that arise from answer abstraction, and to combine with tabled negation and subgoal abstraction. As a first benchmark test, a small EDB of about 10,000 facts about a population of 10,000 persons was generated, and *good_influence/2* (*bound*, *free*) was queried for 200 randomly chosen values for its first argument. If no incremental tabling was used, the combined CPU time for these queries averaged 1.14 seconds and table space was about 233 megabytes — as discussed further below the relatively large cost for this query was almost entirely due to the use of equality. When transparent incremental tabling was used with no abstraction, the cost rose to 3.02 seconds, and 865 megabytes. By applying IDG abstraction the initial query time dropped to 2.73 seconds and 655 megabytes. The purpose of this sets of declarations was only to test the overhead of transparent incremental tabling for queries and updates: they should not necessarily be considered to be “optimal” for these tests

Fig. B 7 shows times to re-evaluate the queries to *good_influence/2* mentioned above after inserting *N* randomly generated facts for a given predicate (the “Asserts” column); and then after retracting these inserted facts (the “Retracts” column). Most of the times in Fig. B 7 are near the level of noise, however recomputation of several of the predicates timed out ¹⁰. Analysis of these timeouts showed that they arose because the additional facts caused a large number of new (sub-)tables to be created for the 200 queries. Usually, this only occurred after 12,500 facts were added, but for *parent_of_edb/2* which strongly affects goals to *equals/2*, the addition of 500 facts led to a timeout, while the addition of 100 facts led to a 5.57 second recomputation time. Although the program is not wholly monotonic, it is largely so, and computations after retractions were always fast. Fig. B 8 shows the times to assert or retract plus the time taken to invalidate affected subgoals via *traverse_affected_nodes()*. Except for updates to *parent_of_edb/2* invalidation did not take a significant amount of time

⁸ The social network programs and supporting data can be found at <http://www.cs.sunysb.edu/~tsswift>.

⁹ Briefly, if an answer *A* with depth greater than a given bound, *A* is rewritten so that terms with depth equal to the bound are replaced by new variables; then the answer *A* is assigned the truth value *undefined*

¹⁰ Timeouts, denoted Tout in Fig. B 7, were triggered after one minute. The short timeout period was to avoid excessive memory consumption on the laptop benchmarking machine. Retracts of bulk inserts could not be measured, and are designated as n/a. As the population size was 10,000, 12,500 distinct facts could not be generated for the unary EDB predicate *has_disease/1*.

Predicate	Asserts				Retracts	
	100	500	2500	12500	2500	12500
friend/2	0.08	0.37	2.36	Tout	0.02	n/a
returned_needle/3	0.01	0.01	0.01	0.01	0.01	0.01
obtained_needle/3	0.01	0.01	0.01	0.02	0.01	0.01
share_needle_report/3	0.03	0.03	0.13	0.55	0.01	0.01
sexual_partner_report/3	0.01	0.02	0.12	Tout	0.01	n/a
has_disease/1	0.01	0.01	0.01	n/a	0.01	n/a
works_for/2	0.01	0.04	0.42	1.76	0.01	0.01
may_have_unprotected_sex/2	0.03	0.08	0.12	0.56	0.02	0.02
pastor/2	0.01	0.01	0.01	0.01	0.01	0.01
parent_of_edb/2	5.57	Tout	Tout	Tout	n/q	n/a
lives_at/2	0.01	0.01	0.07	2.11	0.01	0.01
attends_church/2	0.01	0.01	0.01	0.01	0.01	0.01

Fig. B 7. CPU times to re-evaluate *good_influence/2* for 200 first-argument bindings after batch updates. The program uses non-specialized equality, and the EDB size is $\mathcal{O}(10^4)$. The top group of predicates use depth-0 IDG abstraction; the bottom group has no IDG abstraction.

Predicate	Asserts				Retracts	
	100	500	2500	12500	2500	12500
friend/2	0.01	0.01	0.02		0.03	
returned_needle/3	0.01	0.01	0.03	0.14	0.03	0.14
obtained_needle/3	0.01	0.01	0.03	0.15	0.03	0.20
share_needle_report/3	0.01	0.01	0.02	0.14	0.03	0.17
sexual_partner_report/3	0.01	0.01	0.03		0.03	
has_disease/1	0.01	0.01	0.02	n/a	0.02	n/a
works_for/2	0.01	0.01	0.02	0.10	0.04	0.16
may_have_unprotected_sex/2	0.03	0.08	0.11	0.52	0.04	0.16
pastor/2	0.01	0.01	0.02	0.11	0.03	0.15
parent_of_edb/2	27.8	Tout	Tout	Tout	37.2	Tout
lives_at/2	0.01	0.01	0.02	0.11	0.03	0.17
attends_church/2	0.01	0.01	0.02	0.11	0.03	0.16

Fig. B 8. CPU times to apply updates and to invalidate subgoals created by queries to *good_influence/2* for 200 first-argument bindings. The program uses non-specialized equality, and the EDB size is $\mathcal{O}(10^4)$. The top group of predicates use depth-0 IDG abstraction; the bottom group has no IDG abstraction.

B.2.1 Scalability Analysis on a Program with KRR Features

As a next step, the equality relation in the previously mentioned program of Fig. B 6 was specialized so that it had the form:

```
:- table equals/2 as incremental, subgoal_abstract(3).
equals(X,Y):- atomic(X), Y = parent_of( _),equals(Y,X).
equals(parent_of(X),parent_of(X)).
equals(parent_of(X),Y):- parent_of_edb(X,Y).
equals(parent_of(parent_of(X)),Y):- parent_of_edb(X,Z),equals(parent_of(Z),Y1),Y1 = Y.
```

In this form, the first clause of *equals/2* is changed so that symmetry is applied only if the first argument corresponds to a nominal individual (constant), and the second argument has a functional form. The fourth clause is changed so that subgoals of the form *equals*($\langle bound \rangle$, $\langle bound \rangle$) are not called by this clause, but instead subgoals of the form *equals*($\langle bound \rangle$, $\langle free \rangle$) are called. These changes, which do not affect the semantics of the program, significantly reduce the time and space required for query evaluation, although goals to *equals/2* are still computationally expensive to update.

With this change, a series of 200 queries as described above were tested on EDBs ranging from around 100,000–10,000,000 facts. As shown in Fig. B 9, the space and time for these computations scales roughly linearly. For the EDB of about 10,000,000 facts, various batch updates were timed along with time to re-evaluate queries (Figs. B 10 and B 11). Specifically for $N = 2500, 12500, 62500$ and 312500 , N asserts of each EDB predicate were performed and timed; and then the N asserted facts were retracted and timed. Except for updates to *parent_of_edb/2*, re-evaluation time was low compared to initial query time (even compared to the initial query time for non-incremental tabling). These benchmarks illustrate the scalability of this implementation of transparent incremental tabling even for very large IDGs. In Figs. B 10 and B 11, the IDG contained over 750 million edges; after the update sequences mentioned above were applied, it contained more than 1 billion edges.

EDB Size	Query Time	Table Space	IDG Nodes	IDG Edges	Non-incr Query Time
$\mathcal{O}(10^5)$	3.9	0.51 Gbytes	22,374	7,362,284	1.7
$\mathcal{O}(10^6)$	62.1	5.33 Gbytes	67,106	78,612,966	24.5
$\mathcal{O}(10^7)$	679.8	51.56 Gbytes	505,972	753,798,584	391.9

Fig. B 9. CPU times to initially evaluate *good_influence/2* for 200 first-argument bindings for EDBs of various sizes. The program uses specialized equality.

Predicate	Asserts				Retracts			
	2500	12500	62500	312500	2500	12500	62500	312500
friend/2	3.11	3.16	2.63	3.51	3.11	3.16	2.58	2.91
returned_needle/3	3.11	6.59	2.57	2.96	3.11	3.21	2.57	2.87
obtained_needle/3	3.11	3.16	2.59	2.65	3.11	3.16	2.52	2.52
share_needle_report/3	3.12	3.16	2.52	2.54	3.11	3.16	2.52	2.54
sexual_partner_report/3	3.12	3.16	2.52	2.54	3.11	3.16	2.52	2.55
has_disease/1	3.46	3.51	2.81	2.80	3.46	3.50	2.80	2.81
works_for/2	3.14	3.25	3.34	4.81	3.11	3.16	2.52	2.52
may_have_unprotected_sex/2	4.34	4.37	3.51	3.51	4.33	4.37	3.51	3.51
pastor/2	3.12	3.16	3.34	2.51	3.11	3.16	2.51	2.52
lives_at/2	3.12	3.16	2.52	2.58	3.11	3.16	2.52	2.52
attends_church/2	3.12	3.16	2.52	2.52	3.16	3.16	2.52	2.52

Fig. B 10. CPU times to re-evaluate *good_influence/2* for 200 first-argument bindings after batch updates. The program uses specialized equality, and the EDB size is $\mathcal{O}(10^7)$. The top group of predicates use depth-0 IDG abstraction; the bottom group has no IDG abstraction.

Predicate	Asserts				Retracts			
	2500	12500	62500	312500	2500	12500	62500	312500
friend/2	0.12	0.60	3.01	15.9	0.13	0.67	3.43	18.1
returned_needle/3	0.12	0.60	3.01	16.0	0.13	0.69	3.51	18.4
obtained_needle/3	0.15	0.74	3.74	19.5	0.17	0.83	4.21	22.0
share_needle_report/3	0.12	0.61	2.99	15.9	0.11	0.01	2.98	15.8
sexual_partner_report/3	0.12	0.61	2.99	16.0	0.11	0.59	2.98	15.9
has_disease/1	0.07	0.33	1.67	8.6	0.08	0.39	1.87	10.5
works_for/2	0.12	0.57	0.42	16.1	0.0	0.65	3.34	18.2
may_have_unprotected_sex/2	0.34	1.68	8.45	43.3	0.13	1.75	8.87	45.3
pastor/2	0.07	0.33	1.71	19.5	0.08	0.39	2.04	10.9
parent_of_edb/2	380.9	Tout	Tout	Tout	222.6	Tout	Tout	Tout
lives_at/2	0.11	0.56	2.82	14.7	0.14	0.68	3.45	18.0
attends_church/2	0.07	0.34	1.71	8.9	0.08	0.43	2.15	11.3

Fig. B 11. CPU times to apply updates and to invalidate subgoals created by queries to *good_influence/2* for 200 first-argument bindings. The program uses specialized equality, and the EDB size is $\mathcal{O}(10^7)$. The top group of predicates use depth-0 IDG abstraction; the bottom group has no IDG abstraction.

Appendix C A Note on Usability

The XSB manual contains information on how transparent incremental tabling may be used in practice; however to make this paper self-contained, we provide an outline of some usability and system aspects.

XSB has a variety of tabling mechanisms that are used for different purposes. As seen from Fig. B 6, transparent incremental tabling works properly with subgoal abstraction and with answer abstraction; as discussed in Section 4, transparent incremental tabling works properly with well-founded negation regardless of the tabled negation operator: for instance with *sk_not/1* in Fig. B 6, or with other XSB operators such as *not/1*. It also works properly with tabled attributed variables (supporting tabled constraints). A variety of dynamic code may be used as a basis for transparent incremental tabling including not only regular facts and rules, but also facts that are interned as XSB tries. Incremental tables, of whatever form, may be used alongside non-incremental tables, although special declarations must be made if an incremental table depends on a non-incremental table.

Within the current version of XSB, transparent incremental tabling does not yet work properly with call subsumption, answer subsumption, hash-consed tables, or multi-threaded tables; also, predicates that are tabled as incremental must use static code rather than dynamic code. Attempts to declare a predicate using an unsupported mixture of tabling features causes a compile-time permission error.

There are situations where it is convenient or necessary to abolish an incremental table rather than updating it. An example of this occurs when an exception is thrown. If an exception is thrown over a choice point to a completed table no action need be taken; however if an exception is thrown over a choice point to an incomplete tabled subgoal (including one that is being recomputed), XSB abolishes the table as its computation has become compromised. In transparent incremental tabling, abolishing an incremental table is not problematic. If a table *T* is to be abolished, tables that depend on *T* must be invalidated before actually abolishing *T* itself. When a call is made to a subgoal with an invalidated affected node, portions of the IDG that were removed through abolishing will be reconstructed during the calls made by *incremental_reeval()*, due to the actions of lazy recomputation.