

# Data Representation Synthesis

Peter Hawkins, Alex Aiken, Kathleen Fisher, Martin Rinard,  
Mooly Sagiv

presented by Andreas Schröcker  
Concurrency and Memory Management Seminar  
Prof. Christoph Kirsch, University of Salzburg  
June 14, 2012

## Introduction

Motivation

Example

Approach

## Relational Abstraction

Relational Specification

Notation

Relational Operations

## Decomposition and Decomposition Instances

## Querying and Updateing Decomposed Relations

Queries and Query Plans

Mutation: Empty and Insert Operations

Mutation: Remove and Update Operations

## Autotuner

## Experiments

# Datastructures

- ▶ Non-trivial program represents its data internally using dynamically-allocated data structures.
- ▶ Commit to a particular choice of heap data structures that represent the system's state.
- ▶ Must meet several requirements.
  - ▶ The representation must support all of the operations required by the code.
  - ▶ The data structures must be efficient for the workload.
  - ▶ The implementation must be correct.

# Datastructures

- ▶ Choice of data structures has a pervasive influence on the subsequent code.
- ▶ As requirements evolve it is difficult and tedious to change the data structures.
- ▶ For a data representation to be correct, data structure invariants must be enforced by every piece of code that manipulates the heap.


# Operating System Scheduler

- ▶ Each process has:
  - ▶ *pid* - id of the process
  - ▶ *state* - the state of the process (sleeping, running)
  - ▶ *cpu* - the cpu time consumed by the process
- ▶ Adding support for virtualization:
  - ▶ *ns* - processes with the same *pid* may exist in different namespaces

## Approach - Data Representation Synthesis

- ▶ A data structure client describes and manipulates data at a high level as relations.
- ▶ A data structure designer then provides decompositions which describe how those relations should be represented in memory as a combination of primitive data structures.
- ▶ Our compiler RelC takes a relation and its decomposition and emits correct and efficient low-level code that implements the relational interface.

# Data Representation Synthesis



# Advantages

- ▶ Synthesis allows programmers to describe and manipulate data at a high level as relations, while giving control of how relations are represented physically in memory.
- ▶ By abstracting data from its representation, programmers no longer prematurely commit to a particular representation of data.
- ▶ Synthesized representations are correct by construction; so long as the programmer conforms to the relational specification, invariants on the synthesized data structures are automatically maintained.



# Relational Specification

- ▶ A relational specification is a set of column names  $C$  and functional dependencies  $\Delta$
- ▶ Scheduler Example
  - ▶  $\{ns, pid, state, cpu\}$
  - ▶  $\{ns, pid\} \rightarrow \{state, cpu\}$
- ▶ We use relations to abstract a program's data from its representation.
- ▶ Describing particular representations is the task of the decomposition language

# Values, Tuples, Relations

- ▶ *tuple*  $t = \langle c_1 : v_1, c_2 : v_2, \dots \rangle$  maps a set of columns to values.
- ▶  $\text{dom } t = C$ .
- ▶ A relation  $a$  is a set of tuples  $\{t_1, t_2, t_3, \dots\}$ .
- ▶  $t(c)$  is the value of columns  $c$  in tuple  $t$ .
- ▶ We write  $t \supseteq s$  if the tuple  $t$  extends  $s$ , that is  $t(c) = s(c)$  for all  $c$  in  $\text{dom } s$ .
- ▶ We say tuple  $t$  matches tuple  $s$ , written  $t \sim s$ , if they are equal on all common columns.
- ▶ We say tuple  $t$  matches relation  $r$ , written  $t \sim r$ , if  $t$  matches every tuple in  $r$ .
- ▶ We write  $s \triangleleft t$  for a merge of tuples. Taking values of  $t$  wherever the two disagree on a column's value,

# Functional Dependencies

- ▶ Relation  $r$  has functional dependencies (FD)  $C_1 \rightarrow C_2$  if any pair of tuples that are equal on columns  $C_1$  are equal on columns  $C_2$ .
- ▶ We write  $r \models_{fd} \Delta$  if a set of FDs  $\Delta$  hold on relation  $r$ .
- ▶ We write  $\Delta \vdash_{fd} C_1 \rightarrow C_2$  if FD  $C_1 \rightarrow C_2$  is a consequence of FDs  $\Delta$ .

## Relation for Scheduler Example

$$r_s = \{ \langle ns : 1, pid : 1, state : S, cpu : 7 \rangle, \\ \langle ns : 1, pid : 2, state : R, cpu : 4 \rangle, \\ \langle ns : 2, pid : 1, state : S, cpu : 5 \rangle \}$$

# Relational Algebra

We use the standard notation of relational algebra.

- ▶ Union  $\cup$
- ▶ Set Intersection  $\cap$
- ▶ Set Difference  $\setminus$
- ▶ Symetric Difference  $\ominus$
- ▶ Projection  $\pi_C$
- ▶ Natural Join  $r_1 \bowtie r_2$

# Relational Operations

empty () = ref  $\emptyset$

insert  $r t$  =  $r \leftarrow !r \cup \{t\}$

remove  $r s$  =  $r \leftarrow !r \setminus \{t \in !r \mid t \supseteq s\}$

update  $r s u$  =  $r \leftarrow \{ \text{if } t \supseteq s \text{ then } t \triangleleft u \text{ else } t \mid t \in !r \}$

query  $r s C$  =  $\pi_C \{t \in !r \mid t \supseteq s\}$

# Relational Operations

- ▶ insert  $r \langle ns : 7, pid : 42, state : R, cpu : 0 \rangle$
- ▶ query  $r \langle ns : 7, pid : 42 \rangle \{state, cpu\}$
- ▶ update  $r \langle ns : 7, pid : 42 \rangle \langle state : S \rangle$
- ▶ remove  $r \langle ns : 7, pid : 42 \rangle$

# Relational Operations

```
class scheduler_relation
{
    void insert(tuple_cpu_ns_pid_state const &r);

    void remove(tuple_ns_pid const &pattern);

    void update(tuple_ns_pid const &pattern,
                tuple_cpu_state const &changes);

    void query(tuple_state const &input,
               iterator_state__ns_pid &output);

    ...
};
```




# Decomposition

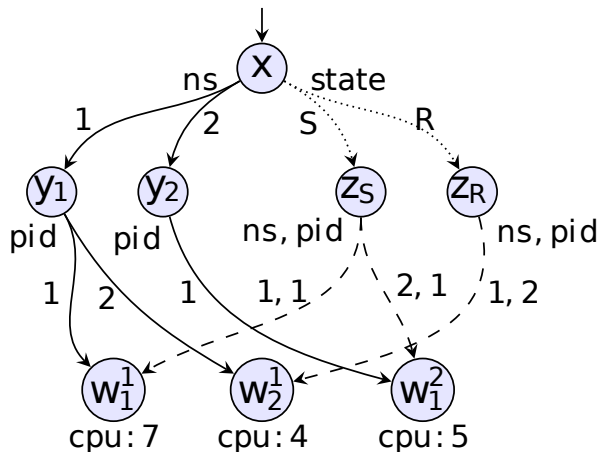
- ▶ Decompositions describe how to represent relations as a combination of primitive data structures.
- ▶ A decomposition is a static description of the structure of data, akin to a type.
- ▶ Its run-time (dynamic) counterpart is the decomposition instance, which describes the representation of a particular relation using the decomposition.

# Decomposition

solid line - hash map  
dotted line - vector  
dashed line - dl. list



## Decomposition Instance



# Decomposition Language

$\hat{p} ::= C \mid C \xrightarrow{\psi} v \mid \hat{p}_1 \bowtie \hat{p}_2$   
 $\hat{d} ::= \text{let } v: C_1 \triangleright C_2 = \hat{p} \text{ in } \hat{d} \mid v$   
 $\psi ::= \text{dlist} \mid \text{htable} \mid \text{vector} \dots$

decomposition primitives  
decomposition  
data structures

# Decomposition Language

$p ::= t \mid \{t \mapsto v_{t'}, \dots\} \mid p_1 \bowtie p_2$  instance primitives  
 $d ::= \text{let } \{v_t = p, \dots\} \text{ in } d \mid v_{\langle \rangle}$  data structures

# Decomposition Language

let  $w : \{ns, pid, state\} \triangleright \{cpu\} = \{cpu\}$  in

let  $y : \{ns\} \triangleright \{pid, cpu\} = \{pid\} \xrightarrow{htable} w$  in

let  $z : \{state\} \triangleright \{ns, pid, cpu\} = \{ns, pid\} \xrightarrow{dlist} w$  in

let  $x : \emptyset \triangleright \{ns, pid, cpu, state\} =$   
 $(\{ns\} \xrightarrow{htable} y) \bowtie (\{state\} \xrightarrow{vector} z)$  in  $x$

# Decomposition Language

$$\text{let } \left\{ \begin{array}{l} w_{\langle ns:1, pid:1, state:S \rangle} = \langle cpu : 7 \rangle, \\ w_{\langle ns:1, pid:2, state:R \rangle} = \langle cpu : 4 \rangle, \\ w_{\langle ns:2, pid:1, state:S \rangle} = \langle cpu : 5 \rangle, \end{array} \right\} \text{ in}$$

$$\text{let } \left\{ \begin{array}{l} y_{\langle ns:1 \rangle} = \left\{ \begin{array}{l} \langle pid : 1 \rangle \mapsto w_{\langle ns:1, pid:1, state:S \rangle}, \\ \langle pid : 2 \rangle \mapsto w_{\langle ns:1, pid:2, state:R \rangle} \end{array} \right\}, \\ y_{\langle ns:2 \rangle} = \left\{ \langle pid : 1 \rangle \mapsto w_{\langle ns:2, pid:1, state:S \rangle} \right\} \end{array} \right\} \text{ in}$$

$$\text{let } \left\{ \begin{array}{l} z_{\langle state:S \rangle} = \left\{ \begin{array}{l} \langle ns : 1, pid : 1 \rangle \mapsto w_{\langle ns:1, pid:1, state:S \rangle}, \\ \langle ns : 1, pid : 2 \rangle \mapsto w_{\langle ns:1, pid:2, state:R \rangle} \end{array} \right\}, \\ z_{\langle state:R \rangle} = \left\{ \langle ns : 2, pid : 1 \rangle \mapsto w_{\langle ns:2, pid:1, state:S \rangle} \right\} \end{array} \right\} \text{ in}$$

$$\text{let } \left\{ \begin{array}{l} x_{\langle \rangle} = \left\{ \begin{array}{l} \langle ns : 1 \rangle \mapsto y_{\langle ns:1 \rangle}, \\ \langle ns : 2 \rangle \mapsto y_{\langle ns:2 \rangle} \end{array} \right\} \bowtie \\ \left\{ \begin{array}{l} \langle state : S \rangle \mapsto z_{\langle state:S \rangle}, \\ \langle state : R \rangle \mapsto z_{\langle state:R \rangle} \end{array} \right\} \end{array} \right\} \text{ in } x_{\langle \rangle}$$

- ▶ Abstraction function
  - ▶ Computes the relation represented by a given decomposition instance.
- ▶ Well-formedness criteria
  - ▶ Check that a decomposition instance is a well-formed instance of a particular decomposition.
- ▶ Adequacy conditions
  - ▶ Which are sufficient conditions for a decomposition to faithfully represent a relation.



# Querying and Updateing Decomposed Relations

Two basic kinds of relational operations:

- ▶ Queries
- ▶ Mutations

# Queries

Queries are implemented in two stages:

- ▶ Query planning
  - ▶ Attempt to find most efficient execution plan for a query.
- ▶ Query execution
  - ▶ Evaluates a particular query plan over a decomposition instance.

This approach is well known in the database literature.

## Query Planer implementatin in RelC

- ▶ Each strategy has a different computational complexity
- ▶ The query planner enumerates the alternatives and chooses the “best” strategy.

# Query Plan

Query plan is a tree of query plan operators.

$$q ::= qunit \mid qscan(q) \mid qllookup(q) \mid qlr(q, lr) \mid qjoin(q_1, q_2lr)$$
$$lr ::= left \mid right$$

## Scheduler Example

- ▶ The query
  - ▶ query  $r\langle ns : 7, pid : 42 \rangle\{cpu\}$
- ▶ Possible query plan
  - ▶  $q_{cpu} = qlr(qlookup(qlookup(qunit)), left)$
- ▶ Perform query  $q_{cpu}$  on an instance  $d$ 
  - ▶  $dexec\ q_{cpu}\ d\ \langle ns : 7, pid : 42 \rangle$

## Scheduler Example

- ▶ The query
  - ▶ query  $r\langle ns : 7, state : R \rangle\{cpu\}$
- ▶ Possible query plan
  - ▶  $q_1 =$   
 $qjoin(qlookup(qscann(qunit)), qlookup(qlookup(qunit)), left)$
  - ▶  $q_2 = qlr(qlookup(qscan(qunit)), rights)$

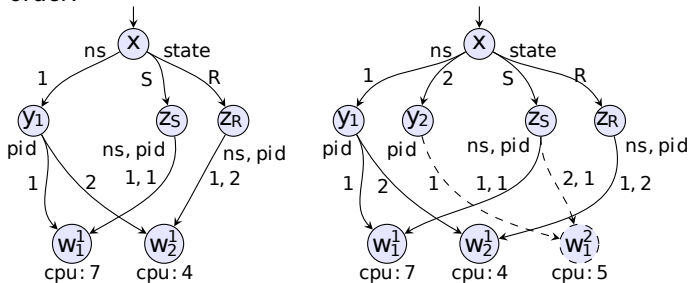
## Query Validity

Not every query plan is a correct strategy for evaluating a query.  
We must check three properties:

- ▶ Query produce all of the columns requested as output.
- ▶ When performing a lookup all necessary key columns are available.
- ▶ Enough columns are computed on each side of a join.

# Insert

Perform insertion over the nodes of a decomposition in topological order.




$$t = \langle ns : 2, pid : 1, state : S, cpu : 5 \rangle$$




# Remove

Create a cut and remove nodes matching the tuple.



cut for columns  $\{ns, pid\}$



cut for columns  $\{state\}$

# Update

- ▶ Semantically, updates are removal followed by an insertion
- ▶ Updates are performed inplace.
- ▶ Only common case is supported - no key columns.

# Autotuner





- ▶ Attempts to find the best decomposition of a relation.
- ▶ Takes as input a benchmark program, that produces as output a cost value, together with the relation to optimize.
- ▶ Constructs all possible decompositions up to a number of edges.

# Experiments

- ▶ Micro-benchmarks
- ▶ Real World Systems.


# Graph Benchmark

Relation  $\{src, dst, weight\}$ ,  $src, dst \rightarrow weight$ .  
 Decompositions up to size 4.



# Real World Benchmarks

System	Original		Synthesis	
	Everything	Module	Decomposition	Module
thttpd	7050	402	42	239
Ipcap	2138	899	55	794
ZTopo	5113	1083	39	1048



## Experiments - Summary

Experiments show that:

- ▶ Different choices of decomposition lead to significant changes in performance
- ▶ The best performance is comparable to existing hand-written implementations
- ▶ The resulting code is concise and the soundness of the compiler guarantees that the resulting data structures are correct by construction.