

M33
1-46

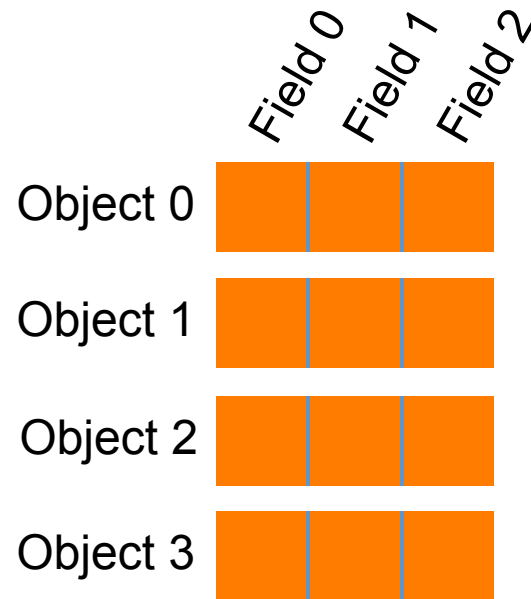
Dualities in Programming Languages

Martin Hirzel and Priya Nagpurkar

IBM Watson Research Center

PLDI-FIT 2010

Row-Based Layout




Values of different **fields** but same **object** contiguous.

Good locality if few hot **objects**.

Field access dereferences **objectAddress** + **fieldOffset**.

Column-Based Layout




Values of different **objects** but same **field** contiguous.

Good locality if few hot **fields**.

Field access dereferences **fieldAddress** + **objectOffset**.

Row-Based \cong Column-Based Layout



Values of different **fields** but same **object** contiguous.

Good locality if few hot **objects**.

Field access dereferences **objectAddress** + **fieldOffset**.

Observations

- Duality =
 pair of concepts
 + terminology substitution.
- Not always perfect
(e.g., memory management
easier for row-based layout.)

Observations

- Duality = pair of concepts + terminology substitution.
- Not always perfect (e.g., memory management easier for row-based layout.)
- Good excuse for fancy formatting.

Dualities in Programming Languages

Martin Hertz Priya Nagarkar
IBM Watson Research Center
(hertz, nagarkar)@ibm.com

Abstract

A duality can be thought of as a pair of concepts and a mapping between their terminologies, with the substituting the concept-specific terminology terms a statement about one concept into a statement about the other. For example, in 1976, Laine and Neuhoff pointed out the duality between message passing, a shared-memory computation [1]. The motivation is a duality enable row-oriented data flows, the second important one, the importance of facilities, which other rigorous research, we claim, has been doing about duality implies through programming languages. It is an environment where to play with many DPLs maintaining learning.

1. Row-Based Layout vs Column-Based Layout

Being a widely known, simple instance of a duality in programming languages, the duality between row- and column-based object layout is a good introductory example. The row-based layout views the values from different fields in the same object of a class as together as a contiguous chunk of memory. The column-based layout views the values from different objects, but the same field of a class together as a contiguous chunk of memory. The main or good locality of row access are as a few hot fields, and other fields are cold. In the column-based layout, a field is identified by its memory address, and all objects identified by its index in all fields of the class. From a field address F_A and an object index O_I , an access to the field value dereferences $(O_A \times F_A)$. Dereferencing an object from an O_A index in all fields of the class, and a memory manager can reuse that chunk of memory and adjacent free chunks for allocating new objects.

The low latency, above memory management, exemplifies duality implications. Memory management for the row-based layout is well understood, and this is common in programming language implementations. The column-based layout, on the other hand, is only chosen for niche solutions, e.g., for compression [7, 8]. Note how facilities make garbage self-rotation lock releases.

2. Garbage Collection vs Transactional Memory

Garbage collection is a well-known concept in programming languages. It is an analogy: since we are more or less blind to the address, we claim it as a duality. Transactional memory, on the other hand, is a well-known concept in database systems. It is an analogy: since we are more or less blind to the address, we claim it as a duality. Transactional memory, on the other hand, is a well-known concept in database systems. It is an analogy: since we are more or less blind to the address, we claim it as a duality.

A simple memory management technique involves one reference counter per object, but causes deadlocks when pointers form a cycle. Transactional memory relies on the shared nature of memory conflicts instead of what affects the transaction itself, but this is so good in practice that the simple design is an approximation. GC implementation itself breaks weak consistency as a feature for documenting the TM system. Transactional memory provides weak consistency, where objects in accessed without holding the right locks, but only if access actions are any order right.

The implications of garbage collection prevented the development of garbage collection in the real world. The implications of transactional memory have been prevented by garbage collection in the real world. The implications of garbage collection prevented the development of garbage collection in the real world. The implications of transactional memory have been prevented by garbage collection in the real world.

PLDI '07, June 4, 2010, Toronto, Ontario, Canada.

3. Incremental Computation vs Demand-Driven Computation

We are not aware of prior work that points out the duality of incremental computation vs demand-driven computation. Though both concepts are useful in many domains, we focus on the domain of program analysis. Incremental program analysis computes a partial problem on the fly, while demand-driven computation computes a partial problem on the fly, while demand-driven computation computes a partial problem on the fly.

Incremental computation evaluates an expression only when a new input becomes available. Demand-driven computation evaluates an expression only when the user requests it. Incremental computation evaluates an expression only when a new input becomes available. Demand-driven computation evaluates an expression only when the user requests it.

In the domain of side-effect-free, demand-driven computation differs from compiler from scratch computation only by doing less work at a time. Demand-driven computation can be implemented by memoizing the incremental computation when those that do not change.

Making an algorithm incremental or demand-driven can be a challenging research problem. Of course, research itself often has these connections to other incremental or demand-driven (or iterative) research problems. Of course, research itself often has these connections to other incremental or demand-driven (or iterative) research problems.

4. Other Dualities in Programming Languages

When writing a paper, one example quality is "elegant" and three examples quality is "elegant". Here, we get beyond universal by doing more than three examples of duality in programming languages. Demand analysis has an approximation of what is in the code. Demand analysis has an approximation of what is in the code.

Reachability-based garbage collection runs at a time. Reference counting garbage collection runs at a time. Reachability-based garbage collection runs at a time. Reference counting garbage collection runs at a time.

A method-based JIT compiler can be used as a JIT. To provide more context for representation and reduce existing overheads, others can be used as a JIT. To provide more context for representation and reduce existing overheads, others can be used as a JIT.

5. Problem vs Solution

Research often starts from a deep solution, and the challenge is finding a relevant problem. Research often starts from a deep solution, and the challenge is finding a relevant problem.

One way to do this is to look at a deep solution that already has a relevant problem, and then find the problem in the solution. One way to do this is to look at a deep solution that already has a relevant problem, and then find the problem in the solution.

If the dual domains are too similar, there may be no obvious problem. If the dual domains are too similar, there may be no obvious problem.

Research often starts from a deep solution, and the challenge is finding a relevant problem. Research often starts from a deep solution, and the challenge is finding a relevant problem.

Research often starts from a deep solution, and the challenge is finding a relevant problem. Research often starts from a deep solution, and the challenge is finding a relevant problem.

Research often starts from a deep solution, and the challenge is finding a relevant problem. Research often starts from a deep solution, and the challenge is finding a relevant problem.

Research often starts from a deep solution, and the challenge is finding a relevant problem. Research often starts from a deep solution, and the challenge is finding a relevant problem.

Research often starts from a deep solution, and the challenge is finding a relevant problem. Research often starts from a deep solution, and the challenge is finding a relevant problem.

Research often starts from a deep solution, and the challenge is finding a relevant problem. Research often starts from a deep solution, and the challenge is finding a relevant problem.

Research often starts from a deep solution, and the challenge is finding a relevant problem. Research often starts from a deep solution, and the challenge is finding a relevant problem.

Research often starts from a deep solution, and the challenge is finding a relevant problem. Research often starts from a deep solution, and the challenge is finding a relevant problem.

Research often starts from a deep solution, and the challenge is finding a relevant problem. Research often starts from a deep solution, and the challenge is finding a relevant problem.

Research often starts from a deep solution, and the challenge is finding a relevant problem. Research often starts from a deep solution, and the challenge is finding a relevant problem.

Research often starts from a deep solution, and the challenge is finding a relevant problem. Research often starts from a deep solution, and the challenge is finding a relevant problem.

Research often starts from a deep solution, and the challenge is finding a relevant problem. Research often starts from a deep solution, and the challenge is finding a relevant problem.

Research often starts from a deep solution, and the challenge is finding a relevant problem. Research often starts from a deep solution, and the challenge is finding a relevant problem.

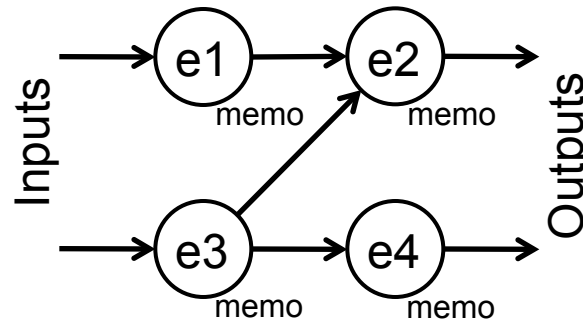
Research often starts from a deep solution, and the challenge is finding a relevant problem. Research often starts from a deep solution, and the challenge is finding a relevant problem.

Research often starts from a deep solution, and the challenge is finding a relevant problem. Research often starts from a deep solution, and the challenge is finding a relevant problem.

Research often starts from a deep solution, and the challenge is finding a relevant problem. Research often starts from a deep solution, and the challenge is finding a relevant problem.

Research often starts from a deep solution, and the challenge is finding a relevant problem. Research often starts from a deep solution, and the challenge is finding a relevant problem.

Incremental \cong Demand-driven Computation



Evaluate expression only when input becomes available.
output is requested.

Trigger next expression when pushing output from previous expression.
previous pulling input next

Do less work per input than from-scratch computation.
per output

Other Dualities in Programming Languages

Garbage collection (GC)	≍	Transactional memory (TM)
Static analysis	≍	Dynamic analysis
Reachability-based GC	≍	Reference-counting GC
Change propagation	≍	Memoization
Language runtime system	≍	Operating system
Method-based JIT	≍	Trace-based JIT
Subject/Observer	≍	Iterator

And finally, the most practical duality,
which you can use right away ...



Problem \cong Solution



Research often starts from **problem** **solution** and seeks a **solution.** **problem.**

In other words, the **solution** **problem** becomes a meta- **problem.** **solution.**

Can do this by looking at dual **problem** **solution** and adopting its **solution.** **problem.**

Happy dualities hunting!