





COMPILING TREES COMBINING DATA LAYOUTS AND THE POLYHEDRAL MODEL



Paul IANNETTA (ENS de Lyon, Inria & LIP) paul.iannetta@ens-lyon.fr

Jury:

CHARLES Henri-Pierre, CLAUSS Philippe, COLLANGE Caroline, KELLER Gabriele,

Rapporteur. Rapporteur. Examinatrice. Examinatrice. GONNORDLaure,RADANNEGabriel,MORELLionel,MEISTERBenoît,

Directrice de thèse. Co-encadrant. Co-encadrant. Guest

What is Compilation?

- Programmers want to write easy to maintain software
- Offload complexity to the compiler (Hardware diversity)
- We don't want that people write hand tuned code.
- ullet ... the compiler has to be smart enough to use the hardware cleverly enough
- let's talk a bit about the hardware...



Are Algebraic Data Types Regular Enough?

How to compile ADTs reshaping efficiently?

Conclusion 000

The Basic Blocks Around a CPU



- Memory contains everything (slow 50 \sim 100ns)
- Cache contains only the bare minimum (fast $0.5 \sim 7 ns$)

Programs Which Play Well With The Hardware

- Programs whose behavior can be accurately predicted (control flow)
 - The CPU always knows what to do

Programs Which Play Well With The Hardware

- Programs whose behavior can be accurately predicted (control flow)
 - The CPU always knows what to do
- Programs with predictable read/write patterns (io patterns)
 - Prefetch and fill the cache from memory
 - Avoid waiting for data

Programs Which Play Well With The Hardware

- Programs whose behavior can be accurately predicted (control flow)
 - The CPU always knows what to do
- Programs with predictable read/write patterns (io patterns)
 - Prefetch and fill the cache from memory
 - Avoid waiting for data
- Programs with independant parts (parallelism opportunities)
 - Share the work between CPUs
 - Better use of system ressources

Introduction	
000000000	

Are Algebraic Data Types Regular Enough?

How to compile ADTs reshaping efficiently?

Conclusion 000

Are Algebraic Data Types Regular Enough?

How to compile ADTs reshaping efficiently?

Conclusion 000







Are Algebraic Data Types Regular Enough?

How to compile ADTs reshaping efficiently?

Conclusion 000







Are Algebraic Data Types Regular Enough?

How to compile ADTs reshaping efficiently?

Conclusion 000







Are Algebraic Data Types Regular Enough?

How to compile ADTs reshaping efficiently?

Conclusion 000







Are Algebraic Data Types Regular Enough?

How to compile ADTs reshaping efficiently?

Conclusion 000



 Statement \leftarrow Dependence



Are Algebraic Data Types Regular Enough?

How to compile ADTs reshaping efficiently?

Conclusion 000







Are Algebraic Data Types Regular Enough?

How to compile ADTs reshaping efficiently?

Conclusion 000



Are Algebraic Data Types Regular Enough?

How to compile ADTs reshaping efficiently?

Conclusion



Are Algebraic Data Types Regular Enough?

How to compile ADTs reshaping efficiently?

Conclusion



Are Algebraic Data Types Regular Enough?

How to compile ADTs reshaping efficiently?

Conclusion



Are Algebraic Data Types Regular Enough?

How to compile ADTs reshaping efficiently?

Conclusion



Are Algebraic Data Types Regular Enough?

How to compile ADTs reshaping efficiently?

Conclusion



Are Algebraic Data Types Regular Enough?

How to compile ADTs reshaping efficiently?

Conclusion



Are Algebraic Data Types Regular Enough?

How to compile ADTs reshaping efficiently?

Conclusion



Are Algebraic Data Types Regular Enough?

How to compile ADTs reshaping efficiently?

Conclusion 000

What About the Real World?

- Linear Algebra kernels
- Simulations
- Machine learning
- Image Processing

• ...





The Polyhedral Model

The Polyhedral Model (French Success Story since 1991!)

Compiles programs with affine loops ranging over arrays accessed with affine functions into efficient C code by: 1) Analysing dependencies, 2) Reschedule the instructions, 3) Generate C code.

Pros:

- The best tool when it comes to polyhedral programs
- 30-year of active research and improvements
- Available in productions compilers (gcc, llvm)

Cons:

• Still limited to arrays

Are Algebraic Data Types Regular Enough?

How to compile ADTs reshaping efficiently?

Conclusion 000

Still Limited to Arrays, the Whys?

ŝ

 $\sqrt[n]{}$

Properties	Array	Pointers
Contiguous in Memory	✓	×
Random Access	~	×



Are Algebraic Data Types Regular Enough?

How to compile ADTs reshaping efficiently?

Conclusion 000

Still Limited to Arrays, the Whys?

ß

 \sum^{2}

Properties	Array	Pointers
Contiguous in Memory	~	×
Random Access	✓	×

Other nice data structures?



Introduction	
0000000	

Research Questions

- Are Algebraic Data Types Regular Enough?
- How to Compile ADTs reshaping efficiently?

ARE ALGEBRAIC DATA TYPES REGULAR ENOUGH?

Introduction	

type suit = Heart Diamond Clubs Spade type value = Ace King Queen Jack Number of int type card = (suit, value) A simple ADT



Introduction 00000000	Are Algebraic Data Types Regular Enough? ○●●○○○○○○○○○○○	How to compile AD
type e x	pr =	

Conclusion





(* 1 + 2 * 3 = 7 *)
eval (Add(Cst(1), Mul(Cst(2), Cst(3)))

The expr ADT and its eval function

This is "regular", isn't it?

Conclusion 000

How Are Encoded ADTs Into Memory?

• The constant part (enumeration with arguments) has a fixed size

Conclusion 000

How Are Encoded ADTs Into Memory?

- The constant part (enumeration with arguments) has a fixed size
- The recursion is captured by using pointers

```
struct tree {
    int val;
    struct tree *left, *right;
};
```



Conclusion 000

How Are Encoded ADTs Into Memory?

- The constant part (enumeration with arguments) has a fixed size
- The recursion is captured by using pointers
- But pointers = no predictability

Conclusion 000

How Are Encoded ADTs Into Memory?

- The constant part (enumeration with arguments) has a fixed size
- The recursion is captured by using pointers
- But pointers = no predictability

What about exploiting the regular properties of ADTs... ...and store them into arrays!

Are Algebraic Data Types Regular Enough?

How to compile ADTs reshaping efficiently?

Conclusion

Breadth-first Layout and Its Properties



Are Algebraic Data Types Regular Enough?

How to compile ADTs reshaping efficiently?

2 3

1

6 7

Conclusion 000

Breadth-first Layout and Its Properties



•	А	notion	of	lavers
•		notion	UI.	layers

Are Algebraic Data Types Regular Enough?

How to compile ADTs reshaping efficiently?

Conclusion 000

Breadth-first Layout and Its Properties



- A notion of layers
- A simple formula for finding parent and children (it's a shift away)
Are Algebraic Data Types Regular Enough?

How to compile ADTs reshaping efficiently?

Conclusion 000

Breadth-first Layout and Its Properties



- A notion of layers
- A simple formula for finding parent and children (it's a shift away)
- The tree can be compressed, if only values matter.

Are Algebraic Data Types Regular Enough?

How to compile ADTs reshaping efficiently?

Conclusion 000

Breadth-first Layout and Its Properties



- A notion of layers
- A simple formula for finding parent and children (it's a shift away)
- The tree can be compressed, if only values matter.
- Insertion happens at the end, or in a hole

Introduction
00000000

The Layer Tower



Figure: A graphical notation

Other Layouts

- Variation without holes
- Obviously, depth-first layout
- Cache oblivious community has designed many byzantine layouts
- Most famous: Van Em Boas Layout (see manuscript)

Other Layouts

- Variation without holes
- Obviously, depth-first layout
- Cache oblivious community has designed many byzantine layouts
- Most famous: Van Em Boas Layout (see manuscript)

The rest of the presentation uses the breadth-first layout

A Data-Structure to Study

We want a data-structure:

- which is actually used
- which would benefit from optimizations
- which is neither to complex nor to simple

A Data-Structure to Study

We want a data-structure:

- which is actually used
- which would benefit from optimizations
- which is neither to complex nor to simple

We chose AVL Trees



 Rotations (Structural Modifications)



5

AVL Implementation

AVL with pointers	Implicit layout (Tarbres)
Rotations are easy and cheap $\mathcal{O}(1)$	Rotations in place?
Bad cache behavior (find/map)	Nice cache behavior (find/map)

How to compile ADTs reshaping efficiently?

Conclusion 000

AVL Rebalancing: Low Level Operations 1/2





Figure: Subtree shifts

Are Algebraic Data Types Regular Enough?

How to compile ADTs reshaping efficiently?

Conclusion 000

AVL Rebalancing: Low Level Operations 1/2



Figure: Subtree pull-ups and pull-downs

Are Algebraic Data Types Regular Enough?

How to compile ADTs reshaping efficiently?

Conclusion 000

AVL Rebalancing: Hands on!



• Original Unbalanced State

Are Algebraic Data Types Regular Enough?

How to compile ADTs reshaping efficiently?

Conclusion 000



- Original Unbalanced State
- \bullet Pull Down (Left) ${\cal A}$

Are Algebraic Data Types Regular Enough?

How to compile ADTs reshaping efficiently?

Conclusion 000



- Original Unbalanced State
- Pull Down (Left) ${\cal A}$
- \bullet Shift Left ${\cal B}$

Conclusion 000



- Original Unbalanced State
- Pull Down (Left) ${\cal A}$
- \bullet Shift Left ${\cal B}$
- Rename

Are Algebraic Data Types Regular Enough?

How to compile ADTs reshaping efficiently?

Conclusion 000



- Original Unbalanced State
- Pull Down (Left) ${\cal A}$
- \bullet Shift Left ${\cal B}$
- Rename
- Move up $(z-\mathcal{B}-\mathcal{C})$

Implementation

- Support for insert/delete/find/map
- Rotation as a combination of low-level operations
- Support for scalar-type and pointer-type values
- The low-level operations can be:
 - Naive: move one layer at a time
 - Pipelined: try to start the next move as soon as possible
 - Parallelized
- Available at: https://gitlab.inria.fr/paiannet/calv/





Figure: Subtree shifts

Are Algebraic Data Types Regular Enough?

How to compile ADTs reshaping efficiently?

Conclusion 000

Benchmark on a Key-Value Store Scenario ($N = 1M \approx 2^{20}$)

- A benchmark adapted from the database community
- Scenario:
 - Emulate a key-value database
 - Check the performance of pointers vs arrays.



Are Algebraic Data Types Regular Enough?

How to compile ADTs reshaping efficiently?

Conclusion 000

Benchmark on a Key-Value Store Scenario (N = 1M $\approx 2^{20}$)



Summary of the Contribution

- Implemented as a small library (calv)
- We tried to exploit the parallelism of the low-level operations
- It works!

Summary of the Contribution

- Implemented as a small library (calv)
- We tried to exploit the parallelism of the low-level operations
- It works!
- Hand Tuned

Summary of the Contribution

- Implemented as a small library (calv)
- We tried to exploit the parallelism of the low-level operations
- It works!
- Hand Tuned
- Improve parallelism
 - Improve granularity of low level operations

Conclusion 000

Summary of the Contribution

- Implemented as a small library (calv)
- We tried to exploit the parallelism of the low-level operations
- It works!
- Hand Tuned
- Improve parallelism
 - Improve granularity of low level operations

Let's try to automate!

How to compile ADTs RESHAPING EFFICIENTLY?

Our Proposition: REW (A DSL to Compile Structural Transformations)

 Rew is small DSL to:

- Algebraic Data Types without sharing
- describe structural transformations on those through pattern matching.

type tree = Empty | Node (tree,int,tree)

```
pull_up (t : tree) : tree = rewrite t {
    | Node(a,i,Node(b,j,c)) -> Node(b,j,c)
    | Node(a,i,Empty) -> Empty
    | Empty -> Empty
}
```

Figure: What REW looks like

Compilation Steps

- Give the specifications (describe the types & the transformations)
- Ø Figure out the movements
- Occupie the dependences between the movements
- Reschedule the movements
- Generate the code

Conclusion



Are Algebraic Data Types Regular Enough?

Conclusion



Are Algebraic Data Types Regular Enough?

Conclusion 000



Are Algebraic Data Types Regular Enough?

Conclusion 000



Introduc	
00000	000

How to compile ADTs reshaping efficiently?

Conclusion



Introduction
00000000

How to compile ADTs reshaping efficiently?

Conclusion



Introduction
00000000

How to compile ADTs reshaping efficiently?

Conclusion 000



Introduction
00000000

How to compile ADTs reshaping efficiently?

Conclusion

Our Notations for Paths



 $.0.0 \equiv (.0)^2$

Introduction
00000000

How to compile ADTs reshaping efficiently?

Conclusion

Our Notations for Paths





 $(.0)^2.1$

Introduction
00000000

How to compile ADTs reshaping efficiently?

Conclusion





.2.1
Introduction
00000000

How to compile ADTs reshaping efficiently?

Conclusion

Our Notations for Paths





Introduction
00000000

How to compile ADTs reshaping efficiently?

Conclusion

Our Notations for Paths





How to compile ADTs reshaping efficiently?

Conclusion



How to compile ADTs reshaping efficiently?

Conclusion



How to compile ADTs reshaping efficiently?

Conclusion



How to compile ADTs reshaping efficiently?

Conclusion 000



How to compile ADTs reshaping efficiently?

Conclusion 000



How to compile ADTs reshaping efficiently?

Conclusion



How to compile ADTs reshaping efficiently?

Conclusion 000

Use Layers to Subdivide Memory Movements



Are Algebraic Data Types Regular Enough?

How to compile ADTs reshaping efficiently?

Conclusion



Are Algebraic Data Types Regular Enough?

Conclusion



Are Algebraic Data Types Regular Enough?

Conclusion



Are Algebraic Data Types Regular Enough?

Conclusion



Are Algebraic Data Types Regular Enough?

How to compile ADTs reshaping efficiently?

Conclusion

Conclusion



Are Algebraic Data Types Regular Enough?

Conclusion





Are Algebraic Data Types Regular Enough?

How to compile ADTs reshaping efficiently?

Conclusion

Step 3: Characterizing Memory Movements



Are Algebraic Data Types Regular Enough?

How to compile ADTs reshaping efficiently?

Conclusion

Step 3: Characterizing Memory Movements



Are Algebraic Data Types Regular Enough?

Conclusion

Step 3: Characterizing Memory Movements



Are Algebraic Data Types Regular Enough?

Conclusion

Step 3: Characterizing Memory Movements



Mon. 2 May 2022

Are Algebraic Data Types Regular Enough?

How to compile ADTs reshaping efficiently?

Conclusion 000

Polyhedral Like Dependencies: Intuition



Let's take the rule
$$(c_0)$$
: $(1.2, 2.2^{k_2}, 1 \rightarrow .2.2^{k_2}, 1)$

Introduction 00000000

How to compile ADTs reshaping efficiently?

Conclusion 000

Polyhedral Like Dependencies: Intuition



Let's take the rule (c_0) : $(.2, 2, 2^{k_2}, 1 \rightarrow .2, 2^{k_2}, 1)$

• When $k_2 = 1$, (c_0) : ((.2)³.1 \rightarrow (.2)².1)

• When
$$k_2 = 2$$
, (c_0) : $((.2)^4.1 \rightarrow (.2)^3.1)$

Introduction 00000000

How to compile ADTs reshaping efficiently?

Conclusion 000

Polyhedral Like Dependencies: Intuition



Let's take the rule (c_0) : $(.2.2.2^{k_2}.1 \rightarrow .2.2^{k_2}.1)$

• When
$$k_2 = 1$$
, (c_0) : ((.2)³.1 \rightarrow (.2)².1)

• When
$$k_2 = 2$$
, (c_0) : $((.2)^4 \cdot 1 \to (.2)^3 \cdot 1)$

Introduction 00000000

How to compile ADTs reshaping efficiently?

Conclusion 000

Polyhedral Like Dependencies: Intuition



Let's take the rule (c_0) : $(.2, 2, 2^{k_2}, 1 \rightarrow .2, 2^{k_2}, 1)$

- When $k_2 = 1$, (c_0) : ((.2)³.1 \rightarrow (.2)².1)
- When $k_2=$ 2, (c_0) : $((.2)^4.1 \rightarrow (.2)^3.1)$
- More generally, $(c_0)(k)$ and $(c_0)(k+1)$ conflict.

Conclusion

Step 5: Polyhedral Model Based Scheduling



Conclusion



Conclusion



How to compile ADTs reshaping efficiently?

Conclusion 000



How to compile ADTs reshaping efficiently?

Conclusion 000



Conclusion



Conclusion



How to compile ADTs reshaping efficiently?

Conclusion 000



Conflict analysis: The Details

Theorem

Let $(s_1 \to d_1)$ and $(s_2 \to d_2)$ be two moves (not necessarily distinct). If we consider d_1 and s_2 as regular expressions. Then,

 $\mathcal{L}(d_1) \cap \mathcal{L}(s_2)$ is the conflict space and it can be characterized by affine inequalities.

Conflict analysis: Compute the dependences for all tuples of movements

How to compile ADTs reshaping efficiently?

Summary of the Contribution

- Promising technique to compile pattern matching-based structural transformations
- Extension of the polyhedral techniques to a different class of regular programs

Summary of the Contribution

- Promising technique to compile pattern matching-based structural transformations
- Extension of the polyhedral techniques to a different class of regular programs

Front end for ADTs which plugs into the polyhedral model

How to compile ADTs reshaping efficiently?

Conclusion & Contributions List

- Are Algebraic Data Types Regular Enough?
 - Yes, but the layout is important
 - Parallelization opportunities
 - Implemented as a small library
 - Better exploited through automation
Conclusion & Contributions List

- Are Algebraic Data Types Regular Enough?
 - Yes, but the layout is important
 - Parallelization opportunities
 - Implemented as a small library
 - Better exploited through automation
- How to Compile ADTs reshaping efficiently?
 - A framework to compile structural transformation
 - Reuse polyhedral ideas
 - Partial implementation by Gabriel

How to compile ADTs reshaping efficiently?

Conclusion & Contributions List

- Are Algebraic Data Types Regular Enough?
 - Yes, but the layout is important
 - Parallelization opportunities
 - Implemented as a small library
 - Better exploited through automation
- How to Compile ADTs reshaping efficiently?
 - A framework to compile structural transformation
 - Reuse polyhedral ideas
 - Partial implementation by Gabriel
- Another Model for the Input Language of the Polyhedral Model?
 - A framework to extend the polyhedral model
 - Based on a modified operational semantics
 - Enabling future non-exact extensions

Future Work

- \bullet Extend REW:
 - Add recursion, guards
 - Add support for functions
- Improve the compilation of REW
 - Generate parallel code
 - Improve parallel cache performance

• Continue to explore reformulation of the polyhedral model input language

THANK YOU FOR YOUR ATTENTION



QUESTIONS?