

# *Description, Implementation, and Evaluation of a Generic Design for Tabled CLP*

Joaquín Arias<sup>1,2</sup>

Manuel Carro<sup>1,2</sup>

joaquin.arias@imdea.org, manuel.carro@{imdea.org,upm.es}

<sup>1</sup>IMDEA Software Institute, <sup>2</sup>Universidad Politécnica de Madrid

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

## Appendix A CLP Trees and TCLP Forests

Prolog and CLP follow a depth-first search strategy with chronological backtracking. The computation rule selects constraints and literals from the resolvent from left to right. Literals are resolved against the clauses of the program, selected from top to bottom. When a literal unifies with a clause head, it is substituted by the body of the clause after applying the unifier obtained from the literal-head unification. If a derivation branch fails because there are no more matching clauses or the constraint store is inconsistent, the evaluation backtracks to the youngest literal that has a candidate matching clause. Depth-first search is incomplete and in general not all answers can be computed. Moreover, there are programs with finite derivations for which logically equivalent programs produce infinite derivations. The use of TCLP can work around this issue in many cases.

We will show the CLP trees and the TCLP forests for the query  $?- D \#< 150, \text{dist}(a, Y, D)$  for two logically equivalent versions of the `dist/3` program: with left recursion (Fig. 1, right) and with right recursion (Fig. 2, right). We use the graph in Fig. A 1, where the length of one of the edges is defined with constraints.




Fig. A 1: Graph definition.  $(25,35)$  is the open interval from 25 to 35.

Fig. A 2 and Fig. A 3 (top) are the CLP trees of the right- / left-recursive programs respectively. Fig. A 3 (bottom) and Fig. A 4 are the TCLP forest of the left- / right-recursive programs respectively. In these figures, the nodes of the trees represent the states (Def. 2) of the computation. A state is a tuple  $\langle R, c \rangle$ , where  $R$  is a sequence of goals,  $[g_1, g_2, \dots, g_n]$  and  $c$  is a conjunction of constraints. The numbers attached to each state indicates the order in which they are created.

On the one hand, Fig. A 2 shows a finite CLP tree which finds all the answers and Fig. A 3 (top) shows an infinite CLP tree caused by the left recursion. On the other hand, Fig. A

3 (bottom) and Fig. A 4 show that the TCLP forest for both programs are finite and all the answers to the query are found, since the use of tabling makes it terminate with left recursion as well.

### A.1 CLP Tree of *dist/3* with Right Recursion

Fig. A 2 shows the CLP tree of the query using the version of *dist/3* with right recursion (Fig. 2, right). We see that the evaluation of the recursive clause generates similar states (s1, s4, s7 and s10), but in each iteration the domain of the constrained variable  $D2_i$  is reduced. As a consequence, the constraint store in state s13 is inconsistent and the evaluation of this derivation fails. The pending branches are evaluated upon backtracking. We explain now how we obtain some of the states; the rest are obtained similarly, so we will skip them:

- s1** the initial state is the representation of the query.
- s2i/ii** are obtained by resolving the literal  $\text{dist}(a, Y, D)$  against the two clauses of the program. The constraints  $Y_1 = Y \wedge D_1 = D$  are added to the constraint store.
- s3** is obtained from the leftmost state s2i by adding the constraints of the resolvent  $[D1_1\#>0, D2_1\#>0, D1\#=D1_1+D2_1]$  to the constraint store.
- s4** is obtained by resolving the literal  $\text{edge}(a, Z_1, D1_1)$ . The constraint  $Z_1 = b \wedge D1_1 = 50$  reduces the domain<sup>1</sup> of  $D2_1$  to  $D2_1 > 0 \wedge D2_1 < 100$ .
- s7** is obtained by resolving the literal  $\text{edge}(b, Z_2, D1_2)$ . The constraint  $Z_2 = a \wedge D1_2 > 25 \wedge D1_2 < 35$  reduces the domain of  $D2_2$  to  $D2_2 > 0 \wedge D2_2 < 75$ .
- s10** is obtained by resolving the literal  $\text{edge}(a, Z_3, D1_3)$ . The constraint  $Z_3 = b \wedge D1_3 = 50$  reduces the domain of  $D2_3$  to  $D2_3 > 0 \wedge D2_3 < 25$ .
- s13** is obtained by resolving the literal  $\text{edge}(b, Z_4, D1_4)$ . The constraint  $Z_4 = a \wedge D1_4 > 25 \wedge D1_4 < 35$  is inconsistent with the current constraint store,  $D < 150 \wedge D > 125 + D1_4 + D2_4 \wedge D2_4 > 0 \wedge \dots$ . Its child is a fail node.
- s14** is obtained, upon backtracking to the state s11b by resolving the literal  $\text{edge}(b, Y, D2_3)$ . However, it is also a failed derivation because the resulting constraint store is inconsistent.
- s15** is a final state of a successful derivation, obtained upon backtracking to the state s8b by resolving the literal  $\text{edge}(a, Y, D2_2)$ . The constraint  $Y = a \wedge D2_3 > 25 \wedge D2_3 < 35$  is consistent with the constraint store.
- a1** is the first answer  $Y = a \wedge D > 125 \wedge D < 315$ , projected onto the variables of the query ( $\text{vars}(Q) = \{Y, D\}$ ).
- s16** is a final state obtained upon backtracking to the state s5b.
- a2** is the second answer,  $Y = a \wedge D > 75 \wedge D < 85$ .
- s17** is a final state obtained upon backtracking to the state s2ii.
- a3** is the third and last answer,  $Y = b \wedge D = 50$

<sup>1</sup> We are considering a linear constraint solver over the rational numbers that from  $D < 150 \wedge D = D1_1 + D2_1 \wedge D1_1 = 50$  it infers that  $D2_1 < 100$

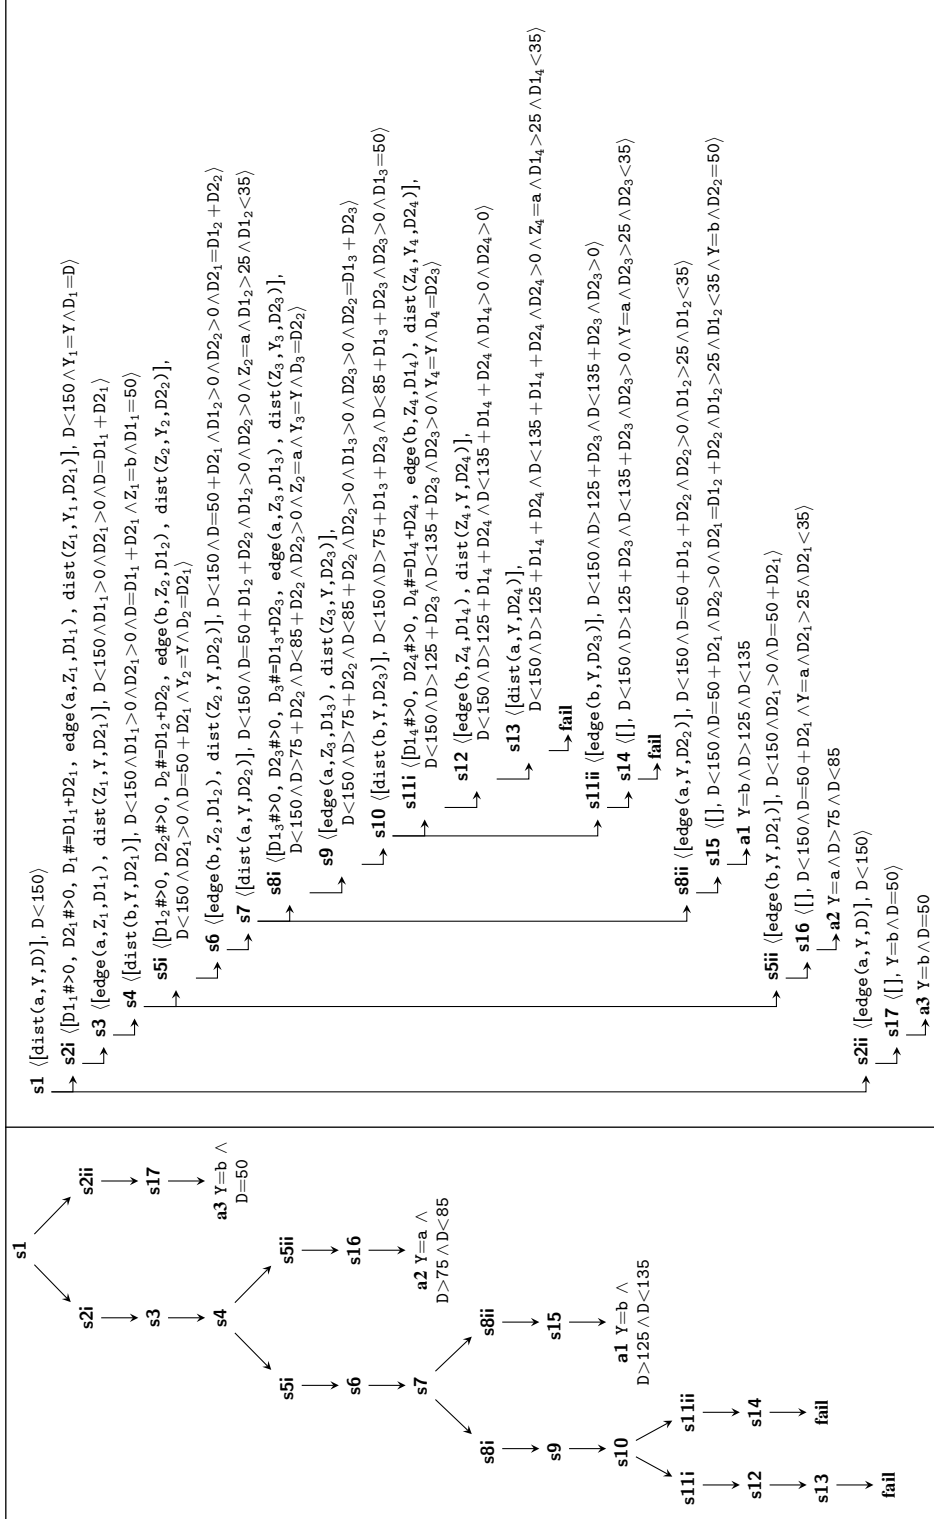


Fig. A2: CLP tree of the query  $?- D \#< 150, \text{dist}(a, Y, D)$  for  $\text{dist}/3$  with right recursion

### A.2 CLP Tree of $\text{dist}/3$ with Left Recursion

Fig. A 3 (top) shows the CLP tree of the query to  $\text{dist}/3$  with left recursion (Fig. 1, right). We see that the recursive clause also generates similar states ( $s_1, s_3, s_5, \dots$ ) but in this example the domain of the constrained variable  $D_1$  remains unchanged, and the evaluation therefore enters a loop. As before, we only explain how we obtain some of the states:

- s3** is obtained from the leftmost state  $s_{2i}$ . The domain of  $D_1$  is  $D_1 > 0 \wedge D_1 < 150$ .
- s5** is obtained from the leftmost node  $s_{4i}$ . The domain of  $D_2$  is  $D_2 > 0 \wedge D_2 < 150$ .
- $\dots$  the evaluation enters a loop.

Although the program that generates this CLP tree is logically equivalent to the previous one, this tree is infinite and no answers are found.

### A.3 TCLP Forest of $\text{dist}/3$ with Left Recursion

Fig. A 3 (bottom) shows the TCLP forest for the query we have been using with the  $\text{dist}/3$  program written using left recursion (Fig. 1, right), where the set of tabled predicates is  $\text{Tab}_P = \{\text{dist}/3\}$ . The main point is that at state  $s_3$  the tabling engine detects that the evaluation of  $\langle \text{dist}(a, Z_1, D_1), D_1 > 0 \wedge D_1 < 150 \rangle$  entails the generator  $\langle \text{dist}(a, V_0, V_1), V_1 < 150 \rangle$  and therefore it suspends the execution and waits until another generator feeds the suspended goal with answers. The evaluation of the state  $s_{2ii}$  generates the first answer  $a_1$  upon backtracking. Then, the tabling engine resumes the consumer with  $a_1$  and generates  $a_2$  which is used to generate  $a_3$ . Finally, the evaluation fails after consuming  $a_3$  and, since all the clauses have been evaluated and there are no more consumers to be resumed or answers to be consumed, the generator is marked as complete and all the answers are returned. We explain below how some of the states are obtained. The rest of the states are obtained similarly, so we skip them for brevity:

- s0** We omit the representation of the TCLP tree for the query  $\tau_P(\text{dist}(a, Y, D), D < 150)$  and its answer resolution.
- s1** the initial state of the TCLP tree  $\tau_P(\text{dist}(a, V_0, V_1), V_1 < 150)$  is the renamed generator. (Def. 3).
- s2i/ii** are obtained by resolving the literal  $\text{dist}(a, V_0, V_1)$  against the two clauses of the program.
- s3** is obtained from the leftmost state  $s_{2i}$  by adding the constraints to the constraint store as in the CLP tree.
- Ans(s1)** the tabled literal  $\text{dist}(a, Z_1, D_1)$  has to be resolved by answer resolution (Def. 3) using the answer from the current TCLP tree  $\tau_P(\text{dist}(a, V_0, V_1), V_1 < 150)$  because, after renaming, the projection of the current constraint store onto the variables of the literal entails the projected constraint store of the generator:  $V_1 > 0 \wedge V_1 < 150 \sqsubseteq V_1 < 150$ . Since the current TCLP tree is under construction and depends on itself, this branch derivation is suspended.
- s4** is a final state of a successful derivation. It is obtained, upon backtracking to the state  $s_{2ii}$ , by resolving with edge  $(a, V_0, V_1)$ . The equations  $V_0 = b \wedge V_1 = 50$  are consistent with the constraint store.



- a1** is the first answer,  $V0 = b \wedge V1 = 50$  (Def. 4). Since is the first one, it is also the more general one.
- s5** is obtained from the state **s3** (because there are no more branches) by answer resolution consuming **a1** (Def. 4).
- s6** is a final state obtained by resolving the literal  $\text{edge}(b, V0, D2_1)$ .
- a2** is the second answer,  $V0 = a \wedge V1 > 75 \wedge V1 < 85$ . It is neither more particular nor more general than **a1**.
- s7** is obtained from the state **s3** by consuming **a2**.
- s8** is a final state.
- a3** is the third answer,  $V0 = b \wedge V1 > 125 \wedge V1 < 135$ . It is neither more particular nor more general than **a1** or **a2**.
- s9** is obtained from the state **s3** by consuming, **a3**.
- s10** is a failed derivation because the resulting constraint store is inconsistent,  $V1 < 150 \wedge \dots \wedge V1 > 125 + D2_1 \wedge D2_1 > 25$ . Its child is a fail node.

Note that the CLP execution entered a loop when resolving the state **s3**. Under TCLP, answer resolution avoids looping and the resulting TCLP forest is finite and complete (i.e., the leaves of the trees are either fail nodes or answers).

#### A.4 TCLP Forest of *dist/3* with Right Recursion

Fig. A 4 shows the TCLP forest corresponding to querying the right recursive *dist/3* program (Fig. 2, right). This example is useful to show how the algorithm works with mutually dependent generators<sup>2</sup> and to see why not all the answers from a generator may be directly used by its consumers.

Unlike the left-recursive version, which shows only one TCLP tree (Fig. A 3, bottom), Fig. A 4 has two TCLP trees (one for each generator). That is because the left recursive version only sought paths from the node **a**, but the right recursive version creates a new TCLP tree at the state **s4** to collect the paths from the node **b**, since  $\text{edge}(a, b)$  had been previously evaluated at state **s3**. As before we only explain how we obtain some of the states:

- s1** the TCLP tree  $\tau_P(\text{dist}(a, V0, V1), V1 < 150)$  is created.
- s4** is obtained by resolving the literal  $\text{edge}(a, Z_1, D1_1)$ .
- Ans(s5)** the tabled literal  $\text{dist}(b, V0, D2_1)$  is a new generator and a new TCLP tree  $\tau_P(\text{dist}(b, V2, V3), V3 > 0 \wedge V3 < 100)$  is created (Def. 3).
- s5** is the root node of the new TCLP tree.
- s6i/ii** are obtained by resolving the literal  $\text{dist}(b, V2, V3)$  against the clauses of the program.
- s8** is obtained by resolving the literal  $\text{edge}(b, Z_1, D1_1)$ .  
In the state **s8**, the call  $(\text{dist}(a, V2, D2_1), D2_1 > 0 \wedge D2_1 < 75)$  is suspended because it entails the former generator  $(\text{dist}(a, V0_1, V1_1), V1_1 < 150)$ .

<sup>2</sup> I.e., generators which consume answers from each other.

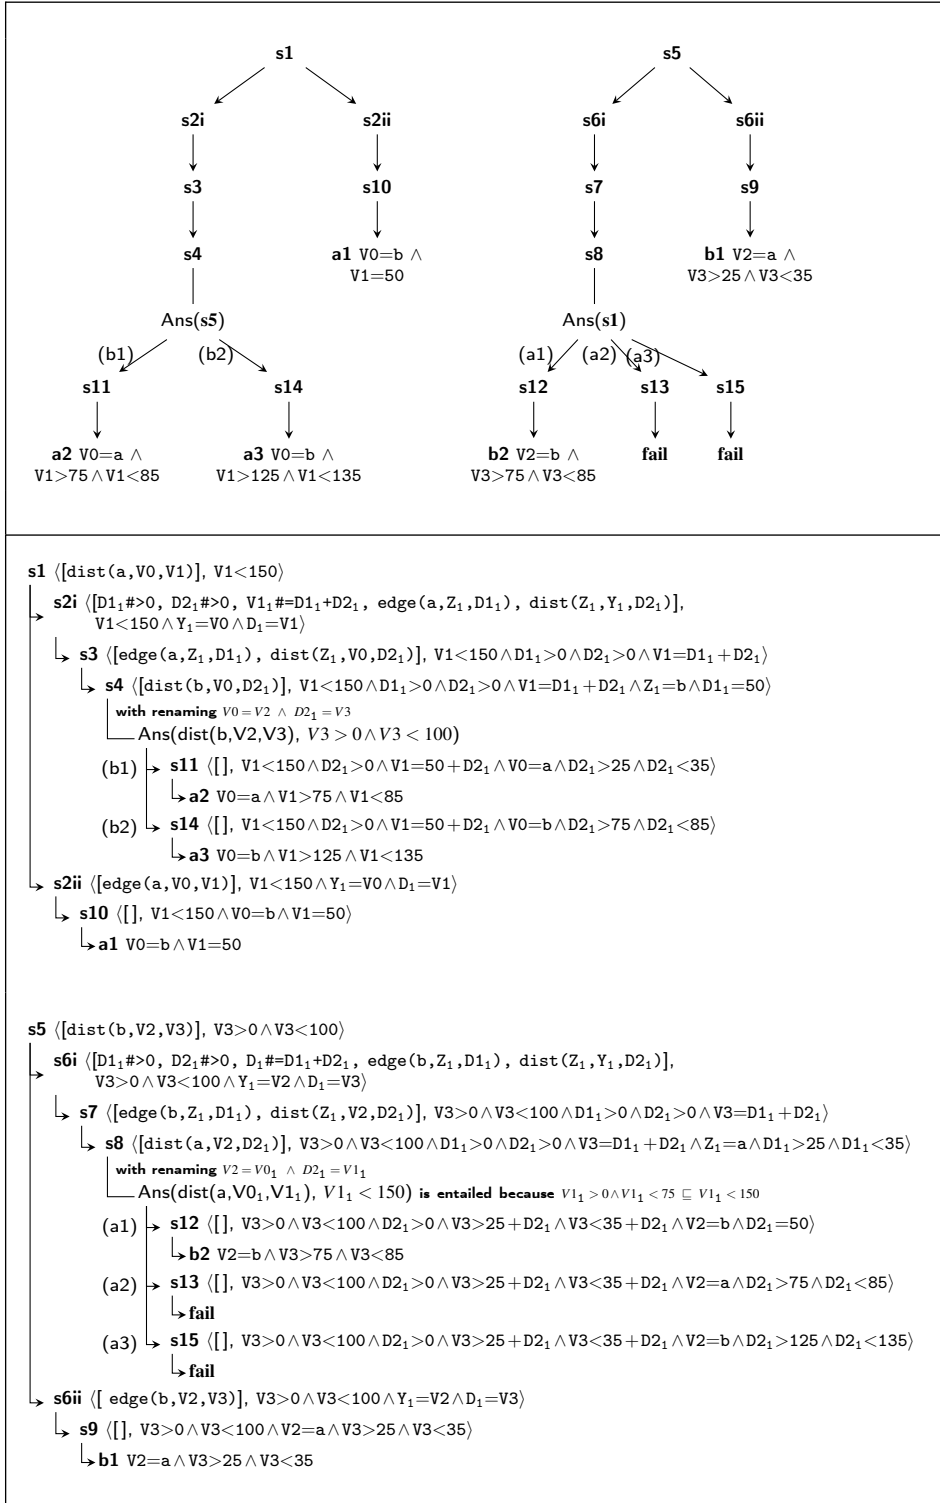


Fig. A 4: TCLP-forest of the query  $?- D \# < 150, \text{dist}(a, Y, D)$  for  $\text{dist}/3$  with right recursion

**Ans(s1)** the tabled literal  $\text{dist}(a, V2, D2_1)$  is resolved with answer resolution (Def. 3) using the answers from the previous TCLP tree  $\tau_P(\text{dist}(a, V0_1, V1_1), V1_1 < 150)$  because the renamed projection<sup>3</sup> of the current constraint store onto the variable of the literal entails the projected constraint store of the generator:  $(V1_1 > 0 \wedge V1_1 < 75) \sqsubseteq V1_1 < 150$ . Since the initial TCLP forest is under construction and depends on itself, the current branch derivation is suspended.

This suspension also causes the former generator to suspend at the state s4.

**s9** is a final state obtained upon backtracking to the state s6ii.

**b1** is the first answer of the second generator.

At this point the suspended calls can be resumed by consuming the answer b1 or by evaluating s2ii. The algorithm first tries to evaluate s2ii and then it will resume s4 consuming b1.

**s10** is a final state obtained upon backtracking to the state s2ii.

**a1** is the first answer of the first generator:  $V0 = b \wedge V1 = 50$ .

**s11** is a final state obtained from the state s4 by consuming b1.

**a2** is the second answer of the first generator:  $V0 = a \wedge V1 > 75 \wedge V1 < 85$ .

**s12** is a final state obtained from the state s8 by consuming a1.

**b2** is the second answer of the second generator.

**s13** is a failed derivation obtained from s8 by consuming a2. It fails because the constraints  $V0 = a \wedge V1 > 75 \wedge V1 < 85$  are inconsistent with the current constraint store. Note that the projection of the constraint store of s8 onto V1 is  $V1 > 0 \wedge V1 < 75$ . Its child is a fail node.

**s14** is a final state obtained from the state s4 by consuming b2.

**a3** is the third answer of the first generator:  $V0 = b \wedge V1 > 125 \wedge V1 < 135$ .

**s15** is a failed derivation obtained from s8 by consuming a3. Its child is a fail node.

This example illustrates why left recursion reduces the execution time and memory requirements when using tabling / TCLP: left recursion will usually create fewer generators. We have also seen that using answers from a more general call, as in the answer resolution of state s8 (i.e., the constraint store of the consumer  $V1_1 > 0 \wedge V1_1 < 75$  is more particular than the constraint store of the generator  $V1_1 < 150$ ), makes it necessary to filter the correct ones (i.e., answer resolution for a2 and a3 failed). This is not required in variant tabling because the answers from a generator are always valid for its consumers.

### Appendix B Step by Step Execution of $\text{dist}/3$ under TCLP(Q)

The trace below shows the step by step execution of the TCLP version of the left recursive distance traversal program in Fig. 5, left, with the query  $?- D \# < 150, \text{dist}(a, Y, D)$ . using the graph in Fig. A 1. In this example we are using the TCLP(Q) interface (Section 4.3). Each step is annotated with the labels used in Fig. 8. The execution starts with the query  $?- D \# < 150, \text{dist}(a, Y, D)$ :

<sup>3</sup> The projection of  $V3 > 0 \wedge V3 < 100 \wedge D1_1 > 0 \wedge D2_1 > 0 \wedge V3 = D1_1 + D2_1 \wedge Z_1 = a \wedge D1_1 > 25 \wedge D1_1 < 35$  onto  $D2_1$  is  $D2_1 > 0 \wedge D2_1 < 75$ . After renaming  $D2_1 = V1_1$ , the resulting projection is  $V1_1 > 0 \wedge V1_1 < 75$ .



- 0** the constraint  $D \# < 150$  in the query is added to the current store (state  $s_0$ ). Then  $\langle \text{dist}(a, Y, D), D < 150 \rangle$  is called and the tabling engine takes the control of the execution calling  $\text{tabled\_call}(\text{dist\_aux}(a, Y, D))$ .
- 1**  $\text{call\_lookup\_table}/3$  initializes and saves (after renaming)  $\text{dist\_aux}(a, V_0, V_1)$ , because it is the first occurrence, and returns  $\text{Vars} = [D]$  and  $\text{Gen} = \$1$ , where  $\$1$  is the reference for this generator.
- 2**  $\text{store\_projection}([D], \text{ProjStore})$  returns  $\text{ProjStore} = ([V_1], [V_1 \# < 150])$ .
- 3**  $\text{member}/2$  fails because the list of projected constraint stores associated to  $\text{Gen} = \$1$  is empty.
- 4**  $\text{save\_generator}/3$  saves  $([V_1], [V_1 \# < 150])$  in the list of projected constraint stores associated to  $\text{Gen} = \$1$  (state  $s_1$ ).
- 7**  $\text{execute\_generator}/2$  evaluates the generator against the first clause of  $\text{dist\_aux}/3$  and adds the body of the clause to the resolvent of the state  $s_{2i}$ . Then the constraints of the resolvent,  $[D_1 \# > 0, D_2 \# > 0, D \# = D_1 + D_2]$ , are added to the constraint store (state  $s_3$ ) and  $\langle \text{dist}(a, Z, D_1), D < 150 \wedge D_1 > 0 \wedge D_2 > 0 \wedge D = D_1 + D_2 \rangle$  is called.
- 0** the tabling engine reenters the tabled execution with  $\text{tabled\_call}(\text{dist\_aux}(a, Z, D_1))$ .
- 1**  $\text{call\_lookup\_table}(\text{dist\_aux}(a, Z, D_1), \text{Vars}, \text{Gen})$  returns  $\text{Vars} = [D_1]$  and  $\text{Gen} = \$1$ , the reference to the previous generator,  $\text{dist\_aux}(a, V_0, V_1)$ .
- 2**  $\text{store\_projection}([D_1], \text{ProjStore})$  returns  $\text{ProjStore} = ([V_1], [V_1 \# > 0, V_1 \# < 150])$ . For clarification, the projection of the current constraint store  $D < 150 \wedge D_1 > 0 \wedge D_2 > 0 \wedge D = D_1 + D_2$  onto  $D_1$  is  $D_1 > 0 \wedge D_1 < 150$ .
- 3**  $\text{member}/2$  retrieves the projected constraint store  $\text{ProjStore}_G = ([V_1], [V_1 \# < 150])$ .
- 5**  $\text{call\_entail}/2$  succeeds because  $(D < 150 \wedge D_1 > 0 \wedge D_2 > 0 \wedge D = D_1 + D_2) \sqsubseteq D_1 < 150$ .
- 6**  $\text{suspend\_consumer}/1$  suspends the current call  $\text{dist\_aux}(a, Z, D_1)$  (state  $s_3$ , waiting for the answer of the current TCLP tree,  $\text{Ans}(s_1)$ ).
- 7** The evaluation of the generator backtracks to evaluate the other clause (state  $s_{2i}$ ). Now the current constraint store is  $D \# < 150$  and the call  $\langle \text{edge}(a, Y, D), D < 150 \rangle$  unifies with  $\text{edge}(a, b, 50)$  (state  $s_4$ ). The first answer is found and  $\text{new\_answer}/0$  is invoked to collect the answer.
- 8**  $\text{answer\_lookup\_table}/2$  stores the Herbrand constraints<sup>4</sup> of the answer,  $Y = b \wedge D = 50$ , returning  $\text{Vars} = []$  and  $\text{Ans} = \$a_1$ , where  $\$a_1$  is the reference for this answer.
- 9**  $\text{store\_projection}/2$  returns  $\text{ProjStore} = ([], [])$ .
- 10**  $\text{member}/2$  fails because the list of projected constraint stores associated to  $\$a_1$  is empty.
- 11**  $\text{save\_answer}/2$  saves  $([], [])$  in the list of the answer constraints associated to  $\$a_1$  (state  $a_1$ ). The first answer is collected.
- 14** the tabling engine resumes the goal suspended at state  $s_3$  and  $\text{member}/2$  retrieves the Herbrand constraints  $Y = b \wedge D = 50$  and the answer constraint  $([], [])$ .

<sup>4</sup> In solvers written in Prolog and implemented using attributed variables, such as CLP(Q) and CLP(R), it is usual that variables lose their association with the constraints where they appeared when these variables become ground. As ground terms do not have attributes attached,  $D = 50$  is handled as part of the Herbrand constraints.

<pre> 1 fib(N,F) :- 2   ⋮ 3   F #= F1 + F2, 4   fib(N1, F1), 5   fib(N2, F2). </pre>		<pre> 1 fib(N,F) :- 2   ⋮ 3   fib(N1, F1), 4   fib(N2, F2), 5   F #= F1 + F2. </pre>
--	--	--

Fig. C 1: Two versions of `fib/2`:  $\mathbb{Q}$  and  $\mathbb{R}$  (left) vs.  $\mathbb{D}_{\leq}$  (right).

- 15** `apply_answer/2` adds the answer to the current constraint store (state s5).
- 7** the execution continues resolving  $\langle \text{edge}(b, Y, D2), \dots \wedge D2 < 100 \rangle$  which unifies with the clause `edge(b, a, D2) :- D2#>25, D2#<35` (state s6). The second answer is found.
- 8** `answer_lookup_table/2` stores the Herbrand constraints of the answer,  $V0 = a$ , returning `Vars=[D]` and `Ans=$a2`.
- 9** `store_projection/2` returns `ProjStore=([V1], [V1#>75, V1#<85])`.
- 10** `member/2` fails because the list of projected constraint stores associated to `Ans=$a2` is empty.
- 11** `save_answer/2` saves  $([V1], [V1#>75, V1#<85])$  in the list of answer constraints associated to `$a2` (state a2). The second answer is collected.
- 14, 15, 7** the tabling engine resumes the suspended goal at state s3 and consumes the second answer following the same steps as with the first one and generating the states s7 and s8. The third answer has been found.
- 8, 9, 10, 11** the answer is collected and  $([V1], [V1#>125, V1#<135])$  is saved in the list of answer constraint associated to `$a3` (state a3).
- 14, 15** the tabling engine resumes the suspended goal at state s3 and consumes the third answer.
- 7** the execution fails resolving  $\langle \text{edge}(b, V0, D2_1), V1 < 150 \wedge \dots \wedge D1_1 < 135 \rangle$  (states s9 and s10)
- 14, 15** the generator has exhausted all the answers and it does not have any more dependencies, so `complete/0` marks the generator as complete. The query retrieves the answers from the generator one by one and returns them.

### Appendix C Comparison of Mod TCLP using $\mathbb{R}$ , $\mathbb{Q}$ and $\mathbb{D}_{\leq}$

This section highlights that the modularity of TCLP makes it possible to choose the most adequate constraint solver for the specific problem, and that decision should not always be based solely on the performance of the constraint solver, but also on its expressiveness and/or precision. Since TCLP, unlike CLP, uses entailment checking extensively to decide whether to suspend and save / discard answers or not, the performance of entailment is more relevant than in CLP. It also makes its soundness (which can be challenged by e.g. numerical accuracy) critical, as incorrect entailment results can lead to non-termination or to unexpected termination.

We use the doubly recursive Fibonacci program (Fig. C 1). It is well-known that tabling reduces its complexity from exponential to linear, but, in addition, CLP makes it possible to run exactly the same program `fib/2 backwards` to find the index of some Fibonacci

	Mod TCLP( $\mathbb{R}$ )	Mod TCLP( $\mathbb{Q}$ )	Mod TCLP( $\mathbb{D}_{\leq}$ )
<code>fib(N, 832040)</code>	<b>25</b>	61	147
<code>fib(N, 28657)</code>	<b>16</b>	40	69
<code>fib(N, 610)</code>	<b>8</b>	19	24
<code>fib(N, 89)</code>	<b>5</b>	12	13

Table C 1: Run time (ms) comparison for the `fib/2` using  $\mathbb{R}$ ,  $\mathbb{Q}$  and  $\mathbb{D}_{\leq}$ .

number by generating a system of equations whose solution is the index of the given Fibonacci number. Under CLP, the size of this system of equations grows exponentially with the index of the Fibonacci number. However, under TCLP, entailment makes redundant equations not to be added and solving them becomes less expensive.

We have run this benchmark using  $\mathbb{R}$ ,  $\mathbb{Q}$  and  $\mathbb{D}_{\leq}$ . Due to the characteristics of  $\mathbb{D}_{\leq}$  (Section 5.1), the program for this constraint system is slightly different from the ones for  $\mathbb{Q}$  and  $\mathbb{R}$  (Section 4.3). In these two, constraints are placed before the recursive calls. However,  $\mathbb{D}_{\leq}$  can have at most two variables per constraint, and therefore we had to move the constraint `F #= F1 + F2` to the end of the clause (Fig. C 1). This can be detrimental to the performance of `Mod TCLP( $\mathbb{D}_{\leq}$ )`, as value propagation in the constraints is less effective.

Table C 1 shows the experimental results. First, note that the `Mod TCLP( $\mathbb{D}_{\leq}$ )` version is slower than any of the other two. While the implementation of `CLP( $\mathbb{D}_{\leq}$ )` is comparatively faster than `CLP( $\mathbb{R}$ )` and `CLP( $\mathbb{Q}$ )`, moving the `F #= F1 + F2` to the end of the clause (which is necessary to satisfy the instantiation requirements of  $\mathbb{D}_{\leq}$ ) reduces its usefulness to prune the generation of redundant constraints.

Additionally, note that although the solvers for  $\mathbb{R}$  and  $\mathbb{Q}$  are practically the same, `Mod TCLP( $\mathbb{R}$ )` is fastest in all cases, since it uses directly CPU floating point numbers while `CLP( $\mathbb{Q}$ )` implements rational numbers by software. However, there is a drawback: floating point arithmetic is not accurate, and when `CLP( $\mathbb{R}$ )` approximates its results, it can cause (depending on the particular program) non-termination. That would be the case for a query such as `?- fib(N, 23416728348467685)`, which terminates correctly with `Mod TCLP( $\mathbb{Q}$ )`, but it does not (in under five minutes) with `Mod TCLP( $\mathbb{R}$ )`, since the termination condition never holds.