

Temporal Logic, μ Kanren, and a Time-Traveling RDF Database

NATHANIEL RUDAVSKY-BRODY*, Independent Researcher, USA

By adding a temporal primitive to the μ Kanren language and adapting its interleaving search strategy to preserve simultaneity in linear time, we show how a system of temporal relational programming can be implemented. This system is then applied to a practical problem in distributed systems and linked data, as we describe a simple data store that can provide incremental solutions (deltas) to complex queries as new records are added to, or removed from the database.

Additional Key Words and Phrases: miniKanren, relational programming, temporal logic, RDF, microservices

1 INTRODUCTION

When using logic programming to model or interact with stateful resources, it is convenient to have a way to reason about time. To this end, temporal logic, and in particular linear temporal logic, has been formalized as a modal extension to predicate logic. Most definitions of linear temporal logic begin with two modal operators, Next (X) and Until (U), which are used to construct a larger set of operators for conveniently reasoning about time. Informally speaking, for the formulas a and b we say that $X b$ holds if b is true in the next time step, and a $U b$ holds if a is always true until b is true, and b is eventually true. For a good review of this work, see [Orgun and Ma 1994].

A number of implementations of linear temporal logic have been proposed as extensions to logic languages such as Prolog. This paper stems from the simple insight that an analogous modeling of linear time can be achieved in the miniKanren family of languages [Friedman et al. 2018], a shallowly embedded logic programming DSL with interleaving depth-first search. Adding a single temporal primitive and adapting the interleaving search strategy to account for the concepts of ‘now’ (simultaneity) and ‘later’ allows us to build tools for temporal reasoning that turn out to integrate quite well with the basic methods of relational programming. The new temporal primitive provides a way of controlling simultaneity and goal construction in miniKanren programs, and be extended in a number of ways that are interesting from both a theoretical and practical standpoint.

This paper is organized in three parts. In the first two sections we go through the low-level implementation, using as our starting point the minimalist μ Kanren language. Then, we sketch out some preliminary ideas on how this operator might be extended to support a more complete linear temporal logic. Finally, we turn to a real-world application in the domain of distributed systems and linked data, and see how the same low-level operator can be used to build a time-aware accessor to a stateful data structure leveraging miniKanren’s interleaving search.¹

The motivation for the last example is as follows. In distributed systems such as a microservice architecture² we often want to send push updates based on *deltas*, or entries that have been added to or removed from the database. In practice, this often means having a service that indiscriminately pushes all deltas to interested subscribers; it is then the responsibility of each subscriber to filter the deltas and determine which, if any, are relevant to its own operations. If we can calculate deltas to specific queries, however, this whole process can be greatly

*Parts of the research for this paper were undertaken while the author was employed as Semantic Applications Developer at Tenforce (Leuven, Belgium). The author thanks Tenforce for its support in undertaking this research. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of Tenforce.

¹The full code for this paper is available at <https://github.com/nathanielrb/time-in-microKanren>.

²For the specific context, see [Versteden and Pauwels 2016] and [Versteden and Pauwels 2018].

Author’s address: Nathaniel Rudavsky-Brody, Independent Researcher, Sebastopol, California, 95472, USA, nathaniel@quince.studio.

2018. 2475-1421/2018/9-ART1 \$15.00

<https://doi.org/>

refined. Ultimately we will describe how a simple RDF database (triple store) can be implemented using temporal μ Kanren as its query language, that calculates deltas to specific queries via a system of incremental indexes.

2 BASIC DEFINITIONS

We begin by briefly reviewing the definitions of μ Kanren as described in [Hemann and Friedman 2013] and [Hemann et al. 2016], deferring to those articles for discussion and motivations. To better expose the core concepts, we choose to work with the earlier implementation³ which has the advantage of using procedures to represent immature streams, leaving us free to use Scheme promises for time. Readers familiar with its implementation may wish to skip to the next section. Though the two versions of μ Kanren run to 39 and 54 lines of code respectively, here we will only repeat the definitions from [Hemann and Friedman 2013] relevant to understanding the changes that will be introduced afterwards. (The omitted definitions can be referred to in the final implementation included in Appendix A.)

A μ Kanren program operates by applying a *goal*, analogous to a predicate in logic programming and implemented as a procedure of one argument, to a *state*, defined as a substitution (an association list of variables and their values) paired with a variable counter. Variables are represented as vectors containing their variable index. The program can succeed or fail, and when it succeeds returns a sequence of new states, called a *stream*, each one extending the original state by new variable substitutions that make the goal succeed.

Goals are built with four basic goal constructors, `==`, `call/fresh`, `disj` and `conj`, to be defined below. In the following example, a goal made of the disjunction of two `==` goals is applied to the empty state. The program returns a stream of two states, each corresponding to one branch of the disjunction.

```
(define empty-state '(() . 0))

((call/fresh
  (lambda (q)
    (disj (== q 4)
          (== q 5))))
  empty-state)
;; => '(((#(0) . 4)) . 1) (((#(0) . 5)) . 1))
```

The walk operator, not defined here, searches for a variable's value in a substitution, while substitutions are extended (without checking for circularities) by consing on the new substitution. `walk` is used by the `unify` operator, which recursively unifies two variables with respect to a given variable substitution, extending the substitution when it succeeds, and returning `#f` when it fails. It is `unify`'s behavior that defines the basic terms of μ Kanren as being variables, objects equivalent under `eq?`, and pairs of such terms.

The fundamental goal constructor is `==`, which builds a goal that succeeds if its two arguments can be unified in the current state, and otherwise returns the empty stream (`mzero`).

```
(define (== u v)
  (lambda (s/c)
    (let ((s (unify u v (car s/c))))
      (if s (unit `(,s . ,(cdr s/c))) mzero))))
```

The `conj` and `disj` goal constructors return goals that succeed if, respectively, both or either of the goals passed as arguments succeed. They are defined in terms of `mplus` and `bind`, which together implement μ Kanren's interleaving search strategy.

```
(define (disj g1 g2) (lambda (s/c) (mplus (g1 s/c) (g2 s/c))))
```

³See <https://github.com/jasonhemann/microKanren> for the code.

```

95 (define (conj g1 g2) (lambda (s/c) (bind (g1 s/c) g2)))
96
97 (define (mplus $1 $2)
98   (cond
99     ((null? $1) $2)
100    ((procedure? $1) (lambda () (mplus $2 ($1)))) ; interleave
101    (else (cons (car $1) (mplus (cdr $1) $2))))))
102
103 (define (bind $ g)
104   (cond
105     ((null? $) mzero)
106     ((procedure? $) (lambda () (bind ($) g)))
107     (else (mplus (g (car $)) (bind (cdr $) g)))))

```

108 The interleaving itself, which guaranties complete search without the performance penalty of breadth-first
109 search, is effectuated by the second case of `mplus` by exchanging the order of `$1` and `$2`. Moreover, the second
110 case in each of the two operators implements *immature streams* as lambda expressions, allowing programs to
111 return infinite streams of results. In μ Kanren it is the user's responsibility to correctly handle immature streams,
112 invoking them if necessary. A regular stream, namely a pair of a state and a stream, is correspondingly termed a
113 *mature stream*.
114

115 3 TIME IN μ KANREN

116 To model linear time in μ Kanren we begin by introducing a third type of stream, *delayed streams*, to represent
117 goals that are delayed until a later point in time. We can represent delayed streams by promises, and construct
118 them using a single temporal goal constructor `next` analogous to the `Next (X)` operator introduced above.
119

```

120 (define-syntax next
121   (syntax-rules ()
122     ((_ g) (lambda (s/c) (delay (g s/c))))))

```

123 A complex goal might have many levels of delays representing different future points in time. It is important
124 that the temporal order of these subsidiary goals be preserved during interleaving search, both with regards to
125 the ordering of solutions and also to guaranty that a goal (which might refer to a stateful resource) is constructed
126 at the right time. We accomplish this by adapting the interleaving process through the addition of two cases to
127 the definition of `mplus`, so that delayed goals (promises) are shunted right, all promises at the same nested level
128 are recombined together. Among other properties, this preserves the order-independence of conjunction and
129 disjunction.

```

130 (define (mplus $1 $2)
131   (cond
132     ((null? $1) $2)
133     ((procedure? $1) (lambda () (mplus $2 ($1))))
134     ((and (promise? $1) (promise? $2))
135      (delay (mplus (force $1) (force $2)))) ; recombine
136     ((promise? $1) (mplus $2 $1)) ; shunt right
137     (else (cons (car $1) (mplus (cdr $1) $2))))))
138

```

139 `bind` also needs to be modified by adding an additional case to delay the binding of delayed streams with a goal.
140 The utility function `forward` is used to 'fast-forward' through the enclosing delayed stream.
141

```

142 (define (forward g)
143   (lambda (s/c)
144     (let rec ((g s/c))
145       (cond ((null? $) '())
146             ((promise? $) (force $))
147             ((procedure? $) (lambda () (rec ($))))
148             (else (cons (car $) (rec (cdr $)))))))
149
150 (define (bind $ g)
151   (cond
152     ((null? $) mzero)
153     ((procedure? $) (lambda () (bind ($) g)))
154     ((promise? $) (delay (bind (force $) (forward g))))
155     (else (mplus (g (car $)) (bind (cdr $) g))))

```

The resulting 53 lines of code make up the functional core of our temporal μ Kanren. To keep the subsequent code simple, we extend the miniKanren wrappers from the original paper to handle delayed streams, saving the code for Appendix B. Using the miniKanren wrappers, and in particular `run*` instead of `call/fresh` to reify solutions, improves the readability of μ Kanren's data representation. We also introduce a few convenience functions for accessing delayed streams, here defined with the help of `take-right` and `drop-right` from SRFI-1.

```

161 (define (promised $)
162   (take-right $ 0))
163
164 (define (current $)
165   (drop-right $ 0))
166
167 (define (advance $)
168   (let ((p (promised $)))
169     (and p (force p))))

```

We are ready for a simple example. Here we use `disj` to specify that `q` equals `*db*` either now or at the next point in time. The first solution in the returned stream is none other than 1, the current value of `*db*`, while the promise contains the delayed stream to be advanced at a future point in time.

```

174 (define *db* 1)
175
176 (define r
177   (run* (q)
178     (fresh (a b)
179       (disj (== q *db*)
180            (next (== q *db*))))))
181 ;; => (1 . #<promise>)

```

Now we can update our database and advance `r` to continue forward in time. At this point the stream is fully mature, since `next` is not used recursively.

```

185 (set! *db* 2)
186 (advance r)
187 ;; => (2)

```

The shunt-right mechanism turns out to have interesting properties for interleaving search even without referring to a stateful resource, for instance in ordering and grouping solutions, as the following contrived example shows. We start by defining `inco` which recursively generates a ‘now or later’ type of goal, incrementing its variable at each future point in time.

```
(define (inco x)
  (let rec ((n 0))
    (disj (== x n)
          (next (rec (+ n 1))))))
```

The program `s` uses `inco` to generate pairs of such incrementing variables.

```
(define s
  (run* (q)
    (fresh (a b)
      (== q `(,a ,b))
      (conj (inco a) (inco b)))))
```

As we advance `s`, we see that the solutions are grouped in increasing order of the maximum value of the two variables.

```
s
;; => ((0 0) . #<promise>))

(advance s)
;; => ((0 1) (1 0) (1 1) . #<promise>))

(advance (advance s))
;; => ((1 2) (2 0) (2 1) (2 2) (0 2) . #<promise>))

(advance (advance (advance s)))
;; => ((0 3) (2 3) (3 0) (3 1) (3 2) (3 3) (1 3) . #<promise>))
```

Indeed, this points up the difference between two basic uses, or two interpretations, of temporal μ Kanren. In the first case, when interacting with a stateful resource, we are doing logic programming *in* time. When reasoning *about* time in a truly relational way, on the other hand, we are more interested in the combinatoric properties of `next`. We will scratch the surface of this distinction in the next section, but a deeper consideration of its implications remains to be done.

4 TOWARDS A TEMPORAL RELATIONAL PROGRAMMING

As we saw in the introduction, linear temporal logic extends predicate logic with temporal modal operators referring to linear time. Several of these operators can be usefully recovered in temporal μ Kanren using the `next` primitive defined above. We will not develop a formally complete system here, but rather provide some heuristic examples to suggest how such a system might developed, noting along the way its potential limitations and specificities.

Most temporal operators need to be defined recursively. When doing this, it is essential to control when the goal is constructed, since the goal might refer to a stateful resource. For temporal μ Kanren we can do this by wrapping goals in a lambda expression which are evaluated at the appropriate time. The definition of `precedes`, modeled on Weak Until (W) and which stipulates that a goal `g` holds at least until a second goal `h` holds, shows how this can be done.

```

236 (define (precedes* g* h*)
237   (let ((g (g*)) (h (h*)))
238     (disj h (conj g (next (precedes* g* h*))))))
239

```

```

240 (define-syntax precedes
241   (syntax-rules ()
242     ((_ g h) (precedes* (lambda () g) (lambda () h)))))
243

```

244 Translating another basic operator, Always (G), reveals one of the pitfalls of naively mapping temporal logic to relational programming. Our first impulse is to define it analogously to `until`.

```

245 (define (always* g*)
246   (let ((g (g*)))
247     (conj g (next (always* g*)))))
248

```

```

249 (define-syntax always
250   (syntax-rules ()
251     ((_ g) (always* (lambda () g)))))
252

```

253 A little experimentation, however, shows how this definition can be problematic. A goal constructed with `always` will return a promise so long as the goal it encloses holds, and the empty stream as soon as it fails. This behavior can be useful for a program that wants to verify whether a state still holds, but `always` will never construct a mature stream of solutions, since clearly the goal will never have been true *forever*, at least not in linear time.

258 There are two ways of addressing this limitation, depending on how we want to use our operator. One is to modify the definition of `next` in such a way as to allow for an `eot` (end-of-time) accessor that would essentially cut off the forward recursion of delayed promises. Another approach is to define the weaker and impure `as-long-as`, which constructs a stream of solutions to the goal `h` that continues as long as the goal `g` still holds. This technique of writing ‘guarded’ goal constructors is reminiscent of the implementation of cuts in `miniKanren`.

```

263 (define (as-long-as* g* h*)
264   (let ((g (g*)) (h (h*)))
265     (lambda (s/c)
266       (let (($ (g s/c))
267           (if (null? $) mzero
268               (bind $ (disj h (next (as-long-as* g* h*))))))))))
269

```

```

270 (define-syntax as-long-as
271   (syntax-rules ()
272     ((_ g h) (as-long-as* (lambda () g) (lambda () h)))))
273

```

274 The `Eventually` (F) operator also has many useful applications, but presents a similar caveat regarding infinite time. Still, the following naive definition will be useful in some contexts.

```

275 (define (eventually* g*)
276   (let ((g (g*)))
277     (disj g (next (eventually* g*)))))
278

```

```

279 (define-syntax eventually
280   (syntax-rules ()
281     ((_ g) (eventually* (lambda () g)))))
282

```

283 The guarded definition, however, provides more intuitive behavior, by cutting off further recursion when the
 284 goal produces a non-empty stream of solutions.

```
285 (define (eventually* g*)
286   (let ((g (g*)))
287     (lambda (s/c)
288       (let (($ (g s/c)))
289         (if (null? $)
290             (bind (list s/c) (next (eventually* g*)))
291             $))))))
```

```
293 (define-syntax eventually
294   (syntax-rules ()
295     ((_ g) (eventually* (lambda () g)))))
```

296 Either one can be used to define an analogue of strong Until, which along with Next is one of the building blocks
 297 of linear temporal logic.

```
299 (define-syntax until
300   (syntax-rules ()
301     ((_ g h) (conj (precedes g h) (eventually h)))))
```

302 Clearly this is only a beginning, and as the above examples suggest, what constitutes a *useful* temporal
 303 relational programming language will be determined in part by the application domain. Still, developing a
 304 formally satisfying version of such a system would be an obvious next step. All this side-stepping of infinite time
 305 also begs the question of what it really means to model it, a question we will not explore further here, since our
 306 primary focus is on using miniKanren to interact with stateful resources in a temporally-aware way. Giving a
 307 satisfactory answer to this question could be a fruitful direction for further work.

308 Another would be recovering temporal μ Kanren's functionality in a version of miniKanren with constraints,
 309 namely cKanren [Alvis et al. 2011] or the extended μ Kanren described in [Hemann and Friedman 2015]. It would
 310 also be interesting to explore the integration of temporal logic with different representations of negation in
 311 miniKanren that have been pursued in both the published and unpublished literature.

313 5 CALCULATING DELTAS WITH INCREMENTAL SEARCH

314 Now we turn to the practical application described in the introduction, and see how our temporal primitive next
 315 can be used to implement a simple data store with temporally-aware incremental search. We begin by describing
 316 the implementation, and then present the full code in the following section.

317 Our goal is to design an RDF database (triple store) using temporal μ Kanren as its query language, that
 318 calculates new or invalidated solutions to specific queries as records are updated in the database. By using the
 319 next constructor and a simple system of incremental indexes, we can store the final search positions for a query.
 320 Running a query will return both the current results and a delayed stream that when advanced will continue
 321 searching at the previous search-tree's leaves. Therefore, advancing the delayed stream after the database has
 322 been updated will return solutions that have been added to, or subtracted from the solution set.

323 An RDF database [Manola and Miller 2004] stores semantic facts, or triples,⁴ made up of a *subject* (a URI), a
 324 *predicate* (a URI), and an *object* (a URI, or a string, boolean or numeric literal).

```
326 <http://ex.org/people/1> <http://xmlns.com/foaf/0.1/name> "John"
```

328 ⁴Here we will only consider triples, though triple stores actually store quads, with the addition of the *graph* element.

A minimalist triple store can be implemented using three indexes, allowing for the retrieval of groups of triples matching a given query pattern: spo, pos, and osp, where s is the subject, p the predicate, and o the object. For the query pattern ?s <P> <O>, for instance, ?s is our only variable so the database would use the pos index for fastest lookup. Here, however, we want to know *incremental indexes*: for a given s, we need the indexed ps in a form that can be easily compared to future ps to determine whether new keys have been added.

Our database is therefore defined as a set of seven incremental indexes (null, s, sp, p, po, o, os) plus the spo index for full triples. Adding a triple to the database is a matter of consing each element to the appropriate incrementals lists and adding a truthy value (here #t, though this could also be a more meaningful identifier or timestamp) to the full spo index. When deleting a triple, we leave the index keys, but update the triple's value in spo to #f. Here are the incremental indexes after adding the triple <A> <C>.

```

340 <∅ [ #f => (<A> ...) ] >
341 <s [ <A> => (<B> ...) ] >
342 <sp [ (<A> <B>) => (<C> ...) ] >
343 <p [ <B> => (<C> ...) ] >
344 <po [ (<B> <C>) => (<A> ...) ] >
345 <o [ <C> => (<A> ...) ] >
346 <os [ (<C> <A>) => (<B> ...) ] >
347 <spo [ (<A> <B> <C>) => #t ] >

```

The main accessor is the goal constructor triple-nolo ('triple-now-or-later') that descends recursively through the incremental indexes one level at a time, constructing a stream with the current indexes and saving the last search positions in a delayed stream. A dynamic parameter is used to specify the current database state.

To see our data store in action, we consider the following query written in SPARQL 1.1 [Harris and Seaborne 2013], the standard query language for RDF data stores. This query will be translated into temporal μ Kanren and run against successive states of the database.

```

355 SELECT ?o
356 WHERE {
357   <S> <P> ?o.
358   <Q> <R> ?o.
359 }

```

First, however, we need to define an empty database.

```
(define db0 (empty-db))
```

Then the SPARQL query is translated into the temporal μ Kanren goal r using triple-nolo. To make it clear what is going on, we keep track of the individual delta flags for each triple goal, though in practice we usually want to know only whether the solution was added (all +s) or removed (at least one -). Since there are no solutions to our query, the program returns a delayed stream.

```

367 (define r
368   (parameterize ((latest-db db0))
369     (run* (q)
370       (fresh (o deltas d1 d2)
371         (== q ` (,deltas ,o)
372           (== deltas ` (,d1 ,d2))
373           (triple-nolo d1 '<S>' '<P>' o)
374           (triple-nolo d2 '<Q>' '<R>' o))))))
375 ;; => #<promise>
376

```


377 Now we can add some triples and advance the delayed stream.

```
378 (define db1
379   (add-triples db0 '((<S> <P> <O1>)
380                     (<S> <P> <O2>)
381                     (<Q> <R> <O1>)
382                     (<A> <B> <C>))))
383
```

```
384 (parameterize ((latest-db db1))
385   (advance r))
386 ;; => (((+ +) <O1>) . #<promise>)
387
```

388 If we delete a triple that contributed to one of our solutions, that solution will be returned with a negative delta
389 flag in the next iteration.

```
390 (define db2
391   (delete-triples db1 '((<S> <P> <O1>))))
392
```

```
393 (parameterize ((latest-db db2))
394   (advance (advance r)))
395 ;; => (((+ -) <O1>) . #<promise>)
396
```

396 Finally we add that triple back along with some new solutions, to get positive deltas.

```
397 (define db3
398   (add-triples db2 '((<S> <P> <O1>)
399                     (<S> <P> <O3>)
400                     (<Q> <R> <O3>)
401                     (<S> <P> <M>)
402                     (<Q> <R> <M>))))
403
404 (parameterize ((latest-db db3))
405   (advance (advance (advance r))))
406 ;; => (((+ +) <O1>) ((+ +) <M>) ((+ +) <O3>) . #<promise>)
407
```

408 Here we have been proceeding linearly, but since the database states are persistent we can calculate deltas
409 over any two states. Keeping track of a stream of states and indexing them on time will therefore give us a truly
410 time-traveling data store.

412 6 IMPLEMENTING THE RDF STORE

413 Now we are ready to walk through the full code for the data store.

414 We start by defining `project`, a standard miniKanren operator [Byrd 2009] that binds a set of variables to their
415 current substitutions. It will be used control `triple-nolo`'s recursive descent through the indexes.

```
416 (define-syntax project
417   (syntax-rules ()
418     ((_ (x ...) g)
419      (lambda (s/c)
420        (let ((x (walk x (car s/c))) ...)
421          (g s/c))))))
422
```

423

424 A database is defined as a set of seven incremental indexes plus the full spo index, as described above. We keep
 425 a separate incremental index for each index level, storing the incrementals as a simple cons list along with a map
 426 for quick existence checking. For each index itself we use a hash array mapped trie⁵ that allows for persistently
 427 storing successive states of the database through path copying.

```
428 (define-record db null s sp p po o os spo)
```

```
429
```

```
430 (define-record incrementals map list)
```

```
431
```

```
432 (define (empty-incrementals)
```

```
433   (make-incrementals (persistent-map) '()))
```

```
434
```

```
435 (define (empty-db)
```

```
436   (apply make-db
```

```
437     (make-list 8 (persistent-map))))
```

```
438
```

Dynamic parameters are used to specify the current database state.

```
439
```

```
(define latest-db (make-parameter (empty-db)))
```

```
440
```

```
441 (define (latest-incrementals accessor key)
```

```
442   (incrementals-list
```

```
443     (map-ref (accessor (latest-db)) key (empty-incrementals))))
```

```
444
```

```
445 (define (latest-triple s p o)
```

```
446   (map-ref (db-spo (latest-db))
```

```
447     (list s p o)))
```

```
448
```

To update the indexes we update the quick-check map to #t and cons the new value to the incrementals list.

```
449 (define (update-incrementals table key val)
```

```
450   (let ((incrementals (map-ref table key (empty-incrementals))))
```

```
451     (map-add table key
```

```
452       (if (map-ref (incrementals-map incrementals) val)
```

```
453         incrementals
```

```
454         (make-incrementals
```

```
455           (map-add (incrementals-map incrementals) val #t)
```

```
456           (cons val (incrementals-list incrementals))))))
```

```
457
```

```
458 (define (update-triple table triple val)
```

```
459   (map-add table triple val))
```

```
460
```

Updating (adding or removing) a triple is therefore a matter of updating all eight indexes.

```
461
```

```
462
```

```
463
```

```
464
```

```
465
```

```
466
```

```
467
```

```
468
```

⁵In the Chicken implementation presented here, we use Persistent Hash Maps [Heidkamp 2013], derived from Ian Price's Hash Array Mapped Tries [Price 2014].

```
469
```

```
470
```

```

471 (define (update-triples DB triples val)
472   (let loop ((triples triples)
473             (i/null (db-null DB))
474             (i/s (db-s DB))
475             (i/sp (db-sp DB))
476             (i/p (db-p DB))
477             (i/po (db-po DB))
478             (i/o (db-o DB))
479             (i/os (db-os DB))
480             (i/spo (db-spo DB)))
481     (if (null? triples)
482         (make-db i/null i/s i/sp i/p i/po i/o i/os i/spo)
483         (match (car triples)
484             ((s p o)
485              (loop (cdr triples)
486                    (update-incrementals i/null #f s)
487                    (update-incrementals i/s s p)
488                    (update-incrementals i/sp (list s p) o)
489                    (update-incrementals i/p p o)
490                    (update-incrementals i/po (list p o) s)
491                    (update-incrementals i/o o s)
492                    (update-incrementals i/os (list o s) p)
493                    (update-triple i/spo (list s p o) val)))))))
494
495 (define (add-triples DB triples) (update-triples DB triples #t))
496
497 (define (delete-triples DB triples) (update-triples DB triples #f))
498
499 (define (add-triple s p o)
500   (add-triples `((,s ,p ,o))))
501
502 (define (delete-triple s p o)
503   (delete-triples `((,s ,p ,o))))

```

At last we can define our main accessor, `triple-nolo`. This operator descends recursively through the indexes, selecting which index to use according to which of its arguments are variables (`var?`) and using `project` in the recursion step to reduce the number of variables by one. The stream of solutions is constructed by `stream` that recurses through the incrementals list (`indexes`), starting at the head and stopping when it reaches the previous head (`ref`), and storing the new head of the list in another recursive call wrapped in `next` as the delayed stream of future solutions.

```

510
511
512
513
514
515
516
517

```

```

518 (define (triple-nolo delta s p o)
519   (let ((mkstrm (lambda (var accessor key)
520               (let* ((get-incrementals (lambda ()
521                                         (latest-incrementals accessor key)))
522                      (initial-indexes (get-incrementals)))
523                 (let stream ((indexes initial-indexes)
524                             (ref '())
525                             (next-ref initial-indexes))
526                   (if (equal? indexes ref)
527                       (next
528                        (let ((vals (get-incrementals)))
529                          (stream vals next-ref vals)))
530                       (disj
531                        (conj (== var (car indexes))
532                             (project (s p o)
533                                       (triple-nolo delta s p o)))
534                        (stream (cdr indexes) ref next-ref))))))))))
535   (cond ((and (var? s) (var? p) (var? o)) (mkstrm s db-null #f))
536         ((and (var? s) (var? p)) (mkstrm s db-o o))
537         ((and (var? s) (var? o)) (mkstrm o db-p p))
538         ((and (var? p) (var? o)) (mkstrm p db-s s))
539         ((var? s) (mkstrm s db-po (list p o)))
540         ((var? p) (mkstrm p db-os (list o s)))
541         ((var? o) (mkstrm o db-sp (list s p)))
542         (else
543          (let leaf ((ref #f))
544            (let ((v (latest-triple s p o)))
545              (cond ((eq? v ref) (next (leaf v)))
546                    (v (disj (== delta '+) (next (leaf v))))
547                    (else (disj (== delta '-') (next (leaf v)))))))))))))

```

548 Though this reduced example is clearly far from a production-ready system, we hope we have demonstrated
549 one of the practical uses of temporal relational programming. Indeed, we intend to pursue these tools in an
550 industrial setting, and the full implementation will be a laboratory for other practical applications of miniKanren,
551 such as drawing on [Byrd et al. 2012] to compile SPARQL to miniKanren and using search ordering as explored
552 in [Swords and Friedman 2013] to aid with query optimization.

554 A APPENDIX: TEMPORAL μ KANREN

```

555 (define (var c) (vector c))
556 (define (var? x) (vector? x))
557 (define (var=? x1 x2) (= (vector-ref x1 0) (vector-ref x2 0)))
558
559 (define (walk u s)
560   (let ((pr (and (var? u) (assp (lambda (v) (var=? u v)) s))))
561     (if pr (walk (cdr pr) s) u)))
562
563
564

```

```

565 (define (ext-s x v s) `((,x . ,v) . ,s))
566
567 (define (= u v)
568   (lambda (s/c)
569     (let ((s (unify u v (car s/c))))
570       (if s (unit `(,s . ,(cdr s/c)) mzero))))))
571
572 (define (unit s/c) (cons s/c mzero))
573 (define mzero '())
574
575 (define (unify u v s)
576   (let ((u (walk u s)) (v (walk v s)))
577     (cond
578       ((and (var? u) (var? v) (var=? u v)) s)
579       ((var? u) (ext-s u v s))
580       ((var? v) (ext-s v u s))
581       ((and (pair? u) (pair? v))
582        (let ((s (unify (car u) (car v) s)))
583          (and s (unify (cdr u) (cdr v) s))))
584       (else (and (eqv? u v) s))))))
585
586 (define (call/fresh f)
587   (lambda (s/c)
588     (let ((c (cdr s/c)))
589       ((f (var c)) `(,(car s/c) . ,(+ c 1))))))
590
591 (define (disj g1 g2) (lambda (s/c) (mplus (g1 s/c) (g2 s/c))))
592 (define (conj g1 g2) (lambda (s/c) (bind (g1 s/c) g2)))
593
594 (define (mplus $1 $2)
595   (cond
596     ((null? $1) $2)
597     ((procedure? $1) (lambda () (mplus $2 ($1))))
598     ((and (promise? $1) (promise? $2))
599      (delay (mplus (force $1) (force $2))))
600     ((promise? $1) (mplus $2 $1))
601     (else (cons (car $1) (mplus (cdr $1) $2))))))
602
603 (define (forward g)
604   (lambda (s/c)
605     (let rec ((g s/c))
606       (cond ((null? $) '())
607             ((promise? $) (force $))
608             ((procedure? $) (lambda () (rec ($))))
609             (else (cons (car $) (rec (cdr $)))))))
610
611

```

```

612 (define (bind $ g)
613   (cond
614     ((null? $) mzero)
615     ((procedure? $) (lambda () (bind ($) g)))
616     ((promise? $) (delay (bind (force $) (forward g))))
617     (else (mplus (g (car $)) (bind (cdr $) g))))))
618
619 (define-syntax next
620   (syntax-rules ()
621     ((_ g) (lambda (s/c) (delay (g s/c))))))
622
623 B APPENDIX: MINIKANREN WRAPPERS
624 Here we adapt the miniKanren control operators described in [Hemann and Friedman 2013] for delayed streams.
625 The main changes are to the definitions of run, run*, pull, take-all and take. Unlike the original paper, we
626 do not use Zzz to wrap goals in conj+ and disj+, since this makes it difficult to control the goal construction
627 time at different levels of nested delays, a problem when referring to stateful resources. This difficulty should be
628 addressed in future implementations of temporal  $\mu$ Kanren.
629
630 (define-syntax conj+
631   (syntax-rules ()
632     ((_ g) g)
633     ((_ g0 g ...) (conj g0 (conj+ g ...))))))
634
635 (define-syntax disj+
636   (syntax-rules ()
637     ((_ g) g)
638     ((_ g0 g ...) (disj g0 (disj+ g ...))))))
639
640 (define-syntax fresh
641   (syntax-rules ()
642     ((_ () g0 g ...) (conj+ g0 g ...))
643     ((_ (x0 x ...) g0 g ...)
644      (call/fresh
645       (lambda (x0)
646         (fresh (x ...) g0 g ...))))))
647
648 (define-syntax conde
649   (syntax-rules ()
650     ((_ (g0 g ...) ...) (disj+ (conj+ g0 g ...) ...))))
651
652 (define-syntax run
653   (syntax-rules ()
654     ((_ n (x ...) g0 g ...)
655      (let r ((k n) ($) (take n (call/goal (fresh (x ...) g0 g ...))))
656        (cond ((null? $) '())
657              ((promise? $) (delay (r (- k 1) (take k (force $))))))
658

```

```

659         (else (cons (reify-1st (car $))
660                     (r (- k 1) (cdr $)))))))))
661
662 (define-syntax run*
663   (syntax-rules ()
664     ((_ (x ...) g0 g ...)
665       (let r (($ (take-all (call/goal (fresh (x ...) g0 g ...))))
666             (cond ((null? $) '())
667                   ((promise? $) (delay (r (take-all (force $))))
668                   (else (cons (reify-1st (car $))
669                               (r (cdr $)))))))))
670
671 (define empty-state '(() . 0))
672
673 (define (call/goal g) (g empty-state))
674
675 (define (pull $)
676   (cond ((procedure? $) (pull ($)))
677         ((promise? $) $)
678         (else $)))
679
680 (define (take-all $)
681   (let (($ (pull $)))
682     (cond ((null? $) '())
683           ((promise? $) $)
684           (else (cons (car $) (take-all (cdr $))))))
685
686 (define (take n $)
687   (if (zero? n) '()
688       (let (($ (pull $)))
689         (cond ((null? $) '())
690               ((promise? $) $)
691               (else (cons (car $) (take (- n 1) (cdr $))))))
692
693 (define (reify-1st s/c)
694   (let ((v (walk* (var 0) (car s/c))))
695     (walk* v (reify-s v '()))))
696
697 (define (walk* v s)
698   (let ((v (walk v s)))
699     (cond
700       ((var? v) v)
701       ((pair? v) (cons (walk* (car v) s)
702                        (walk* (cdr v) s)))
703       (else v))))
704
705

```

```

706 (define (reify-s v s)
707   (let ((v (walk v s)))
708     (cond
709       ((var? v)
710        (let ((n (reify-name (length s))))
711          (cons `(,v . ,n) s)))
712       ((pair? v) (reify-s (cdr v) (reify-s (car v) s)))
713       (else s))))
714
715 (define (reify-name n)
716   (string->symbol
717    (string-append "_" "." (number->string n))))
718
719 (define (fresh/nf n f)
720   (letrec
721     ((app-f/v*
722      (lambda (n v*)
723        (cond
724          ((zero? n) (apply f (reverse v*)))
725          (else (call/fresh
726                 (lambda (x)
727                   (app-f/v* (- n 1) (cons x v*))))))))))
728    (app-f/v* n '()))
729

```

730 REFERENCES

- 731 Claire E. Alvis, Jeremiah J. Willcock, Kyle M. Carter, William E. Byrd, and Daniel P. Friedman. 2011. cKanren: miniKanren with Constraints. *Proceedings of the 2011 Workshop on Scheme and Functional Programming*.
- 732 William E. Byrd. 2009. *Relational programming in miniKanren: techniques, applications, and implementations*. Ph.D. Dissertation. Bloomington, IN, USA.
- 733 William E. Byrd, Eric Holk, and Daniel P. Friedman. 2012. miniKanren, Live and Untagged: Quine Generation via Relational Interpreters (Programming Pearl). In *Proceedings of the 2012 Annual Workshop on Scheme and Functional Programming (Scheme '12)*. ACM, New York, NY, USA, 8–29. <https://doi.org/10.1145/2661103.2661105>
- 734 Daniel P. Friedman, William E. Byrd, Oleg Kiselyov, and Jason Hemann. 2018. *The Reasoned Schemer, Second Edition*. The MIT Press, Cambridge, MA, USA.
- 735 Steve Harris and Andy Seaborne. 2013. *SPARQL 1.1 Query Language*. Technical Report. <https://www.w3.org/TR/sparql11-query/>
- 736 Moritz Heidkamp. 2013. persistent-hash-map. (2013). <http://wiki.call-cc.org/eggref/4/persistent-hash-map>
- 737 Jason Hemann and Daniel P. Friedman. 2013. μ Kanren: A Minimal Functional Core for Relational Programming. *Proceedings of the 2013 Workshop on Scheme and Functional Programming*.
- 738 Jason Hemann and Daniel P. Friedman. 2015. A Framework for Extending microKanren with Constraints. *Proceedings of the 2015 Workshop on Scheme and Functional Programming*.
- 739 Jason Hemann, Daniel P. Friedman, William E. Byrd, and Might Matthew. 2016. A Small Embedding of Logic Programming with a Simple Complete Search. In *Proceedings of the 12th Symposium on Dynamic Languages (DLS 2016)*. ACM, New York, NY, USA, 96–107. <https://doi.org/10.1145/2989225.2989230>
- 740 Frank Manola and Eric Miller. 2004. *RDF Primer*. Technical Report. <https://www.w3.org/TR/2004/REC-rdf-primer-20040210/>
- 741 Mehmet A. Orgun and Wanli Ma. 1994. An Overview of Temporal and Modal Logic Programming. In *Proc. First Int. Conf. on Temporal Logic - LNAI 827*. Springer-Verlag, 445–479.
- 742 Ian A. Price. 2014. Purely Functional Data Structures in Scheme. (2014). <https://github.com/ijp/pfds/blob/master/hamts.sls>
- 743 Cameron Swords and Daniel Friedman. 2013. rKanren: Guided search in miniKanren. In *Proceedings of the 2013 Workshop on Scheme and Functional Programming*.

753 Aad Versteden and Erika Pauwels. 2016. State-of-the-art Web Applications using Microservices and Linked Data. In *Proceedings of the 4th*
754 *Workshop on Services and Applications over Linked APIs and Data*.

755 Aad Versteden and Erika Pauwels. 2018. mu.semte.ch: State-of-the-art Web Applications fuelled by Linked Data. (2018). Presentation at the
756 Extended Semantic Web Conference 2018, Heraklion, Greece.

757

758

759

760

761

762

763

764

765

766

767

768

769

770

771

772

773

774

775

776

777

778

779

780

781

782

783

784

785

786

787

788

789

790

791

792

793

794

795

796

797

798

799