Composable Rewrites for Layouts of Algebraic Data Types

CHRISTOPHE GYURGYIK, Stanford University, USA
ALEXANDER J ROOT, Stanford University, USA
BENJAMIN DRISCOLL, Stanford University, USA
GABRIEL RADANNE, ENS Lyon, France
FREDRIK KJOLSTAD, Stanford University, USA
CLÉMENT PIT-CLAUDEL, EPFL, Switzerland

Algebraic data types (ADTs) decouple the logical structure of data from physical layout choices that affect performance. For example, packing small fields into a word improves spatial locality, while storing co-accessed fields contiguously can enable vectorization. Such choices reinforce or undermine one another, with ramifications for compiler optimizations and memory management. Their interplay creates a combinatorial space of layouts with different runtime costs and memory footprints. The most favorable layout depends on the program's construction and destruction patterns, the characteristics of input data, and the target machine.

Exploring this space requires control over layout decisions and a way to compose them while preserving program semantics. We provide both through rewrite rules over layouts, independent of any particular source language or bit-level encoding. Starting from a canonical layout derived from the ADT, these rules transform its representation while leaving the source program intact. We identify a local refinement criterion for proving each rewrite rule correct, so that layouts continue to represent every value and keep distinct values apart. Adding a new rule therefore only requires a proof satisfying this local criterion.

We realize these rules in Isomer, a new pure functional programming language whose compiler specializes construction, pattern matching, and memory management to the chosen layouts. Using a default layout for all ADTs, Isomer achieves competitive performance against established functional compilers that automatically select layouts. Our rewrites further improve performance by tailoring layouts to each workload, permitting representations that prior work cannot attain or can only achieve through explicit changes to the source code. Finally, Isomer can replicate the layouts and peak memory of a language with state-of-the-art reference counting, and additional rewrites further reduce peak memory.

1 Introduction

Algebraic data types (ADTs) describe data by composing products and sums: a product type holds a value of each of its component types, and a sum type holds a value of exactly one of its alternatives. This abstraction supports type-safe construction and pattern matching independently of any physical representation, and thus the same logical values can be represented using many different memory layouts.

The choice among these layouts can substantially affect runtime performance and memory footprint, motivating extensive research on ADT layout optimizations in functional programming languages [6, 14, 23–25, 28, 33, 50, 61, 70, 78]. For example, packing small fields into a word can save space but may require additional instructions or unaligned loads when they are read, while boxing a large, rarely accessed field places it behind a pointer, keeping frequently accessed data compact at the cost of a pointer dereference for infrequently accessed data. These choices are interdependent; the benefits of boxing depend in turn on how the remaining fields are packed and ordered.

Furthermore, layouts are necessarily entangled with the language implementation. Field widths and alignment affect instruction selection and vectorization [7, 28, 96, 97], while the sizes of allocated objects constrain opportunities for memory reuse [56, 57, 73]. A layout therefore cannot

Authors' Contact Information: Christophe Gyurgyik, Stanford University, Stanford, USA, cpg@cs.stanford.edu; Alexander J Root, Stanford University, Stanford, USA, ajroot@cs.stanford.edu; Benjamin Driscoll, Stanford University, Stanford, USA, bdrisc@cs.stanford.edu; Gabriel Radanne, ENS Lyon, Lyon, France, gabriel.radanne@inria.fr; Fredrik Kjolstad, Stanford University, Stanford, USA, kjolstad@cs.stanford.edu; Clément Pit-Claudel, EPFL, Lausanne, Switzerland, clement.pit-claudel@epfl.ch.

be evaluated in isolation: its performance depends on which optimizations it enables and which of those optimizations the compiler applies. More broadly, a program's construction and destruction patterns, characteristics of its input data, and the target architecture shape the Pareto frontier between runtime performance and memory footprint, leaving no universally best layout [28].

Language implementations presently address these costs through both automatic layout selection and explicit representation control. MLton and Morphic select representations through compiler analyses [9, 10, 95]. Haskell and OCaml expose representation choices through type system extensions [11, 24, 51, 70]. More recently, layout description languages make the separation of data type and data layout explicit and programmable [6, 14, 28].

We build on the decoupled approach of prior work [6, 14, 28], introducing composable rewrite rules over a layout language for deriving ADT representations independently of a particular source language or bit-level encoding. Rather than writing each candidate layout in full, as required by layout description languages [6, 14, 28], a programmer or an automated search procedure starts from a canonical layout induced by the source type, and transforms it through incremental rewrites. A layout is *faithful* when every source value has a representation and no representation corresponds to two distinct values. Our key insight is that a local criterion is enough to preserve faithfulness: each rewrite rule need only satisfy it on the part of the layout it changes. For a core calculus, we derive construction and destruction for rewritten layouts and prove that valid rewrite sequences preserve the result of terminating programs. Adding a new rewrite rule requires only a reconstruction function that establishes the local correctness criterion.

We implement these rewrite rules in Isomer, a pure functional programming language and compiler that embodies a separation of concerns: changing a layout never requires modifying the program. The compiler specializes construction, destruction, and memory management (reference counting) to the rewritten layouts. To support different layouts across program phases, Isomer also provides a `coerce` scheduling directive, which inserts a change of representation at a selected variable binding while leaving the source code unchanged. Our primary contributions are:

- Composable layout rewrite rules that, independent of a particular source language or machine, express a broad range of layouts and support extension with new rules.
- A formal model of layouts that separates structural representation from any particular bit-level encoding. Per-rule proofs compose to show that valid rewrites keep layouts faithful and preserve the scalar results of terminating programs.
- Isomer, a new pure functional programming language that decouples layout rewrites from the source program, and a compiler that automatically specializes with respect to rewritten layouts. Isomer additionally supports transformations of ADT layouts at arbitrary program points.

Our evaluation demonstrates that Isomer is competitive with state-of-the-art functional languages, and our decoupled abstraction incurs no performance overhead. In three case studies, we examine composing rewrites, memoizing compute without editing the source code, and changing layouts between program phases. With an automatically selected layout, Isomer achieves geomean speedups of 1.42× over Morphic [9] and 1.53× over MLton [95], whose compilers automatically select representations. We also perform layout design space exploration, where the fastest measured layout per benchmark gives a 1.27× (geomean) speedup over the default, and the layout with the lowest peak memory reduces memory by 23%. We compare Isomer with Ribbit [6], a layout description language for general-purpose ADTs and the closest prior work. Isomer's code generation yields a 1.58× (geomean) speedup when matching Ribbit's layouts. Lastly, we show layout control complements automatic memory management: with equivalent layouts, Isomer uses at most the peak memory of Koka's fully-in-place programs [57], and further rewrites lower peak memory by an additional 20%.

Finally, we clarify the scope of this work. First, we assume finite trees and monomorphization in our implementation. Although monomorphization typically bars features such as polymorphic recursion and separate compilation, runtime specialization [41] could relax this restriction. Second, we focus on data representation, which is orthogonal to the execution strategies that user-schedulable languages decouple from algorithms [30, 32, 71]. Combining the two in a general-purpose functional language would be a natural direction for future work. Third, our formalism covers a formal model of layouts, program preservation under layout rewrites, and recovery from physical encodings under explicit conditions. It does not establish end-to-end correctness of the Isomer compiler.

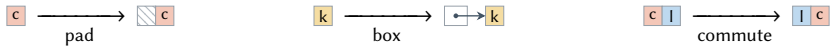
2 Optimizing Layouts of Red-Black Trees

Consider a programmer implementing an ordered set as a red-black tree [65]:

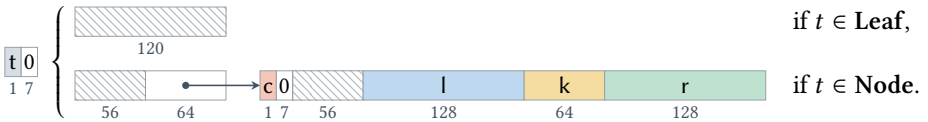
$$\text{Color} = \textbf{Red} \mid \textbf{Black} \qquad \text{Tree} = \textbf{Leaf} \mid \textbf{Node}(\mathsf{c} : \text{Color},\ \mathsf{l} : \text{Tree},\ \mathsf{k} : \mathsf{u}_{64},\ \mathsf{r} : \text{Tree})$$

Throughout this section, the type and source code stay fixed while the physical layout changes. The layout of Tree is defined recursively and each child field, l and r, uses the same layout. We assume little-endian byte order, 64-bit pointers and 8-byte aligned allocations. Our diagrams place storage offsets in increasing order from left to right, with each field's least significant bit (LSB) on the left. Adjacency ( ) represents contiguous memory, braces ( ) distinguish arms discriminated by a value t, arrows ( ) lead from pointers to heap cells, and zero bits ( ) are marked separately from padding ( ), whose contents are undefined.

A *rewrite* makes one elementary change to a layout, e.g., inserting unused bits before a field (pad), moving a field into a heap cell (box), or swapping two adjacent fields (commute):

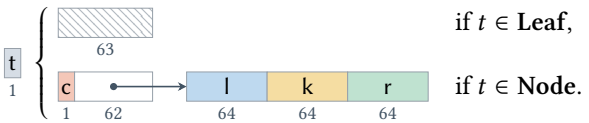


Each layout below is realized with a sequence of such rewrites. We start with the layout that Morphic [9, 10] automatically selects for this type and explore alternative layouts. This initial layout follows a C-style tagged union with natural alignment, in which a tree value is a byte tag (or discriminant) beside a pointer to a separately allocated heap cell. Fields are aligned, and those that are less than one word are padded to a full word (8 bytes). A tree value therefore occupies 16 bytes, and the node's fields occupy 48 bytes, including both 16-byte children.



(1)

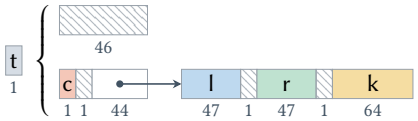
While this layout permits aligned memory loads, nearly half of the bits in the **Node**'s fields are zero or padding. Systems programmers reclaim such bits by hand. The Linux kernel's implementation of a red-black tree, for example, stores a **Node**'s color in an unused low bit of its parent pointer [87]. This decision also exposes clever instruction selection. Red is encoded as zero, so a red **Node**'s parent field stores exactly its parent's address, and tree insertion can dereference it without masking [52]. Similarly, a pointer to a black node can be effectively encoded as $p + 1$ where p is the pointer address, resulting in a single instruction, `lea`, on x86 [27]. The same opportunity exists in our tree: moving each child's tag into a low bit of its pointer shrinks each to one word. We also hoist the color of each **Node** into a second low bit of the pointer that references it, resulting in a 24-byte cell:



(2)

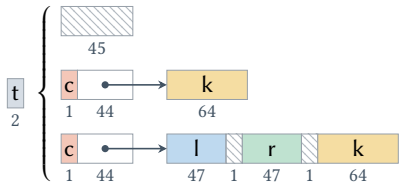
The layout applies recursively to both children: each occupies one word, consisting of a one-bit tag followed by either a color bit and a 62-bit pointer for a **Node** or 63 bits of padding for a **Leaf**. An additional benefit is that red-black tree rebalancing can test children's colors directly without first dereferencing a pointer.

But we can go further: by default, x86-64 Linux places user allocations below 2^{47} [84], leaving the upper 17 bits of their 64-bit addresses zero. An 8-byte-aligned address thus needs only bits $[3, 47)$. This enables a *double-ended* bit stealing technique [25], where a pointer is stored in just 44 bits. Consequently, a child now occupies only 47 bits of space, and some additional padding keeps each cell's fields at byte boundaries:



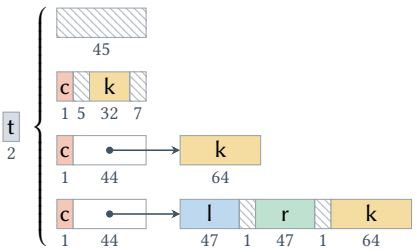
(3)

So far, these rewrites reclaim unused bits in an effort to reduce the memory footprint. The tag can *also* distinguish more interesting tree shapes, which lets a layout store such shapes more compactly. Consider a **Node** whose children are both **Leaf**, such as the one each insertion creates for its new key. Layout 3 still reserves 47 + 47 = 94 bits for these two children, even though they provide no additional information. We can instead give this shape an arm of its own, **Node**(**Leaf**, **Leaf**), similar to prior work in acceleration structure design [7]. Within this arm, an unfold rewrite expands each child's layout, and, since the layout of **Leaf** is empty, l and r do not need to be stored. With the color already in the reference, this arm's cell then holds only its key, in exactly 8 bytes. With three arms, the tag widens to two bits, and every other **Node**(•, •) shape falls back to the general **Node** layout. Its cell is unchanged, and the wider tag takes the place of the padding bit before its pointer:



(4)

Further rewrites can exploit properties of the data. Keys may be drawn from a much smaller distribution than the type admits. A refine rewrite gives the specialized **Node** whose keys are below 2^{32} an arm of their own, which stores the key in 32 bits. As the color already sits in the reference, an unbox rewrite can store the key there too, in place of the pointer, leaving this arm with no cell at all. The layout maintains the full u64 domain, but incurs a cost of checking whether the key fits in 32 bits when constructing a node:



(5)

Finally, layouts also interact with memory management. Rebalancing a tree consumes and rebuilds nodes, and the compiler can reuse the allocation of a consumed cell for a rebuilt node *if* they share the same size [73, 89]. In Layout 5, a **Node**(**Leaf**, **Leaf**) with a 64-bit key outgrows its 8-byte cell when an insertion replaces one of its **Leaf** children, forcing a fresh allocation. A pad

rewrite restores reuse by growing these cells to 20 bytes, which prior work on reference counting achieves by modifying the source program itself [57].

We express optimizations beyond this example, such as array inlining [96] and memoizing values that are expensive to recompute [67]. Each layout exploits a semantics-preserving transformation of the program, and the best combination depends on the workload in ways that fixed compiler heuristics cannot easily realize. A functional programmer *today* can express only a handful of these layouts, and even those often require manipulating types or the source program itself [57, 58, 70].

3 Overview

We now present a general framework for expressing and composing such layout transformations. Section 4 introduces a layout language that can describe each layout above, together with rewrite rules that each make a local change, e.g., swapping two fields or placing a field behind a pointer. Layout rewriting starts from a source type's canonical layout, which mirrors its type declaration.

Sections 5 to 7 develop the formalism. Section 5 defines the *structural representations* a layout permits for each source value, which describe how the layout organizes the value before it is encoded as bits. Section 6 proves each rule correct on the part of the layout it changes and lifts this guarantee through enclosing and recursive layouts, so every layout reached from the canonical one by valid rewrites is faithful, and terminating programs in a core calculus compute the same scalar results. Section 7 gives conditions under which a model encoder and decoder recover source values from memory. The development is mechanized in Lean[1] [64], and references identify Lean declarations corresponding to formal statements and grammars, e.g., Lean: `WF`. Appendix I indexes the complete set of definitions, theorems, and worked examples.

Section 8 describes the implementation, Isomer, a new functional language that uses our rewrites to derive data layouts. Its compiler specializes construction, destruction, and reference counting to the resulting layouts, and its `coerce` directive changes a value's layout between program phases that favor different layouts. Finally, we evaluate the runtime performance and memory footprint of Isomer programs in Section 9, through case studies and comparisons with state-of-the-art functional languages.

4 Layout Language

A *layout* describes how the values of a source type are represented (grammar in Figure 1). Before defining the individual constructs, we express Layout 3 from Section 2 in the layout language:

$$t : \mathsf{u}_1 \times \bigsqcup^t \{(0, \mathbf{Leaf}) \Rightarrow \mathsf{pad}_{46},\\ (1, \mathbf{Node}) \Rightarrow \mathsf{c} : \mathsf{ref\ Color} \times \mathsf{pad}_1 \times \&_{44}(\mathsf{l} : \mathsf{ref\ Tree} \times \mathsf{pad}_1 \times \mathsf{r} : \mathsf{ref\ Tree} \times \mathsf{pad}_1 \times \mathsf{k} : \mathsf{u}_{64})\}$$

Read from left to right, a tree stores a one-bit tag t, followed by the arm that t selects, code 0 for a **Leaf** and 1 for a **Node**. A **Leaf** stores 46 bits of padding, while a **Node** stores its color, a bit of padding, and a 44-bit pointer to its remaining fields. The color c refers to the layout of Color, which stores only a one-bit tag. Each child refers back to the layout of Tree, thus each occupies $1 + 1 + 1 + 44 = 47$ bits. The rest of this section defines these constructs and the rewrite rules that derive this layout from the one induced by the source type.

4.1 Values and Paths

Source values are unsigned scalars and finite constructor trees $\mathbf{K}(v_1, \ldots, v_n)$, with fields in the order originally given by the source type. We write $\mathcal{V}$ for all source values and $\mathcal{V}_\tau \subseteq \mathcal{V}$ for all source values of an algebraic data type τ. We call the source value that a layout represents its *subject*. A *path* π selects a subvalue, or a range of bits of a scalar, and write $v.\pi$ for the value that π

[1]The mechanization relies upon two axioms, propositional extensionality and quotient soundness.

$x, \mathsf{T}, \mathbf{K} \in \mathsf{Id}, \quad n, i, j, c \in \mathbb{N}$

$\ell ::= \epsilon \mid \pi : \ell$	empty, name	$\pi ::= \varepsilon \mid x.\pi \mid \pi[i, j)$	names and bit subsets
$\mid \mathsf{pad}_n \mid \mathsf{zero}_n \mid \mathsf{u}_n \mid \langle e \rangle$	leaves	$e ::= n \mid \pi \mid e \oplus e$	memoized expressions
$\mid \&_n(\ell)$	pointer	$\phi ::= \top \mid \bot$	provenance
$\mid \mathsf{ref}\ \mathsf{T}$	recursive reference	$\mid \pi \mapsto \mathbf{K} \mid \pi \mapsto c$	constructor, scalar
$\mid \ell_0 \times \ell_1$	product	$\mid \phi \cap \phi \mid \phi \cup \phi$	intersection, union
$\mid \bigcup\{\phi_j \Rightarrow \ell_j\}$	sum	$\Delta ::= \varnothing \mid \Delta, \mathsf{T} \mapsto \ell$	layout environment
$\mid \bigsqcup^{\pi}\{(c_j, \phi_j) \Rightarrow \ell_j\}$	disjoint sum		

Fig. 1. Grammar of the layout language. Lean: `Layout, Path, Expr, Prov, LayoutEnv`

selects from v, when it exists. The empty path ε selects v itself, the path x selects its field x, and $x[i, j)$ selects bits $[i, j)$ of that field, counting from the least significant bit. For example, in node $v = \mathbf{Node}(\mathbf{Black}, \mathbf{Leaf}, 5, \mathbf{Leaf}) \in \mathcal{V}_{\mathsf{Tree}}$, the path k selects the key, $v.\mathsf{k} = 5$, and $\mathsf{k}[0, 2)$ selects its two low bits, $v.\mathsf{k}[0, 2) = 1$. For simplicity, we assume that a field name shared by constructors always has the same type.

4.2 Provenances

A layout can exploit knowledge of a value's constructors to specialize its representation [7, 78]. Following Baudon et al. [6], we call this knowledge a *provenance*. We depart from their patterns by taking each provenance to be the set of values it denotes, so that provenances form a lattice. The provenance $\pi \mapsto \mathbf{K}$ denotes the source values whose subvalue at path π has constructor $\mathbf{K}$, and we abbreviate $\varepsilon \mapsto \mathbf{K}$ as $\mathbf{K}$. Since a path can be arbitrarily long, a provenance can test a subvalue at any depth, such as $\mathsf{l.r} \mapsto \mathbf{Leaf}$ for a node whose left child's right child is a $\mathbf{Leaf}$. A scalar guard $\pi \mapsto c$ similarly selects the values with scalar c at path π, e.g., Layout 5 guards its small keys by $\bigcup_{i=0}^{2^{32}-1} \mathsf{k} \mapsto i$. As sets of values, provenances form a lattice bounded by $\top = \mathcal{V}$ and $\bot = \varnothing$, with intersection as meet, union as join, and inclusion as order. Two provenances are equal when they denote the same values. Appendix A provides the complete depth-one lattice of the red-black Tree. Such provenances readily model patterns on ADTs, e.g.,

$$\mathbf{Node}(_, \mathbf{Node}(\mathbf{Black}, _, _, _), _, _) \;\equiv\; (\varepsilon \mapsto \mathbf{Node} \cap l \mapsto \mathbf{Node} \cap l.c \mapsto \mathbf{Black}).$$

Each layout must represent a set $A \subseteq \mathcal{V}$ of *admitted subjects*, the source values that can reach it. An arm has a condition ϕ, its *guard*, and a layout for the values in $A \cap \phi$, its *payload*. A sum offers several arms, whose guards must together cover A. We additionally allow overlapping guards as an explicit design decision, since some optimizations rely on this [7, 31].

4.3 Layouts

To let rewrites target individual layout choices, the language makes explicit how a value is stored. A named layout $\pi : \ell$ makes $v.\pi$ the subject of ℓ, and ϵ is the empty layout, which stores nothing. Leaves describe scalar storage: pad_n reserves n bits with unconstrained contents, and zero_n requires them to be zero. Layout 1, for instance, follows its one-bit tag with zero_7 to fill a byte. u_n stores its scalar subject in n bits. A *memoized expression* $\langle e \rangle$ stores the result of evaluating e on the subject, at the width of its resulting type. Expressions use integer literals, scalar path reads, and supplied binary operations, but can't use arbitrary layout components. A pointer $\&_n(\ell)$ places the representation described by ℓ behind an n-bit address. A *layout environment* Δ maps type names to layouts. An entry $\Delta(\mathsf{T}) = \ell$ is a *layout definition*, with body ℓ, and a reference $\mathsf{ref}\ \mathsf{T}$ stands for that body.

Products $\ell_0 \times \ell_1$ are binary and ordered; we omit the parentheses of right-associated products. Their components occupy increasing bit offsets from left to right, and offsets restart at zero in each pointee. The cell of Layout 3 places l at offset 0, r at 48, and k at 96. An ordinary sum $\bigcup\{\phi_j \Rightarrow \ell_j\}$ has arms with distinct guards and represents a subject by the payload of any arm whose guard it

$$
\begin{array}{rcl}
\ell & \xrightarrow{\mathsf{pad}_n} & \mathsf{pad}_n \times \ell \\
\ell & \xrightarrow{\mathsf{zero}_n} & \mathsf{zero}_n \times \ell \\
\ell & \xrightarrow{\mathsf{ptr}_n} & \&_n(\epsilon) \times \ell \\
\&_n(\ell_0) \times \ell_1 & \xrightarrow{\mathsf{box}} & \&_n(\ell_0 \times \ell_1) \\
\&_n(\ell) & \xrightarrow{\mathsf{unbox}} & \ell \\
\ell_0 \times \ell_1 & \xrightarrow{\mathsf{commute}} & \ell_1 \times \ell_0 \\
(\ell_0 \times \ell_1) \times \ell_2 & \xleftrightarrow{\mathsf{assoc}} & \ell_0 \times (\ell_1 \times \ell_2) \\
\epsilon \times \ell & \xleftrightarrow{\mathsf{unit}} & \ell \\
\pi : \mathsf{u}_n & \xrightarrow{\mathsf{split}_k} & \pi[0,k) : \mathsf{u}_k \times \pi[k,n) : \mathsf{u}_{n-k}
\end{array}
\qquad
\begin{array}{rcl}
\bigcup\{\phi_j \Rightarrow \ell_j\} & \xrightarrow{\mathsf{tag}} & t : \mathsf{u}_n \times \bigsqcup^t\{(j,\phi_j) \Rightarrow \ell_j\} \\
\ell_0 \times \bigcup\{\phi_j \Rightarrow \ell_j\} & \xrightarrow{\mathsf{distribute}} & \bigcup\{\phi_j \Rightarrow \ell_0 \times \ell_j\} \\
\bigcup\{\phi_j \Rightarrow \ell \times \ell_j\} & \xrightarrow{\mathsf{factor}} & \ell \times \bigcup\{\phi_j \Rightarrow \ell_j\} \\
\ell & \xrightarrow{\mathsf{refine}_\phi} & \bigcup\{\phi \Rightarrow \ell, \top \Rightarrow \ell\} \\
\bigcup\{\phi \Rightarrow \ell, \Omega\} & \xrightarrow{\mathsf{prune}} & \bigcup\{\Omega\} \\
\bigcup\{\phi \Rightarrow \ell\} & \xrightarrow{\mathsf{collapse}} & \ell \\
\bigcup\{\phi \Rightarrow \bigcup\{\psi_j \Rightarrow \ell_j\}\} & \xrightarrow{\mathsf{flatten}} & \bigcup\{\phi \cap \psi_j \Rightarrow \ell_j\} \\
\mathsf{ref}\,\mathsf{T} & \xrightarrow{\mathsf{unfold}} & \Delta(\mathsf{T}) \\
\ell & \xrightarrow{\mathsf{memoize}_e} & \langle e \rangle \times \ell
\end{array}
$$

Fig. 2. Layout rewrites, with premises omitted. See Appendix C for rules with premises (Lean: `Core.Step`).

satisfies; we name each arm by its guard. A *disjoint sum* $\bigsqcup^\pi\{(c_j,\phi_j) \Rightarrow \ell_j\}$ instead selects its arm by the scalar at path π, using arm j for every subject whose scalar at π equals the code c_j. Its codes are distinct, and each guard ϕ_j must contain every admitted subject with code c_j.

The double-ended layout of Layout 3 starts as an ordinary sum over **Leaf** and **Node**, and tagging turns it into a disjoint sum keyed by its tag field t (Figure 3). A *tag* under path t names the scalar function that computes each value's code. The tag lists k guards, $\mathsf{guards}(t)$, and a selector s_t that maps each value v in their union to the index j of a guard ϕ_j containing v. The path t selects this index, $v.t = s_t(v)$, and is undefined on every other value. The selector resolves overlapping guards, so their order imposes no priority. A tag field $t : \mathsf{u}_{|t|}$ stores the code in $|t| = \lceil \log_2 k \rceil$ bits.

Finally, source types induce a straightforward *canonical* layout environment Δ_{can}. It binds each declared type to a sum with one arm per constructor **K**, whose payload is the product of the constructor's named fields, or ϵ if there are none, and a field of a declared type refers to that type. For the original red-black Tree, we write

$$\Delta_{\mathsf{can}}(\mathsf{Tree}) = \bigcup\{\mathbf{Leaf} \Rightarrow \epsilon, \quad \mathbf{Node} \Rightarrow \mathsf{c} : \mathsf{ref}\,\mathsf{Color} \times \mathsf{l} : \mathsf{ref}\,\mathsf{Tree} \times \mathsf{k} : \mathsf{u}_{64} \times \mathsf{r} : \mathsf{ref}\,\mathsf{Tree}\}.$$

4.4 Layout Rewrites

Figure 2 lists a representative subset of rewrite rules used in our implementation, each of which makes one local change to a layout. Commuting fields places data that is read together side by side, and boxing moves rarely read fields out of line behind a pointer. Tagging stores the choice of arm in a tag field that subsequent rewrites can transform like any other field, and memoizing an expression stores a computed result within the layout so that a program reads the result in place of recomputing it. Each rule carries premises, which Appendix C states in full. For example, splitting an n-bit field at k requires $k \le n$, and pruning an arm requires the remaining arms to cover every admitted subject. Crucially, the rule set is extensible because its proofs are modular. As Section 6.3 shows, a new rewrite rule only needs a proof for the part of the layout it changes.

A rewrite applies a rule to one occurrence of a sublayout, its *focus*, and leaves the rest of the layout fixed. A *layout context* $C[\Box]$ is the surrounding layout with the focus replaced by a hole, and filling the hole with ℓ gives $C[\ell]$. The grammar is

$$C ::= \Box \mid \pi : C \mid \&_n(C) \mid C \times \ell \mid \ell \times C \mid \bigcup\{\phi \Rightarrow C, \Omega\} \mid \bigsqcup^\pi\{(c,\phi) \Rightarrow C, \Omega\}.$$

For example, $C[\Box] = \mathsf{k} : \mathsf{u}_{64} \times \Box$ keeps the field k fixed while rewriting the component in the hole. A rule that replaces ℓ_0 by ℓ_1 takes $C[\ell_0]$ to $C[\ell_1]$, and its premises are checked on the subjects that

$$
\begin{aligned}
& \bigcup\{\textbf{Leaf} \Rightarrow \epsilon, \quad \textbf{Node} \Rightarrow \mathsf{c} \times \mathsf{l} \times \mathsf{k} \times \mathsf{r}\} \\
\xrightarrow{\mathsf{pad}_{46},\ \text{commute, unit}} \quad & \bigcup\{\textbf{Leaf} \Rightarrow \mathsf{pad}_{46}, \quad \textbf{Node} \Rightarrow \mathsf{c} \times \mathsf{l} \times \mathbf{k} \times \mathbf{r}\} \\
\xrightarrow{\mathsf{pad}_1,\ \text{commute, pad}_1} \quad & \bigcup\{\textbf{Leaf} \Rightarrow \mathsf{pad}_{46}, \quad \textbf{Node} \Rightarrow \mathsf{c} \times \mathbf{l} \times \mathbf{pad}_1 \times \mathbf{r} \times \mathbf{pad}_1 \times \mathbf{k}\} \\
\xrightarrow{\mathsf{ptr}_{44},\ \text{box, unit}} \quad & \bigcup\{\textbf{Leaf} \Rightarrow \mathsf{pad}_{46}, \quad \textbf{Node} \Rightarrow \mathsf{c} \times \&_{44}(\mathbf{l} \times \mathbf{pad}_1 \times \mathbf{r} \times \mathbf{pad}_1 \times \mathbf{k})\} \\
\xrightarrow{\mathsf{pad}_1} \quad & \bigcup\{\textbf{Leaf} \Rightarrow \mathbf{pad}_{46}, \quad \textbf{Node} \Rightarrow \mathbf{c} \times \mathbf{pad}_1 \times \&_{44}(\mathbf{l} \times \mathbf{pad}_1 \times \mathbf{r} \times \mathbf{pad}_1 \times \mathbf{k})\} \\
\xrightarrow{\mathsf{tag}} \quad & t : \mathsf{u}_1 \times \bigsqcup^{t}\{(0, \textbf{Leaf}) \Rightarrow \mathsf{pad}_{46}, \quad (1, \textbf{Node}) \Rightarrow \mathsf{c} \times \mathsf{pad}_1 \times \&_{44}(\mathsf{l} \times \mathsf{pad}_1 \times \mathsf{r} \times \mathsf{pad}_1 \times \mathsf{k})\}
\end{aligned}
$$

Fig. 3. Rewrites from the canonical body of Tree to Layout 3, abbreviating each field by its name. Each arrow rewrites the preceding layout at its **bolded** focus.

reach the hole through the context's field selections and guards. We write $\ell_0 \longrightarrow^* \ell_k$ for a finite sequence of rewrites.

Rewriting the body of a type updates the layout environment. An *update* $\Delta' = \Delta[\mathsf{T} \mapsto \ell_1]$ replaces the body assigned to T and leaves other entries unchanged, so every reference ref T, including recursive references, then uses ℓ_1. Figure 3 derives Layout 3 this way, rewriting the canonical body of Tree one focus at a time, and Section 6.2 gives the conditions under which replacements in contexts and updates to definitions preserve faithfulness.

5 Structural Representation Semantics

To reason about rewrites locally without tying correctness to a particular bit-level encoding, we first reason about *structural representations*: finite trees that retain the layout's components and their organization. An encoder then places these components in bits and heap blocks, and a decoder recovers them. Because a structural representation keeps boundaries between components explicit, we show rewrites can be proved correct without reasoning about memory in Section 6. In Section 7, we encode these representations in bits, where boundaries are implicit, and discuss limitations.

5.1 Structural Representations

Structural representations $\eta \in \mathbb{S}$ are defined by the grammar

$$\eta ::= k \mid () \mid (\eta_0, \eta_1) \mid \mathsf{cons}_\phi(\eta) \mid \mathsf{ptr}(\alpha, \eta). \quad \text{Lean: Enc}$$

Scalars $k \in \mathbb{N}$ hold source data, padding, tags or the results of memoized expressions, at widths supplied by the layout. A product contributes a pair, and a pointer records an address α together with its pointee's representation. An ordinary sum contributes a marker $\mathsf{cons}_\phi(\eta)$ naming the selected arm by its guard. The marker has no physical counterpart; it keeps arms distinguishable until a tag rewrite stores the choice in a field. A disjoint sum contributes the pair (c, η) of its selected code and payload; physical encoding stores the code only in the tag field.

For a layout ℓ under environment Δ, $\mathcal{W}_\ell^\Delta(v)$ is the set of structural representations of subject v (defined in Figure 4). Sets accommodate unspecified padding and addresses, as well as overlapping arms. The canonical layout gives a node exactly one structural representation, e.g.,

$$\mathcal{W}\big(\textbf{Node}(\textbf{Black}, \textbf{Leaf}, 5, \textbf{Leaf})\big) = \big\{\mathsf{cons}_{\textbf{Node}}\big(\big(\mathsf{cons}_{\textbf{Black}}(()),\ \mathsf{cons}_{\textbf{Leaf}}(()),\ 5,\ \mathsf{cons}_{\textbf{Leaf}}(())\big)\big)\big\}.$$

5.2 Faithfulness

To preserve programs over different layouts, a layout must be *total*, representing every source value, and *absolutely injective*, giving distinct values disjoint sets of representations. Totality ensures that a valid construction policy exists, one that chooses a permitted representation for every value a program builds. Absolute injectivity ensures that destruction has a unique answer.

Definition 1 (Faithfulness Lean: `Faithful`*).* A layout ℓ *faithfully represents* τ under Δ if, for all $v, w \in \mathcal{V}_\tau$,

$$\mathcal{W}^\Delta_\ell(v) \neq \varnothing \qquad \text{(totality)},$$
$$v \neq w \implies \mathcal{W}^\Delta_\ell(v) \cap \mathcal{W}^\Delta_\ell(w) = \varnothing \qquad \text{(absolute injectivity)}.$$

Notably, absolute injectivity is analogous to the "no confusion" property of Diatchki et al. [22]. Since $\mathcal{W}^\Delta_\ell$ is a relation, one value may have several representations, and absolute injectivity only forbids two values from sharing one. A representation η *determines* v when it represents v and no other value of its type. Under a faithful layout, each permitted representation determines its source value. We write $\mathcal{M}^\Delta_\ell(\eta) = v$ for this partial recovery function, reminiscent of the representation function of Hoare [34].

5.3 Well-Formedness

Injectivity alone cannot keep values from being lost, e.g., pruning the **Leaf** arm from the body of Tree maintains absolute injectivity, yet every tree ends in leaves, so every tree loses its representation. The judgment $A \vdash \ell$ wf establishes totaliy syntactically, checking each component only on the subjects that reach it. Its rules, defined in Appendix B, require, e.g., that field selections are defined, memoized expressions fit their widths, and sum arms cover every admitted subject.

A reference, on the other hand, defers its check to the corresponding layout definition. The environment Δ is *well formed* when each body is well formed on all values of its type, $\mathcal{V}_\mathsf{T} \vdash \Delta(\mathsf{T})$ wf. We require layout definitions to be *structurally recursive*, so each reference applies to a smaller source value. Syntactically, every reference in a body must occur within a named layout whose path begins with a field name, e.g., `l : ref Tree` applies the ref Tree layout to the left child. We package the assumptions as one invariant, Inv, and prove that well-formedness suffices for totality.

Definition 2 (Environment invariant Lean: `Core.Inv`*).* We write $\mathsf{Inv}(\Delta)$ when Δ is well formed and structurally recursive.

Lemma 3 (Well-formedness implies totality Lean: `WF.totality`). *If* $\mathsf{Inv}(\Delta)$ *and* $A \vdash \ell$ wf*, then* $\mathcal{W}^\Delta_\ell(v) \neq \varnothing$ *for every* $v \in A$.

6 Correctness of Layout Rewrites

We prove each rewrite locally, then lift its guarantee through the surrounding layout and recursive definitions. The local requirement compares the focus before and after the rewrite. Importantly, a focus need not faithfully represent its whole subject, e.g., the field `l : ref Tree` alone cannot distinguish nodes with equal left children but different keys.

$$\mathcal{W}^\Delta_\epsilon(v) = \{()\}, \quad \mathcal{W}^\Delta_{\mathsf{u}_n}(k) = \{k\} \qquad \mathcal{W}^\Delta_{\pi:\ell}(v) = \mathcal{W}^\Delta_\ell(v.\pi)$$
$$\mathcal{W}^\Delta_{\mathsf{pad}_n}(v) = \mathcal{V}_{\mathsf{u}n} \qquad \mathcal{W}^\Delta_{\ell_0 \times \ell_1}(v) = \mathcal{W}^\Delta_{\ell_0}(v) \times \mathcal{W}^\Delta_{\ell_1}(v)$$
$$\mathcal{W}^\Delta_{\mathsf{zero}_n}(v) = \{0\} \qquad \mathcal{W}^\Delta_{\&_n(\ell)}(v) = \{\mathsf{ptr}(\alpha, \eta) \mid \alpha \in \mathcal{V}_{\mathsf{u}_n},\ \eta \in \mathcal{W}^\Delta_\ell(v)\}$$
$$\mathcal{W}^\Delta_{\langle e \rangle}(v) = \{[\![e]\!](v)\} \qquad \mathcal{W}^\Delta_{\bigcup\{\phi_j \Rightarrow \ell_j\}}(v) = \{\mathsf{cons}_{\phi_j}(\eta) \mid v \in \phi_j,\ \eta \in \mathcal{W}^\Delta_{\ell_j}(v)\}$$
$$\mathcal{W}^\Delta_{\mathsf{ref}\,\mathsf{T}}(v) = \mathcal{W}^\Delta_{\Delta(\mathsf{T})}(v) \qquad \mathcal{W}^\Delta_{\bigsqcup^\pi\{(c_i, \phi_i) \Rightarrow \ell_i\}}(v) = \{(c_i, \eta) \mid v.\pi = c_i,\ \eta \in \mathcal{W}^\Delta_{\ell_i}(v)\}$$

Fig. 4. Structural interpretation of layouts. $\mathcal{W}^\Delta_\ell(v)$ is the set of structural representations permitted by layout ℓ for source value v. The layout environment Δ maps declared types to layout bodies. We write $\mathcal{V}_{\mathsf{u}_n}$ to represent the set $\{0, \ldots, 2^n - 1\}$, and $\mathcal{W}^\Delta_\ell(v) = \varnothing$ when evaluation is undefined. Lean: `Rep`

6.1 Local Refinement

Write $v \approx_\ell^\Delta w$ when v and w have a common representation under ℓ, that is, $\mathcal{W}_\ell^\Delta(v) \cap \mathcal{W}_\ell^\Delta(w) \neq \varnothing$. Layout ℓ' under Δ' *refines* ℓ under Δ on subjects A when

$$\forall v, w \in A. \quad v \approx_{\ell'}^{\Delta'} w \implies v \approx_\ell^\Delta w.$$

Refinement thus preserves absolute injectivity.

Definition 4 (Local refinement Lean: `Core.LocalRefinement`*).* Under Δ, layout ℓ_1 *locally refines* ℓ_0 on subjects A, written $\ell_0 \preceq \ell_1$, if both layouts are well formed on A and ℓ_1 refines ℓ_0 on A.

Each primitive rule proves refinement, a reflexive and transitive relation, with a *reconstruction function* $\mathsf{f} : \mathbb{S} \rightharpoonup \mathbb{S}$. For every $v \in A$ and $\eta' \in \mathcal{W}_{\ell_1}^\Delta(v)$, it returns $\mathsf{f}(\eta') \in \mathcal{W}_{\ell_0}^\Delta(v)$. Refinement then follows at once, i.e., if v and w share η' after the rewrite, they share $\mathsf{f}(\eta')$ before it. For example, padding reconstructs by dropping the added component, commutation by swapping the pair, and split_k by joining its two scalars as $(b_0, b_1) \mapsto b_0 + 2^k b_1$. Composing a derivation's reconstruction functions with the canonical decoder gives a decoder for the rewritten layout, although such a decoder may be inefficient if naively implemented.

6.2 Lifting Local Refinement

Since a rewrite leaves the rest of the layout unchanged, the distinctions that local refinement preserves at the focus carry over to the whole layout. Updating a layout definition is subtler: local refinement is proved with the old body at recursive references, but the update changes those references too.

Theorem 5 (Contextual replacement Lean: `Core.LocalRefinement.context`). *Let $C[\ell_0]$ be well formed on B, and let A be the subjects that C assigns to its hole from B. If $\ell_0 \preceq \ell_1$ on A, then $C[\ell_0] \preceq C[\ell_1]$ on B.*

Lemma 6 (Recursive binding Lean: `Core.LocalRefinement.bind`). *Let Δ wf, $\Delta(\mathsf{T}) = \ell_0$, and $\Delta' = \Delta[\mathsf{T} \mapsto \ell_1]$. If $\ell_0 \preceq \ell_1$ on $\mathcal{V}_\mathsf{T}$ under Δ and ℓ is well formed on B, then ℓ under Δ' refines ℓ under Δ on B.*

The proof inducts on one subject's representation under Δ' and compares the other subject's derivation of the same representation. At a reference to T, the induction hypothesis yields a shared representation of ℓ_1 under Δ, and local refinement gives a shared representation of ℓ_0. A *valid rewrite* satisfies local refinement at its focus and either replaces that focus under the premises of Theorem 5, or updates a definition under the premises of Lemma 6 while preserving structural recursion.

Theorem 7 (Soundness of local rewrites Lean: `Core.Rewrite.sound`). *Suppose $\mathsf{Inv}(\Delta)$ and $A \vdash \ell$ wf under Δ. Every valid rewrite from (Δ, ℓ) to (Δ', ℓ') establishes $\mathsf{Inv}(\Delta')$ and $A \vdash \ell'$ wf under Δ', and ℓ' under Δ' refines ℓ under Δ on A.*

If ℓ faithfully represents τ and $A = \mathcal{V}_\tau$, the preserved invariant and well-formedness give totality by Lemma 3, and refinement preserves absolute injectivity. Thus the new layout remains faithful. By transitivity, every layout reached from a canonical layout by valid rewrites is faithful.

6.3 Adding a Rewrite Rule

The rewrite rules of Figure 2 are sound but not complete, i.e., not every *faithful* layout can be derived. The rule set can be extended to express novel layouts, provided each new rule establishes local refinement (Definition 4) with a reconstruction function. Theorems 7 and 9 then apply to it

unchanged. As a concrete example, we realize an optimization used in Scion for storing an interval as an offset and a count [28, Fig. 5]. When $k < n$, an interval $[\mathsf{lo}, \mathsf{hi})$ can store its start and length,

$$\mathsf{lo} : \mathsf{u}_n \times \mathsf{hi} : \mathsf{u}_n \longrightarrow \mathsf{lo} : \mathsf{u}_n \times \langle (\mathsf{hi} - \mathsf{lo}) : \mathsf{u}_k \rangle \quad \text{Lean: } \texttt{Core.Interval.rewrite}$$

on admitted subjects with $\mathsf{lo} \le \mathsf{hi} < 2^n$ and $\mathsf{hi} - \mathsf{lo} < 2^k$. The reconstruction $\mathsf{f}(b, d) = (b, b + d)$ recovers the original pair, and, together with well-formedness, this establishes local refinement.

6.4 Program Preservation

Layout rewrites should change how programs represent data while preserving what they compute. We prove this for a pure language with constructors, matching, let bindings, and recursive calls, under call-by-value big-step semantics [40]. We compare terminating programs by their scalar results, the observations on which all layouts agree (details in Appendix G); this does not verify generated code.

A program evaluates over source values, $P; \rho \vdash e \Downarrow v$, or over represented values, $\Delta, b, \iota; P; \rho' \vdash e \Downarrow r$, where a construction policy b picks a representation for each new value and a destruction policy ι reads one back. The policies are *valid* when b succeeds exactly on values of the requested type and returns a permitted representation, and ι recovers each such value from every permitted representation. We write $r \sim_\Delta v$ when r represents v, relate environments that bind the same variables to related values, and call Δ *faithful* when ref T faithfully represents T for every type T.

Theorem 8 (Representation independence Lean: `Programs.representation_independence`). *Suppose Δ_0 and Δ_1 are faithful. For valid policies (b_0, ι_0) at Δ_0 and (b_1, ι_1) at Δ_1 and related inputs $\rho_0 \sim_{\Delta_0} \rho$ and $\rho_1 \sim_{\Delta_1} \rho$, every program P, expression e, and scalar $k \in \mathbb{N}$ satisfies*

$$\Delta_0, b_0, \iota_0; P; \rho_0 \vdash e \Downarrow k \quad \Longleftrightarrow \quad \Delta_1, b_1, \iota_1; P; \rho_1 \vdash e \Downarrow k. \tag{6}$$

Theorem 9 (Program preservation Lean: `Core.BodyRewriteSeq.{observations,roots}`). *Suppose $\mathsf{Inv}(\Delta_0)$, Δ_0 binds every declared type, and Δ_0 is faithful. If $\Delta_0 \longrightarrow^* \Delta_1$ is a finite sequence of valid updates to layout definitions, then Δ_1 is faithful, and (6) holds for all valid policies at the two endpoints and all related inputs.*

For rewritten layouts, the final layout yields a valid construction policy b (under Lemma 3). A valid destruction policy ι is recovered by inverting the rewrite rules through their reconstruction functions (Section 6.1).

7 From Structural Representation to Bits

Structural representations retain component boundaries and pointee contents, but a bit-level encoding that discards this information can be ambiguous even when the layout is faithful. Recovery from memory therefore requires a separate guarantee from an encoder and decoder, together called a *codec*. Prior work has studied the construction and verification of binary codecs extensively [4, 21, 26, 90, 100]. We connect this work to layout rewriting by stating the recovery property a codec must satisfy, then give a simple codec satisfying it under explicit premises.

7.1 Recovery from Physical Storage

A physical encoding $m = (\beta, H) \in \mathcal{M}$ consists of a finite top-level bit vector β and a finite heap H from addresses to bit vectors. For a layout ℓ under Δ, an encoder and a decoder are partial functions

$$\mathsf{encode}_\ell : \mathbb{S} \rightharpoonup \mathcal{M}, \qquad \mathsf{decode}_\ell : \mathcal{M} \rightharpoonup \mathbb{S},$$

and we leave Δ implicit in their notation. The *storage recovery property* requires decoding to preserve the represented source value. For all $v \in \mathcal{V}_\tau$, $\eta \in \mathbb{S}$, and $m \in \mathcal{M}$,

$$\eta \in \mathcal{W}_\ell^\Delta(v) \ \wedge \ \mathsf{encode}_\ell(\eta) = m \quad \Longrightarrow \quad \exists \eta' \in \mathcal{W}_\ell^\Delta(v). \quad \mathsf{decode}_\ell(m) = \eta'. \tag{7}$$

The decoder receives only the memory; it may return another representation of v, so recovery does not require identical padding, and it is unconstrained on memories that no encoding produces.

Theorem 10 (Recovery after rewriting and encoding Lean: `Core.storage_recovery`). *Suppose* $\mathsf{Inv}(\Delta)$ *and* ℓ *is well formed on* $\mathcal{V}_\tau$ *and faithfully represents* τ *under* Δ. *Let* (Δ', ℓ') *be obtained by a finite sequence of valid rewrites. If an encoder and a decoder for* ℓ' *under* Δ' *satisfy Equation* (7), *then for every* $v \in \mathcal{V}_\tau$, $\eta \in \mathcal{W}_{\ell'}^{\Delta'}(v)$, *and successful encoding* $\mathsf{encode}_{\ell'}(\eta) = m$, *the decoder returns a representation that determines* v, *so* $\mathcal{M}_{\ell'}^{\Delta'}(\mathsf{decode}_{\ell'}(m)) = v$.

The theorem guarantees recovery of the source value provided that encoding succeeded. For an end-to-end instance, the mechanization supplies a simple forward encoder and decoder, which place scalars, products, and pointers in a bit vector and heap and reads them back left to right. Encoding succeeds when every ordinary sum has been tagged, every tag field is stored whole in its scope, and there are enough addresses for the pointers. Decoding inverts every successful encoding when, in addition, each tag field precedes its payload. Recovery of the source value then follows from Theorem 10; the mechanization carries this out for Layout 2, and Appendix F elaborates on the conditions.

7.2 Decoding Limitations

Recovery from storage also depends on factors beyond faithfulness. For example, commuting a tag after the arm it selects preserves faithfulness, but the forward decoder is required to read the arm before learning which arm it is. On the other hand, a backward decoder, i.e., one that reads fields in reverse order, could still recover the values. Moreover, variable-width arms can make the encoding itself ambiguous. Consider the type $\mathsf{S} = \mathbf{A}(x : \mathsf{u}_1) \mid \mathbf{B}(z : \mathsf{u}_2)$ with initial layout $t : \mathsf{u}_1 \times \lfloor\rfloor^t\{(0, \mathbf{A}) \Rightarrow x : \mathsf{u}_1,\ (1, \mathbf{B}) \Rightarrow z : \mathsf{u}_2\}$, and then store a pair of S values with the first component's tag after its payload and the second component's tag before. Then,

$$(\mathbf{A}(1), \mathbf{B}(1)) \longmapsto \mathtt{10} \parallel \mathtt{101} = \mathtt{10101},$$
$$(\mathbf{B}(2), \mathbf{A}(1)) \longmapsto \mathtt{101} \parallel \mathtt{01} = \mathtt{10101}$$

so no decoder can recover both pairs. Both tags lie at the boundary between the components, but only a tag reveals where that boundary is. Narcissus resolves such ambiguities by fixing boundaries from the surrounding format [21, Sec. 3.1]. Local rewrites thus preserve faithfulness but not physical recoverability, which depends on the layout as a whole.

8 The Isomer Programming Language

The preceding sections define how a layout organizes data in memory. The compiler must abide by this specification, but, importantly, is free to choose how construction, destruction, and memory management are realized. We next describe how Isomer exploits this to generate efficient code.

8.1 Core Language and Layout Selection

$$x, f, \mathsf{T}, \mathbf{K} \in \mathsf{Id} \qquad c \in \mathsf{Lit} \qquad n \in \mathbb{N} \qquad t \in \{\mathsf{u}_n, \mathsf{i}_n, \mathsf{f}_n\}$$

$$\begin{array}{lll} P ::= \overline{d}\ \overline{F} & F ::= \mathtt{fn}\ f(\overline{x : \tau}) : \tau = e & e ::= x \mid c \mid \mathbf{K}(\overline{e}) \mid f(\overline{e}) \mid e_1 \oplus e_2 \\ d ::= \mathtt{data}\ \mathsf{T} = \kappa_1 \mid \cdots \mid \kappa_n; & \tau ::= t \mid \mathsf{T} \mid [x]\tau & \quad \mid\ \mathtt{let}\ x : \tau = e_1\ \mathtt{in}\ e_2 \\ \kappa ::= \mathbf{K}(\overline{x : \tau}) & p ::= x \mid _ \mid n \mid \mathbf{K}(\overline{p}) & \quad \mid\ \mathtt{match}\ e\ \{\overline{p \Rightarrow e}\} \end{array}$$

Fig. 5. Monomorphic, first-order core of Isomer.

Figure 5 gives the grammar of the monomorphic, first-order core of Isomer. The frontend monomorphizes generic definitions and defunctionalizes closures into ADTs, so every value has a

concrete type. Scalars are signed integers, unsigned integers, and floating-point numbers. Variable-length arrays $[x]\tau$ are first-class, with an in-scope scalar length x [37, 68]. The formalism models a finite array as a sum with one constructor per length, indexed by its length; the implementation provides an explicit array construct for simplicity.

For types without rewrites applied, Isomer uses a *default* layout. Fields keep their declared order, fixed-size fields are inlined, recursively defined fields sit behind a pointer, and tags are rounded up to whole bytes. When a sum is stored inline within another value its arms are padded to a common size, so that following fields have statically known offsets. For heap-allocated values, we size cells per constructor rather than padding them to a common size; each cell includes a 32-bit reference count. Sections 9.3 and 9.4 show that this simple policy strikes a balance between layout compactness and runtime performance.

8.2 Layout-Directed Compilation

Well-formedness. Before generating code, the compiler verifies the machine can realize each layout. This *machine well-formedness* requires that tags and array lengths are recoverable, fields admit an acyclic reading order satisfying their offset dependencies, every cycle in the storage graph crosses a pointer, and storage sizes meet allocation constraints. Pointer compression, for example, must comply with the requirements of the target architecture. These conditions are in part intended to exclude the ambiguous encodings of Section 7.2.

Construction and destruction. Unlike layout description languages [6, 14, 28] which decode a value by inverting its layout specification, Isomer generates code by following a layout's rewrites from a canonical layout. Each rewrite updates two artifacts, (1) the layout's *decoder*, a function from a value's stored form to the source value, and (2) a *storage plan* that records each stored component's physical metadata, e.g., its offset and width. Construction proceeds by reading the storage plan, while destruction, i.e., pattern matching and field access, queries the decoder. A new rule thus requires supplying only its reconstruction function and its changes to the storage plan.

Producer and consumer specialization. Compiling pattern matches and constructions field by field, as the source writes them, can make a program asymptotically slower under a rewritten layout. Each field the source binds or supplies is a value of its own, while the layout may keep that value only inside a larger component. For example, when an unfolded list stores two elements per cell, matching `Cons(h, t)` binds `t` to a tail with no standalone representation. Rebuilding that tail means repacking the entire suffix, and doing so at every step turns a linear traversal quadratic. Isomer instead specializes consumers and producers to the layout through a source-to-source translation that combines deforestation [94] and partial evaluation [39] techniques to avoid this asymptotic slowdown when possible. The algorithms are discussed further in Appendix E.

Specialized pattern matching is then lowered to decision trees [60] whose tests read physical encodings, and branch provenance removes tests whose outcome is already known. Unlike the formalism, the implementation treats the arms of a sum as an ordered list, and a rule that commutes arms lets programmers choose which guard construction tests first[2].

Instruction selection. Naively executing a decoder for every query would rebuild an entire source value just to read a single field. The compiler instead optimizes the decoder, using the storage plan to simplify instruction selection. Knowledge of field positions and widths lets the compiler combine representation operations with the instructions that use their results. For example, once a branch establishes that a pointer carries the low-bit tag t, untagging folds into each field access,

[2]We further explored decoupled programmer control of decision-tree compilation but saw little runtime benefit for our workloads, consistent with Scott and Ramsey [76]. Isomer defaults to the necessity heuristic of Maranget [60].

as in GHC's pointer tagging [61], so a field at byte offset d from the untagged pointer is read at offset $d - t$ from the tagged word, without a separate mask instruction. Finally, Isomer lowers the specialized program to SSA [20, 75], applies standard compiler optimizations [1, 18, 48, 69], and emits LLVM IR [45].

8.3 Reference Counting

Isomer manages memory with ownership-based reference counting, following the compiler-based approach of prior work [57, 73, 89]. In Isomer, however, the count is itself part of the layout. The rewrite rule `rc n` prepends an unsigned n-bit count to the focused layout, a field that the compiler manages but other rewrites treat like any other, so a count can be narrowed, moved, or removed entirely. Omitting the count asserts that each instance of the type has a single owner, which Isomer verifies after reference count optimizations are applied. As in Koka [73, Sec. 2.7.2], counts are *sticky*, so an n-bit count saturates once its high bit is set, and its cell is then neither reclaimed nor reused.

Layout control lets programmers choose cell sizes that enable reuse and omit reference counts when ownership permits. Dynamic reuse (`rc-1`) updates a dying cell in place when its count is one at run time [10, 73]. It enables cells of the same size to be reused between constructors of the same or different types, and arrays with a single owner to be updated in place. Layout rewrites let programmers choose between per-arm allocation sizes and a common cell size. Per-arm sizes avoid padding smaller arms, while padding all arms to a common size enables cell reuse across all constructors of the type. Additionally, the programmer can entirely omit the reference count field from a layout, which the compiler accepts only when static ownership analysis verifies unique ownership; we call this optimization `rc-0`. The k-d tree case study, the layout design space exploration, and the Ribbit comparison exercise `rc-0` (Sections 9.2, 9.4 and 9.5).

8.4 Layout Coercion

One layout rarely suits every phase of a program, since construction and traversal can favor different representations of the same ADT. The scheduling directive `coerce` changes a value's representation at a chosen binding without changing the program. Logically, the conversion is a catamorphism h [62] that converts a value's children and rebuilds each constructor in the target representation,

$$h\big(\mathbf{K}_0(\bar{a}, t_1, \ldots, t_k)\big) = \mathbf{K}_1\big(\bar{a}, h(t_1), \ldots, h(t_k)\big).$$

Here $\bar{a}$ are the nonrecursive fields, and subscripts 0 and 1 mark the source and target representations. Because both sides share one ADT, the fold's constructor functions are fixed, and the compiler implements them by specializing construction to the target layout and matching to the source layout. Conversion thus reuses the machinery of Section 8.2. The 2-3 tree and k-d tree case studies evaluate `coerce` between their construction and query phases (Section 9.2).

8.5 Using Memoized Expressions

A layout can also store what a program would otherwise recompute. A memoized expression moves work from queries to construction. For example, given the type $\mathsf{T} = \mathbf{T}(a : \mathsf{i}_{32}, b : \mathsf{i}_{32})$, the rewrite `memoize (a + b)` stores the sum of fields `a` and `b` in a new field, which the generated constructors compute. During destruction, the compiler annotates each cache load with the expression it represents. After canonicalization, common-subexpression elimination (CSE) [18] replaces matching computations with the cached value. Memoized expressions may also be recursive; the interval-tree case study (Section 9.2) evaluates such an instance.

9 Evaluation

Our evaluation demonstrates that Isomer achieves competitive performance against existing implementations using automatically selected default layouts, and that composing rewrites can improve runtime and reduce peak memory further. Layout control also complements automatic memory management by shrinking allocations that reuse does not eliminate. We substantiate these findings with measurements of runtime and peak memory.

9.1 Experimental Methodology

All benchmarks run single-threaded on an Apple M2 Pro with 10 cores and 16 GiB of memory, running macOS 26.0. Isomer emits LLVM IR, which Clang 19.1.7 compiles with `-O3 -mcpu=native`. We compare against MLton [95], Morphic [9, 10], Ribbit [5, 6], and Koka [47, 57][3]. Every timed system measures wall-clock time, and each configuration runs in fresh processes, with 2 discarded warm-ups followed by 14 measured runs. We report the mean of the 10 runs left after discarding the 2 fastest and 2 slowest, with approximate 95% error bars. For memory, we report the median peak resident set size (RSS). Unless noted, we use mimalloc 3.6.1 [49], and Appendix H.4 repeats the results with the system allocator. We verify every configuration computes the same answer and Isomer has no memory leaks. Experiments below describe any remaining setup.

AI Use. AI [3, 66] assisted with implementation of the Isomer compiler from an initial prototype, proof mechanization, the benchmark harness, paper figures, and the layout search of Section 9.4.

9.2 Case Studies

Three case studies measure what layout rewrites gain in runtime and memory within a single program, by composing successive rewrites on a k-d tree, memoizing an expression in an interval tree, and changing layouts between the insertions and range scans of a 2-3 tree.

9.2.1 *Composing Rewrites on k-d Trees.*

$\text{KdTree} = \textbf{Leaf}(x, y, z : \text{f}_{32}) \mid \textbf{Node}(min_x, min_y, min_z, max_x, max_y, max_z : \text{f}_{32}, l : \text{KdTree}, r : \text{KdTree})$

We explore the interaction of layout rewrites through two-point correlation counting [63]. The program counts points within a query radius, using a k-d tree [8] to accelerate the query. We build a tree over 2^{20} points and issue 2^{20} queries. Starting from Isomer's default layout, each configuration in Figure 6 adds one optimization to the layout before it, realized through a short sequence of rewrites. The results show that effective layouts can be reached incrementally by composing simple rewrites, and additionally expose a tradeoff between runtime performance and memory use.

Because the tree is built once and then only read, `1-unique` can drop the reference count, which ownership analysis confirms no cell needs (`rc-0`, Section 8.3), and every leaf moves to an allocation bin half the size. `2-hoist` then moves the tag into low bits of each pointer [6, 25, 61, 95], so a traversal dispatches on bits it already holds instead of loading a byte, and a leaf is just the three coordinates. Next, most nodes near the bottom of the tree have two leaf children, and `3-LL` unfolds both into their parent [7, 31, 78], turning three cells into one 48-byte cell and eliminating two pointer dereferences. `4-NN` unfolds two **Node** children the same way, which barely moves either metric but places two levels of the tree in one cell for later steps to reshape. `5-soa` then reorders the inlined leaf coordinates and the bounding-box fields for vectorized arithmetic [7, 97, 101], so one 2-wide comparison tests both leaves, and `6-LT` inlines a lone **Leaf** child while keeping its sibling behind a pointer. The last two incremental steps trade runtime for memory. `7-narrow` refines pointers to 48 bits, moving two-**Node** cells from the 112- to the 96-byte allocation bin [15, 46] at the cost of a slower load for every child, and `8-LN` also inlines the **Node** sibling of a lone leaf.

[3] MLton `on-20241230-release` (https://github.com/MLton/mlton), Morphic `da679c4` (https://github.com/morphic-lang/morphic), Ribbit `35a2225` (https://gitlab.inria.fr/ribbit/ribbit), and Koka `v3.2.9` (https://github.com/koka-lang/koka)

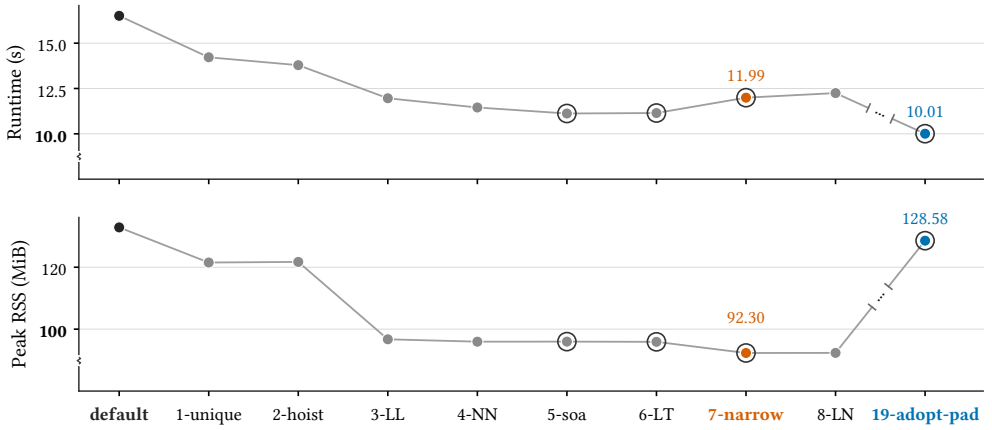


Fig. 6. Lower is better. Runtime and peak RSS across k-d tree layouts. The line denotes successive composition of rewrite rules. • marks minimum runtime, • minimum RSS, and ○ the observed Pareto frontier.

The final layout `19-adopt-pad` performs a `coerce` step (included in the timed region) between the construction and query phases. This scheduling directive builds the tree under a layout that keeps children boxed and steals bits for the tag, then converts it to an unfolded, alignment-friendly layout for better instruction selection during the queries. It also builds the point buffer's array inline in the counted cell, and then pads leaves in the read layout within their existing allocation bin. Of the layouts shown, `19-adopt-pad` is the fastest, 1.65× faster than the default, and `7-narrow` has the smallest peak RSS, 31% below the default's. Importantly, the fastest and smallest layouts differ, thus no single layout is optimal across both runtime and peak memory use.

9.2.2 *Memoized Expressions in Interval Trees.*

Tree = **Leaf** | **Node**(l : Tree, r : Tree, lo : i_{64}, hi : i_{64})

An interval-tree query uses the maximum upper endpoint in a subtree to choose its next branch [19, Sec. 15.3]. An unaugmented tree recomputes this maximum by traversing the left subtree of each node that a query visits, so a query can take time linear in the number of intervals. Augmenting the maximum at every tree node instead makes a query take time proportional to the tree's height. The cached maximum is defined recursively,

$$\mu(t) = \begin{cases} 0, & t \in \textbf{Leaf} \\ \max\big(\mu(\mathrm{l}), \mu(\mathrm{r}), \mathrm{hi}\big), & t \in \textbf{Node}(\mathrm{l}, \mathrm{r}, \mathrm{lo}, \mathrm{hi}). \end{cases}$$

Memoized expressions cache this value directly in the ADT, analogous to the reduction *annotations* in Bonsai [74, Sec. 4]. Isomer realizes this augmentation through layout rewrites over the unaugmented source. MLton and Morphic reach the same asymptotic cost only by changing the source, so all three systems also run a hand-augmented program that stores the maximum in every constructor. We time the queries alone, issuing 128,000 point queries to trees of 2^n intervals for $n \in [6, 19)$, with the compiler configurations of Section 9.3. Across all tree sizes, the memoized Isomer layout is 1.58× faster than hand-augmented Morphic and 1.10× faster than hand-augmented MLton (geomean) as shown in Figure 7 (left). The memoized layout also performs similarly to the hand-augmented Isomer program at every size, showing that expressing the augmentation as a layout rewrite adds no overhead.

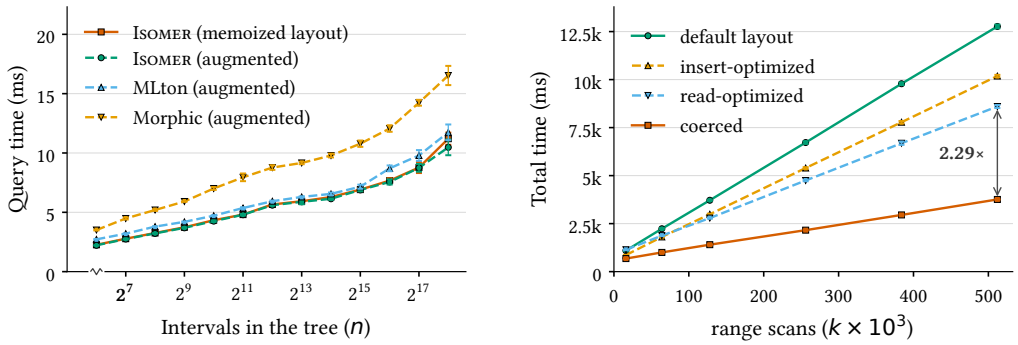


Fig. 7. Lower is better. Interval-tree query time with memoized expressions in Isomer and with a hand-augmented source in Isomer, MLton, and Morphic (left), and total runtime for 2–3 tree insertion and range scans, including data structure translation (right).

9.2.3 *Changing Layouts in 2-3 Trees.*

Tree = **Leaf** | **Two**(l : Tree, k : u_{64}, r : Tree) | **Three**(l : Tree, k_1 : u_{64}, m : Tree, k_2 : u_{64}, r : Tree)

The interval-tree cache moves recursive work from queries to construction; a complementary opportunity arises when construction and traversal favor different representations. The `coerce` directive lets a program convert once between phases and use a compact layout for repeated scans. A 2-3 tree [42, Sec. 6.2.4] is a balanced search tree whose internal nodes hold one or two keys and have two or three children, respectively, with all leaves at the same depth. We insert 1.6×10^6 keys, then issue $10^3 \times k$ range scans for $k \in \{16, 64, 128, 256, 384, 512\}$ and a final reduction. Figure 7 (right) compares the default layout with three configurations. `insert-optimized` pads **Two** cells to match **Three** cells for reuse, `read-optimized` inlines empty **Leaf** children in leaf-level nodes, and `coerced` builds with `insert-optimized` and then translates via `coerce` to `read-optimized` before scanning. The best fixed layout depends on the workload, switching from `insert-optimized` to `read-optimized` between 64,000 and 128,000 scans. `coerced` pays for one conversion but is faster than both throughout the sweep. At 512,000 scans, it is 2.29× faster than the best fixed layout. The `coerce` directive thus turns layout choice from a program-wide decision [28] into a phase-specific one.

9.3 Comparison with Monomorphizing Languages

MLton and Morphic are whole-program compilers that monomorphize and, unlike Isomer, do not offer user control of layouts. Morphic computes strongly connected components (SCCs) over datatype-size dependencies and boxes an arm when its payload refers to a type in its own SCC. MLton's representation emerges from interacting whole-program passes, including monomorphization, unboxing, argument flattening, and representation selection [95]. Figure 8 shows the layout each system chooses for `tmap`'s `Tree` ADT. Morphic opts for aligned fields at the cost of unused bits, while MLton employs more advanced layouts automatically, e.g., bit stealing techniques.

Tree = **Leaf** | **Bin**(l : Tree, x : i_{32}, r : Tree)

Morphic	$t : \mathrm{u}_1 \times \mathrm{zero}_7 \times \mathrm{pad}_{56} \times \bigsqcup^t \{(0,\mathbf{Leaf}) \Rightarrow \mathrm{pad}_{64},\ (1,\mathbf{Bin}) \Rightarrow \&_{64}(\mathrm{rc} : \mathrm{u}_{64} \times l : \mathrm{refTree} \times x : \mathrm{i}_{32} \times \mathrm{pad}_{32} \times r : \mathrm{refTree})\}$
MLton	$t : \mathrm{u}_2 \times \bigsqcup^t \{(1,\mathbf{Leaf}) \Rightarrow \mathrm{zero}_{62},\ (0,\mathbf{Bin}) \Rightarrow \&_{62}(\mathrm{gc_metadata} : \mathrm{u}_{64} \times x : \mathrm{i}_{64} \times l : \mathrm{refTree} \times r : \mathrm{refTree})\}$
Isomer (default)	$\&_{64}(\mathrm{rc} : \mathrm{u}_{32} \times t : \mathrm{u}_8 \times \bigsqcup^t \{(0,\mathbf{Leaf}) \Rightarrow \epsilon,\ (1,\mathbf{Bin}) \Rightarrow l : \mathrm{refTree} \times x : \mathrm{i}_{32} \times r : \mathrm{refTree}\})$
Isomer (optimized)	$\bigsqcup^t \{(0,\mathbf{Leaf}) \Rightarrow \mathrm{pad}_{56},\ (1,\mathbf{Bin}) \Rightarrow \&_{56}(l : \mathrm{refTree} \times r : \mathrm{refTree} \times \mathrm{rc} : \mathrm{u}_{32} \times x : \mathrm{i}_{32})\} \times \mathrm{zero}_7 \times t : \mathrm{u}_1$

Fig. 8. Layouts for Tree in `tmap`. Optimized Isomer is `paired-high-tags` discussed in Section 9.4.

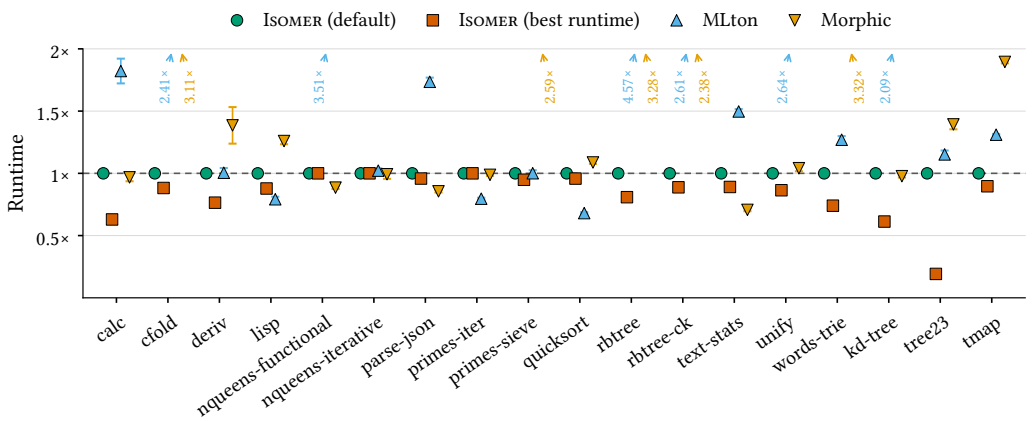


Fig. 9. Lower is better. Runtime relative to default Isomer. The best-runtime layout is selected independently for each benchmark. Arrows indicate ratios beyond the displayed range.

The suite is the union of the 15 benchmarks evaluated in both Morphic papers [9, 10] in addition to `kd-tree`, `tree23` (Section 9.2), and `tmap` (Section 9.6). The benchmark `primes-iter` declares no ADTs but remains in the suite for completeness. Following the AST-based translation of Morphic [9], we generate each MLton and Morphic program from the same Isomer source, so the comparison measures code generation and data representation rather than how each program was written. Both reference compilers check every array access, so we enable Isomer's optional bounds checks.

Isomer's arrays are immutable, whereas SML's are mutable. A direct AST translation to MLton therefore cannot safely update an array in place unless no other reference observes the original. For each benchmark, we time four MLton array implementations and use the fastest whose output agrees (Appendix H.5). Morphic uses borrow-based reference counting [10] with a 64-bit count. We acknowledge that cross-language comparison is inherently contrived, and provide these results as a baseline that places Isomer relative to the state of the art.

Default layouts. With default layouts, Isomer is 1.53× faster than MLton and 1.42× faster than Morphic (geomean, Figure 9). Against MLton, both extremes come from garbage collection, whose copying is costly on the large live sets of `rbtree` and `rbtree-ck`, but whose bump allocation beats Isomer on programs with many short-lived objects, e.g., `lisp`. Against Morphic, which updates unique arrays in place but never reuses a constructor's cell, dynamic reuse likely contributes to Isomer's lead. Appendix H.2 compares peak RSS, which for MLton includes its garbage-collected heap.

Parity layouts. To separate layout from code generation, we also run Isomer at Morphic's own layouts, derived with our rewrite rules and matched bit for bit, in 18 comparisons (Appendix H.3). Isomer remains 1.18× faster (geomean), suggesting its lead comes from additional compiler optimizations across instruction selection and memory management.

9.4 Layout Design Space Exploration

We now turn to *exploration* of layouts. Holding each program's source and the compiler fixed, we measure 152 layout configurations across the 18 benchmarks of Section 9.3. Following Helion's approach of combining language-model proposals with search [16], we use language models to propose rewrite sequences and further refine the resulting layouts with a straightforward search

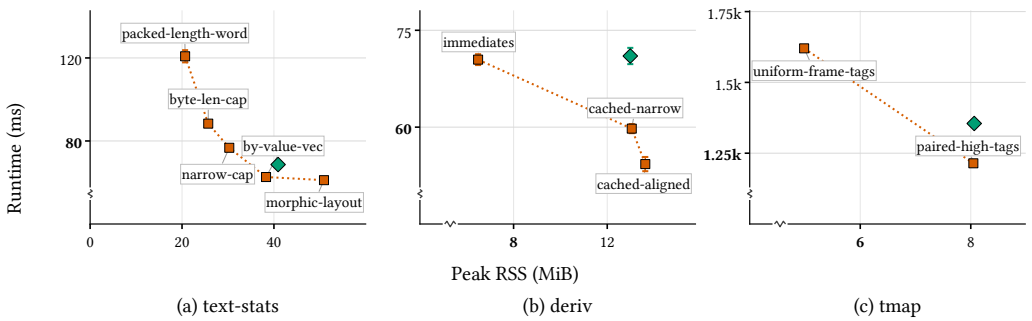


Fig. 10. Lower is better. Runtime–RSS tradeoffs among Isomer layouts for selected benchmarks. Axis break marks indicate nonzero origins. ◆ marks the default Isomer layout, and ··■·· marks a rewritten Isomer layout on the Pareto frontier.

procedure [59]. Selecting the fastest measured layout for each benchmark gives a 1.27× geometric-mean speedup over Isomer with a default layout, ranging from none on `primes-iter` (no ADTs) to 5.21× on `tree23`. Independently selecting the lowest-RSS layouts reduces peak RSS by 23% (geomean). Figure 10 shows the frontiers of three benchmarks, and Appendix H.1 shows the frontier of every benchmark with a non-default Pareto-optimal layout.

The three frontiers of Figure 10 come from different mechanisms. In `text-stats`, inlining data of vectors avoids an indirection, whereas narrower capacities and lengths save space for short words. Packing lengths into reference words saves further space, but also adds a decoding step when reading. In `deriv`, the fast layouts cache each subexpression's leaf count, as the interval tree caches its maximum, so counting reads one field, and `immediates` keeps small constants inside the pointer word, which saves a heap cell per constant and reduces peak memory by 50%. Finally, in `tmap`, `uniform-frame-tags` refines small keys from 32 to 24 bits and matches tree and continuation-frame cell sizes for reuse. The faster `paired-high-tags` layout (displayed in Figure 8) keeps child references adjacent and places the tree tag in the highest bit of the reference word.

9.5 Comparison with a Data-Description Language

Ribbit [5, 6] is a layout description language for general-purpose ADTs. We pair Ribbit's fastest measured layout on the three benchmarks of Baudon et al. [6] with the corresponding Isomer layout, and compile both to LLVM IR. Ribbit has no automatic memory management, while Isomer's layouts drop reference counting through the `rc-0` optimization for two benchmarks, and uses reference counting for third. Isomer achieves a 1.58× geometric-mean speedup (Figure 11); the gap lies in code generation. For example, Ribbit compiles every function in destination-passing style [5, Sec. 5.2] so every iteration makes an out-of-line call that returns through a pointer, while Isomer keeps these values in registers. We therefore report the speedup as a comparison of the two implementations, rather than a claim about the performance potential of Ribbit's layout language.

9.6 Allocation and Reuse with Koka

We test whether layout control composes with the reference counting optimizations of Koka [47], in which a constructor reuses the cell of a same-sized owned value that has just been destructed [73]. We run 4 of the 5 benchmarks of Lorenzen et al. [57] with $n = 100{,}000$, omitting `ftree` because it requires polymorphic recursion. We compile Koka with `-O3`, and both systems use Koka's allocator. Using Koka's exact layouts, Isomer's peak memory is at most that of Koka's fully in-place versions, `koka-fip`, on every benchmark (9% lower peak RSS, geomean; Figure 11). Notably, the fully in-place

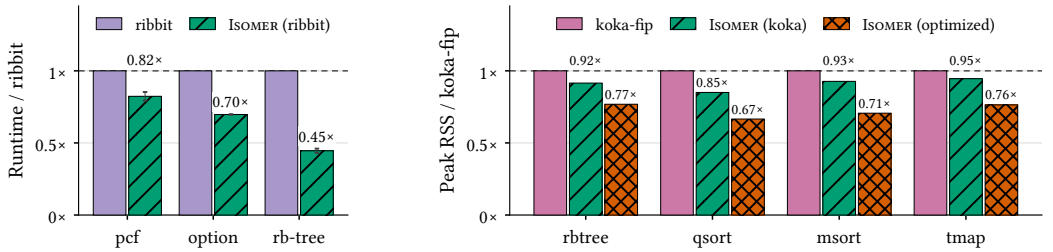


Fig. 11. Lower is better. Runtime relative to Ribbit with corresponding layouts (left), and peak RSS relative to `koka-fip` (right). In Isomer (ℓ), ℓ denotes the layout used, e.g., Isomer (koka) uses Koka's layout.

versions of `msort` and `qsort` give different constructors the same size by adding `Pad` fields to the source [57, Sec. 4.3], whereas Isomer changes cell sizes solely through layout rewrites. The `qsort` program shows the largest gap, which partly is a result of tail recursion modulo cons (TRMC) [48], which Isomer applies to mutually recursive functions while Koka does not. Further rewrites, such as the bit stealing of `tmap` in Section 9.4, lower peak memory by an *additional* 20% (geomean), demonstrating that our layout rewrites can complement reuse by reducing the size of allocations that are not optimized away.

10 Related Work

Automatic and type-driven layouts. MLton [95] and Morphic [9, 10] automatically select ADT layouts through whole-program analyses. Gibbon and its successors [43, 80, 92, 93] compile recursive ADTs to dense, serialized buffers. Such analyses could select among the layouts our rewrites express. OCaml [11, 51] and Rust [85, 86] annotations select from a fixed set of layouts without changing application code. Haskell [70] and OxCaml [38, 58] expose unboxed types, with levity polymorphism [24] allowing Haskell code to abstract over representation. These choices often require changes to the source code.

Decoupled layouts and schedules. Ribbit [6], Dargent [14], and Scion [28] separate layouts from type declarations but specify the target layout directly. Recently, Rig [88] proposed layout transformations for finite ADTs from type isomorphisms. We instead optimize layouts of recursive ADTs via composable rewrite rules. A layout description language checks each target layout as a whole and decodes it by inverting its final form, so it admits only encodings its checker can validate and its compiler can invert. Our rules carry local proofs and reconstruction functions that compose; a new rule need only satisfy a local correctness criterion rather than modification of a monolithic validation checker. Rewrites also let a programmer or search procedure move between valid layouts incrementally. Dense [29, 44, 79, 83, 98, 102] and sparse [2, 17, 54, 55, 91, 99] array languages expose layouts for n-dimensional arrays. In the same spirit, Halide and related systems separate algorithms from execution schedules for a specific domain [12, 13, 71, 77]. Hall et al. [32] identify layout and data movement as further directions. Lift [81], Elevate [30], Exo [35, 36], and ATL [53, 54] compose rewrites for verifiability and extensibility, a principle we apply to data representation.

Data representations and codecs. Hoare [34] proves correctness through a function from concrete to abstract values, the role of our recovery function. Data description languages generate parsers, printers, and other tools from a format description [4, 26], and a verified Protocol Buffer compiler proves every serializer and deserializer it generates correct [100]. EverParse proves its generated parsers correct [72, 82] and each message's encoding unique [72]. Narcissus [21] derives verified codecs via local proofs; our destruction policy is likewise assembled from local proofs over structural representations instead of bits. Its formats, like our layouts, may admit several representations per

value and must remain injective. These systems verify executable codecs, each for a given format. Similar to our rewrite rules, van Geest and Swierstra [90] chain verified conversions toward binary encodings, for formats built from fixed-width words, products, and dependent fields.

References

[1] Alfred V. Aho, Monica S. Lam, Ravi Sethi, and Jeffrey D. Ullman. 2006. *Compilers: Principles, Techniques, and Tools* (2nd ed.). Addison-Wesley, Boston, MA, USA.

[2] Willow Ahrens, Daniel Donenfeld, Fredrik Kjolstad, and Saman Amarasinghe. 2023. Looplets: A Language for Structured Coiteration. In *Proceedings of the 21st ACM/IEEE International Symposium on Code Generation and Optimization* (Montréal, QC, Canada) *(CGO '23)*. Association for Computing Machinery, New York, NY, USA, 41–54. doi:10.1145/3579990.3580020

[3] Anthropic. 2026. Claude. https://claude.ai. Large language model (Opus 5.5, Fable 5.1)..

[4] Julian Bangert and Nickolai Zeldovich. 2014. Nail: a practical tool for parsing and generating data formats. In *Proceedings of the 11th USENIX Conference on Operating Systems Design and Implementation* (Broomfield, CO) *(OSDI'14)*. USENIX Association, USA, 615–628.

[5] Thaïs Baudon. 2024. *Algebraic Data Types for High-Performance Computing*. Theses. Ecole normale supérieure de lyon - ENS LYON. doi:10.70675/7bb9b51ezf689z4cb5z8d80z794397b5f052

[6] Thaïs Baudon, Gabriel Radanne, and Laure Gonnord. 2023. Bit-Stealing Made Legal: Compilation for Custom Memory Representations of Algebraic Data Types. *Proc. ACM Program. Lang.* 7, ICFP, Article 216 (2023), 34 pages. doi:10.1145/3607858

[7] Carsten Benthin, Ingo Wald, Sven Woop, and Attila T Áfra. 2018. Compressed-leaf bounding volume hierarchies. In *Proceedings of the Conference on High-Performance Graphics*. 1–4. doi:10.1145/3231578.3231581

[8] Jon Louis Bentley. 1975. Multidimensional binary search trees used for associative searching. *Commun. ACM* 18, 9 (Sept. 1975), 509–517. doi:10.1145/361002.361007

[9] William Brandon, Benjamin Driscoll, Frank Dai, Wilson Berkow, and Mae Milano. 2023. Better Defunctionalization through Lambda Set Specialization. *Proc. ACM Program. Lang.* 7, PLDI, Article 146 (2023), 24 pages. doi:10.1145/3591260

[10] William Brandon, Benjamin Driscoll, Frank Dai, Jonathan Ragan-Kelley, Mae Milano, and Alex Aiken. 2026. Fully-Automatic Type Inference for Borrows with Lifetimes. *Proc. ACM Program. Lang.* 10, OOPSLA1 (2026), 597–623. doi:10.1145/3798221

[11] Nicolas Chataing, Stephen Dolan, Gabriel Scherer, and Jeremy Yallop. 2024. Unboxed Data Constructors: Or, How cpp Decides a Halting Problem. *Proc. ACM Program. Lang.* 8, POPL, Article 51 (2024), 31 pages. doi:10.1145/3632893

[12] Chun Chen, Jacqueline Chame, and Mary Hall. 2008. *CHiLL: A framework for composing high-level loop transformations*. Technical Report 08-897. University of Southern California.

[13] Tianqi Chen, Thierry Moreau, Ziheng Jiang, Lianmin Zheng, Eddie Yan, Meghan Cowan, Haichen Shen, Leyuan Wang, Yuwei Hu, Luis Ceze, Carlos Guestrin, and Arvind Krishnamurthy. 2018. TVM: an automated end-to-end optimizing compiler for deep learning. In *Proceedings of the 13th USENIX Conference on Operating Systems Design and Implementation* (Carlsbad, CA, USA) *(OSDI'18)*. USENIX Association, USA, 579–594.

[14] Zilin Chen, Ambroise Lafont, Liam O'Connor, Gabriele Keller, Craig McLaughlin, Vincent Jackson, and Christine Rizkallah. 2023. Dargent: A Silver Bullet for Verified Data Layout Refinement. *Proc. ACM Program. Lang.* 7, POPL, Article 47 (2023), 27 pages. doi:10.1145/3571240

[15] Trishul M Chilimbi, Bob Davidson, and James R Larus. 1999. Cache-conscious structure definition. In *Proceedings of the ACM SIGPLAN 1999 conference on Programming language design and implementation*. 13–24. doi:10.1145/301618.301635

[16] Jongsok Choi, Ethan Che, Jason Ansel, and Oguz Ulgen. 2026. From Minutes to Seconds: LLM-Guided Autotuning for Helion Kernels. PyTorch Blog. https://pytorch.org/blog/from-minutes-to-seconds-llm-guided-autotuning-for-helion-kernels/

[17] Stephen Chou, Fredrik Kjolstad, and Saman Amarasinghe. 2018. Format Abstraction for Sparse Tensor Algebra Compilers. *Proc. ACM Program. Lang.* 2, OOPSLA, Article 123 (2018), 30 pages. doi:10.1145/3276493

[18] John Cocke. 1970. Global Common Subexpression Elimination. In *Proceedings of a Symposium on Compiler Optimization*. Association for Computing Machinery, New York, NY, USA, 20–24. doi:10.1145/800028.808480

[19] Thomas H. Cormen, Charles E. Leiserson, and Ronald L. Rivest. 1990. *Introduction to Algorithms*. MIT Press, Cambridge, MA, USA.

[20] Ron Cytron, Jeanne Ferrante, Barry K. Rosen, Mark N. Wegman, and F. Kenneth Zadeck. 1991. Efficiently Computing Static Single Assignment Form and the Control Dependence Graph. *ACM Trans. Program. Lang. Syst.* 13, 4 (Oct. 1991), 451–490. doi:10.1145/115372.115320

[21] Benjamin Delaware, Sorawit Suriyakarn, Clément Pit-Claudel, Qianchuan Ye, and Adam Chlipala. 2019. Narcissus: correct-by-construction derivation of decoders and encoders from binary formats. *Proc. ACM Program. Lang.* 3, ICFP,

Article 82 (July 2019), 29 pages. doi:10.1145/3341686
[22] Iavor S. Diatchki, Mark P. Jones, and Rebekah Leslie. 2005. High-level views on low-level representations. In *Proceedings of the Tenth ACM SIGPLAN International Conference on Functional Programming* (Tallinn, Estonia) *(ICFP '05)*. Association for Computing Machinery, New York, NY, USA, 168–179. doi:10.1145/1086365.1086387
[23] Paul Downen, Zena M. Ariola, Simon Peyton Jones, and Richard A. Eisenberg. 2020. Kinds are calling conventions. *Proc. ACM Program. Lang.* 4, ICFP, Article 104 (Aug. 2020), 29 pages. doi:10.1145/3408986
[24] Richard A. Eisenberg and Simon Peyton Jones. 2017. Levity Polymorphism. In *Proceedings of the 38th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI 2017)*. Association for Computing Machinery, New York, NY, USA, 525–539. doi:10.1145/3062341.3062357
[25] Martin Elsman. 2024. Double-Ended Bit-Stealing for Algebraic Data Types. *Proc. ACM Program. Lang.* 8, ICFP, Article 239 (Aug. 2024), 33 pages. doi:10.1145/3674628
[26] Kathleen Fisher and Robert Gruber. 2005. PADS: a domain-specific language for processing ad hoc data. *SIGPLAN Not.* 40, 6 (June 2005), 295–304. doi:10.1145/1064978.1065046
[27] Noah Goldstein. 2023. [PATCH v1] lib/rbtree: Use '+' instead of '|' for setting color. Linux kernel mailing list. https://lists.openwall.net/linux-kernel/2023/04/04/1522
[28] Christophe Gyurgyik, Alexander J Root, and Fredrik Kjolstad. 2026. Decoupling Data Layouts from Bounding Volume Hierarchies. *Proc. ACM Program. Lang.* 10, PLDI, Article 175 (June 2026), 26 pages. doi:10.1145/3808253
[29] Bastian Hagedorn, Archibald Samuel Elliott, Henrik Barthels, Rastislav Bodik, and Vinod Grover. 2020. Fireiron: A Data-Movement-Aware Scheduling Language for GPUs. In *Proceedings of the ACM International Conference on Parallel Architectures and Compilation Techniques* (Virtual Event, GA, USA) *(PACT '20)*. Association for Computing Machinery, New York, NY, USA, 71–82. doi:10.1145/3410463.3414632
[30] Bastian Hagedorn, Johannes Lenfers, Thomas Kœhler, Xueying Qin, Sergei Gorlatch, and Michel Steuwer. 2020. Achieving High-Performance the Functional Way: A Functional Pearl on Expressing High-Performance Optimizations as Rewrite Strategies. *Proc. ACM Program. Lang.* 4, ICFP, Article 92 (2020), 29 pages. doi:10.1145/3408974
[31] Cordelia V Hall. 1994. Using Hindley-Milner type inference to optimise list representation. In *Proceedings of the 1994 ACM conference on LISP and functional programming*. 162–172. doi:10.1145/182409.156781
[32] Mary Hall, Cosmin E. Oancea, Anne C. Elster, Ari Rasch, Sameeran Joshi, Amir Mohammad Tavakkoli, and Richard Schulze. 2025. Scheduling Language Chronology: Past, Present, and Future. *ACM Trans. Archit. Code Optim.* 22, 3, Article 100 (Sept. 2025), 31 pages. doi:10.1145/3743135
[33] Robert Harper and Greg Morrisett. 1995. Compiling polymorphism using intensional type analysis. In *Proceedings of the 22nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (San Francisco, California, USA) *(POPL '95)*. Association for Computing Machinery, New York, NY, USA, 130–141. doi:10.1145/199448.199475
[34] C. A. R. Hoare. 1972. Proof of correctness of data representations. *Acta Informatica* 1, 4 (1972), 271–281. doi:10.1007/BF00289507
[35] Yuka Ikarashi, Gilbert Louis Bernstein, Alex Reinking, Hasan Genc, and Jonathan Ragan-Kelley. 2022. Exocompilation for Productive Programming of Hardware Accelerators. In *Proceedings of the 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation (PLDI 2022)*. Association for Computing Machinery, New York, NY, USA, 703–718. doi:10.1145/3519939.3523446
[36] Yuka Ikarashi, Kevin Qian, Samir Droubi, Alex Reinking, Gilbert Louis Bernstein, and Jonathan Ragan-Kelley. 2025. Exo 2: Growing a Scheduling Language. In *Proceedings of the 30th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 1* (Rotterdam, Netherlands) *(ASPLOS '25)*. Association for Computing Machinery, New York, NY, USA, 426–444. doi:10.1145/3669940.3707218
[37] International Organization for Standardization 1999. *ISO/IEC 9899:1999: Programming Languages – C*. International Organization for Standardization, Geneva, Switzerland.
[38] Jane Street. 2025. OxCaml: Unboxed Types. https://oxcaml.org/documentation/unboxed-types/intro/. Accessed 2026-10-16.
[39] Neil D. Jones, Carsten K. Gomard, and Peter Sestoft. 1993. *Partial evaluation and automatic program generation*. Prentice-Hall, Inc., USA.
[40] G. Kahn. 1987. Natural semantics. In *STACS 87*, Franz J. Brandenburg, Guy Vidal-Naquet, and Martin Wirsing (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 22–39. doi:10.1007/BFb0039592
[41] Andrew Kennedy and Don Syme. 2001. Design and Implementation of Generics for the .NET Common Language Runtime. In *Proceedings of the ACM SIGPLAN 2001 Conference on Programming Language Design and Implementation (PLDI '01)*. Association for Computing Machinery, New York, NY, USA, 1–12. doi:10.1145/378795.378797
[42] Donald E. Knuth. 1998. *Sorting and Searching* (2nd ed.). The Art of Computer Programming, Vol. 3. Addison-Wesley, Reading, MA, USA.
[43] Chaitanya Koparkar, Mike Rainey, Michael Vollmer, Milind Kulkarni, and Ryan R. Newton. 2021. Efficient tree-traversals: reconciling parallelism and dense data representations. *Proc. ACM Program. Lang.* 5, ICFP, Article 91 (Aug.

2021), 29 pages. doi:10.1145/3473596

[44] Mahesh Lakshminarasimhan, Oscar Antepara, Tuowen Zhao, Benjamin Sepanski, Protonu Basu, Hans Johansen, Mary Hall, and Samuel Williams. 2024. Bricks: A high-performance portability layer for computations on block-structured grids. *Int. J. High Perform. Comput. Appl.* 38, 6 (Nov. 2024), 549–567. doi:10.1177/10943420241268288

[45] Chris Lattner and Vikram Adve. 2004. LLVM: A Compilation Framework for Lifelong Program Analysis & Transformation. In *Proceedings of the International Symposium on Code Generation and Optimization: Feedback-Directed and Runtime Optimization* (Palo Alto, California) *(CGO '04)*. IEEE Computer Society, USA, 75–86. doi:10.1109/CGO.2004.1281665

[46] Chris Lattner and Vikram S Adve. 2005. Transparent pointer compression for linked data structures. In *Proceedings of the 2005 workshop on Memory system performance*. 24–35. doi:10.1145/1111583.1111587

[47] Daan Leijen. 2014. Koka: Programming with Row Polymorphic Effect Types. In *Proceedings 5th Workshop on Mathematically Structured Functional Programming, Grenoble, France, 12 April 2014 (Electronic Proceedings in Theoretical Computer Science, Vol. 153)*, Paul Levy and Neel Krishnaswami (Eds.). Open Publishing Association, 100–126. doi:10.4204/EPTCS.153.8

[48] Daan Leijen and Anton Lorenzen. 2023. Tail Recursion Modulo Context: An Equational Approach. *Proc. ACM Program. Lang.* 7, POPL, Article 40 (Jan. 2023), 30 pages. doi:10.1145/3571233

[49] Daan Leijen, Benjamin Zorn, and Leonardo De Moura. 2019. Mimalloc: Free list sharding in action. In *Asian Symposium on Programming Languages and Systems*. Springer, 244–265. doi:10.1007/978-3-030-34175-6_13

[50] Xavier Leroy. 1992. Unboxed Objects and Polymorphic Typing. In *Proceedings of the 19th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '92)*. Association for Computing Machinery, New York, NY, USA, 177–188. doi:10.1145/143165.143205

[51] Xavier Leroy, Damien Doligez, Alain Frisch, Jacques Garrigue, Didier Rémy, KC Sivaramakrishnan, and Jérôme Vouillon. 2025. *The OCaml System Release 5.4: Documentation and User's Manual*. Inria. https://ocaml.org/manual/5.4/

[52] Michel Lespinasse. 2012. [PATCH v2 08/12] rbtree: low level optimizations in `rb_insert_color()`. Linux kernel mailing list. https://lkml.iu.edu/1207.1/03044.html

[53] Amanda Liu, Gilbert Louis Bernstein, Adam Chlipala, and Jonathan Ragan-Kelley. 2022. Verified Tensor-Program Optimization via High-Level Scheduling Rewrites. *Proc. ACM Program. Lang.* 6, POPL, Article 55 (2022), 28 pages. doi:10.1145/3498717

[54] Amanda Liu, Gilbert Louis Bernstein, Shoaib Kamil, Adam Chlipala, and Jonathan Ragan-Kelley. 2026. A Mechanized Algebra of Verified Data Structures for Optimizing Sparse Tensor Programs. *Proc. ACM Program. Lang.* 10, PLDI, Article 183 (June 2026), 23 pages. doi:10.1145/3808261

[55] Jie Liu, Zhongyuan Zhao, Zijian Ding, Benjamin Brock, Hongbo Rong, and Zhiru Zhang. 2024. UniSparse: An Intermediate Language for General Sparse Format Customization. *Proc. ACM Program. Lang.* 8, OOPSLA1, Article 99 (2024), 29 pages. doi:10.1145/3649816

[56] Anton Lorenzen and Daan Leijen. 2022. Reference Counting with Frame Limited Reuse. *Proc. ACM Program. Lang.* 6, ICFP, Article 93 (2022), 24 pages. doi:10.1145/3547634

[57] Anton Lorenzen, Daan Leijen, and Wouter Swierstra. 2023. FP^2: Fully in-Place Functional Programming. *Proc. ACM Program. Lang.* 7, ICFP, Article 198 (Aug. 2023), 30 pages. doi:10.1145/3607840

[58] Anton Lorenzen, Leo White, Stephen Dolan, Richard A. Eisenberg, and Sam Lindley. 2024. Oxidizing OCaml with Modal Memory Management. *Proc. ACM Program. Lang.* 8, ICFP, Article 253 (Aug. 2024), 30 pages. doi:10.1145/3674642

[59] Bruce T. Lowerre. 1976. *The HARPY Speech Recognition System*. Ph. D. Dissertation. Carnegie Mellon University. https://www.csd.cs.cmu.edu/academics/doctoral/degrees-conferred/bruce-lowerre

[60] Luc Maranget. 2008. Compiling pattern matching to good decision trees. In *Proceedings of the 2008 ACM SIGPLAN Workshop on ML* (Victoria, BC, Canada) *(ML '08)*. Association for Computing Machinery, New York, NY, USA, 35–46. doi:10.1145/1411304.1411311

[61] Simon Marlow, Alexey Rodriguez Yakushev, and Simon Peyton Jones. 2007. Faster Laziness Using Dynamic Pointer Tagging. In *Proceedings of the 12th ACM SIGPLAN International Conference on Functional Programming (ICFP '07)*. Association for Computing Machinery, New York, NY, USA, 277–288. doi:10.1145/1291151.1291194

[62] Erik Meijer, Maarten Fokkinga, and Ross Paterson. 1991. Functional programming with bananas, lenses, envelopes and barbed wire. In *Conference on functional programming languages and computer architecture*. Springer, 124–144. doi:10.1007/3540543961_7

[63] Andrew W. Moore, Andy J. Connolly, Chris Genovese, Alex Gray, Larry Grone, Nick Kanidoris II, Robert C. Nichol, Jeff Schneider, Alex S. Szalay, Istvan Szapudi, et al. 2001. Fast algorithms and efficient statistics: N-point correlation functions. In *Mining the Sky: Proceedings of the MPA/ESO/MPE Workshop Held in Garching, Germany, July 31–August 4, 2000*. Springer, 71–82. doi:10.1007/10849171_5

[64] Leonardo de Moura and Sebastian Ullrich. 2021. The Lean 4 Theorem Prover and Programming Language. In *Automated Deduction – CADE 28: 28th International Conference on Automated Deduction, Virtual Event, July 12–15,*

2021, Proceedings. Springer-Verlag, Berlin, Heidelberg, 625–635. doi:10.1007/978-3-030-79876-5_37

[65] Chris Okasaki. 1999. Red-Black Trees in a Functional Setting. *Journal of Functional Programming* 9, 4 (1999), 471–477. doi:10.1017/S0956796899003494

[66] OpenAI. 2026. ChatGPT. https://chatgpt.com. Large language model (GPT-6.1)..

[67] OpenJDK. 2025. `java.util.HashMap`. OpenJDK JDK 25. `src/java.base/share/classes/java/util/HashMap.java`. https://github.com/openjdk/jdk/blob/jdk-25-ga/src/java.base/share/classes/java/util/HashMap.java

[68] Adam Paszke, Daniel D. Johnson, David Duvenaud, Dimitrios Vytiniotis, Alexey Radul, Matthew J. Johnson, Jonathan Ragan-Kelley, and Dougal Maclaurin. 2021. Getting to the Point: Index Sets and Parallelism-Preserving Autodiff for Pointful Array Programming. *Proc. ACM Program. Lang.* 5, ICFP, Article 88 (Aug. 2021), 29 pages. doi:10.1145/3473593

[69] Mikael Pettersson. 1992. A term pattern-match compiler inspired by finite automata theory. In *International Conference on Compiler Construction*. Springer, 258–270. doi:10.1007/3-540-55984-1_24

[70] Simon L. Peyton Jones and John Launchbury. 1991. Unboxed Values as First Class Citizens in a Non-Strict Functional Language. In *Functional Programming Languages and Computer Architecture (Lecture Notes in Computer Science, Vol. 523)*. Springer, Berlin, Heidelberg, 636–666. doi:10.1007/3540543961_30

[71] Jonathan Ragan-Kelley, Connelly Barnes, Andrew Adams, Sylvain Paris, Frédo Durand, and Saman Amarasinghe. 2013. Halide: A Language and Compiler for Optimizing Parallelism, Locality, and Recomputation in Image Processing Pipelines. In *Proceedings of the 34th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI '13)*. Association for Computing Machinery, New York, NY, USA, 519–530. doi:10.1145/2491956.2462176

[72] Tahina Ramananandro, Antoine Delignat-Lavaud, Cédric Fournet, Nikhil Swamy, Tej Chajed, Nadim Kobeissi, and Jonathan Protzenko. 2019. EverParse: Verified Secure Zero-Copy Parsers for Authenticated Message Formats. In *28th USENIX Security Symposium (USENIX Security 19)*. USENIX Association, Santa Clara, CA, 1465–1482. https://www.usenix.org/conference/usenixsecurity19/presentation/delignat-lavaud

[73] Alex Reinking, Ningning Xie, Leonardo de Moura, and Daan Leijen. 2021. Perceus: Garbage Free Reference Counting with Reuse. In *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation (PLDI 2021)*. Association for Computing Machinery, New York, NY, USA, 96–111. doi:10.1145/3453483.3454032

[74] Alexander J Root, Christophe Gyurgyik, Purvi Goel, Kayvon Fatahalian, Jonathan Ragan-Kelley, Andrew Adams, and Fredrik Kjolstad. 2026. Bonsai: Compiling Queries to Pruned Tree Traversals. *Proc. ACM Program. Lang.* 10, PLDI, Article 178 (June 2026), 27 pages. doi:10.1145/3808256

[75] Barry K. Rosen, Mark N. Wegman, and F. Kenneth Zadeck. 1988. Global Value Numbers and Redundant Computations. In *Proceedings of the 15th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '88)*. Association for Computing Machinery, New York, NY, USA, 12–27. doi:10.1145/73560.73562

[76] Kevin Scott and Norman Ramsey. 2000. *When Do Match-compilation Heuristics Matter?* Technical Report. USA. doi:10.18130/v3gb4m

[77] Ryan Senanayake, Changwan Hong, Ziheng Wang, Amalee Wilson, Stephen Chou, Shoaib Kamil, Saman Amarasinghe, and Fredrik Kjolstad. 2020. A sparse iteration space transformation framework for sparse tensor algebra. *Proc. ACM Program. Lang.* 4, OOPSLA, Article 158 (Nov. 2020), 30 pages. doi:10.1145/3428226

[78] Zhong Shao, John H. Reppy, and Andrew W. Appel. 1994. Unrolling lists. In *Proceedings of the 1994 ACM Conference on LISP and Functional Programming* (Orlando, Florida, USA) *(LFP '94)*. Association for Computing Machinery, New York, NY, USA, 185–195. doi:10.1145/182409.182453

[79] Jun Shirako and Vivek Sarkar. 2022. An Affine Scheduling Framework for Integrating Data Layout and Loop Transformations. In *Languages and Compilers for Parallel Computing: 33rd International Workshop, LCPC 2020, Virtual Event, October 14-16, 2020, Revised Selected Papers*. Springer-Verlag, Berlin, Heidelberg, 3–19. doi:10.1007/978-3-030-95953-1_1

[80] Vidush Singhal, Chaitanya Koparkar, Joseph Zullo, Artem Pelenitsyn, Michael Vollmer, Mike Rainey, Ryan Newton, and Milind Kulkarni. 2024. Optimizing Layout of Recursive Datatypes with Marmoset: Or, Algorithms + Data Layouts = Efficient Programs. In *38th European Conference on Object-Oriented Programming (ECOOP 2024) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 313)*, Jonathan Aldrich and Guido Salvaneschi (Eds.). Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 38:1–38:28. doi:10.4230/LIPIcs.ECOOP.2024.38

[81] Michel Steuwer, Toomas Remmelg, and Christophe Dubach. 2017. LIFT: A functional data-parallel IR for high-performance GPU code generation. In *2017 IEEE/ACM International Symposium on Code Generation and Optimization (CGO)*. 74–85. doi:10.1109/CGO.2017.7863730

[82] Nikhil Swamy, Tahina Ramananandro, Aseem Rastogi, Irina Spiridonova, Haobin Ni, Dmitry Malloy, Juan Vazquez, Michael Tang, Omar Cardona, and Arti Gupta. 2022. Hardening attack surfaces with formally proven binary format parsers. In *Proceedings of the 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation (PLDI 2022)*. Association for Computing Machinery, New York, NY, USA, 31–45. doi:10.1145/3519939.3523708

[83] Amir Mohammad Tavakkoli, Cosmin E. Oancea, and Mary Hall. 2026. LEGO: A Layout Expression Language for Code Generation of Hierarchical Mapping. In *2026 IEEE/ACM International Symposium on Code Generation and Optimization (CGO)*. 228–241. doi:10.1109/CGO68049.2026.11394846

[84] The Linux Kernel Documentation. 2026. 5-Level Paging. The Linux Kernel Documentation, x86-64. Accessed: 2026-10-28. https://docs.kernel.org/arch/x86/x86_64/5level-paging.html

[85] The Rust Project Developers. 2026. Attribute `repr`. The Rust Standard Library Documentation. Nightly documentation. Accessed: 2026-10-23. https://doc.rust-lang.org/nightly/std/attribute.repr.html

[86] The Rust Project Developers. 2026. `repr(Rust)`. The Rustonomicon. Accessed: 2026-10-23. https://doc.rust-lang.org/nomicon/repr-rust.html

[87] Linus Torvalds. 2026. *Red-Black Trees in Linux*. The Linux Kernel, version 7.2. `include/linux/rbtree_types.h`. https://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git/tree/include/linux/rbtree_types.h?h=v7.2

[88] Emma Tye. 2026. A Rig of Transformations. Extended abstract presented at the Workshop on Type-Driven Development (TyDe 2026). arXiv:2608.12409 [cs.PL] doi:10.48550/arXiv.2608.12409

[89] Sebastian Ullrich and Leonardo de Moura. 2019. Counting immutable beans: Reference counting optimized for purely functional programming. In *Proceedings of the 31st Symposium on Implementation and Application of Functional Languages*. 1–12. doi:10.1145/3412932.3412935

[90] Marcell van Geest and Wouter Swierstra. 2017. Generic packet descriptions: verified parsing and pretty printing of low-level data. In *Proceedings of the 2nd ACM SIGPLAN International Workshop on Type-Driven Development* (Oxford, UK) *(TyDe 2017)*. Association for Computing Machinery, New York, NY, USA, 30–40. doi:10.1145/3122975.3122979

[91] Anand Venkat, Mary Hall, and Michelle Strout. 2015. Loop and data transformations for sparse matrix code. In *Proceedings of the 36th ACM SIGPLAN Conference on Programming Language Design and Implementation* (Portland, OR, USA) *(PLDI '15)*. Association for Computing Machinery, New York, NY, USA, 521–532. doi:10.1145/2737924.2738003

[92] Michael Vollmer, Chaitanya Koparkar, Mike Rainey, Laith Sakka, Milind Kulkarni, and Ryan R. Newton. 2019. LoCal: a language for programs operating on serialized data. In *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation* (Phoenix, AZ, USA) *(PLDI 2019)*. Association for Computing Machinery, New York, NY, USA, 48–62. doi:10.1145/3314221.3314631

[93] Michael Vollmer, Sarah Spall, Buddhika Chamith, Laith Sakka, Chaitanya Koparkar, Milind Kulkarni, Sam Tobin-Hochstadt, and Ryan R. Newton. 2017. Compiling Tree Transforms to Operate on Packed Representations. In *31st European Conference on Object-Oriented Programming (ECOOP 2017) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 74)*, Peter Müller (Ed.). Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 26:1–26:29. doi:10.4230/LIPIcs.ECOOP.2017.26

[94] Philip Wadler. 1990. Deforestation: transforming programs to eliminate trees. *Theoretical Computer Science* 73, 2 (1990), 231–248. doi:10.1016/0304-3975(90)90147-A

[95] Stephen Weeks. 2006. Whole-Program Compilation in MLton. In *Proceedings of the 2006 Workshop on ML (ML '06)*. Association for Computing Machinery, New York, NY, USA, 1. doi:10.1145/1159876.1159877

[96] Christian Wimmer and Hanspeter Mössenböck. 2008. Automatic array inlining in java virtual machines. In *Proceedings of the 6th Annual IEEE/ACM International Symposium on Code Generation and Optimization* (Boston, MA, USA) *(CGO '08)*. Association for Computing Machinery, New York, NY, USA, 14–23. doi:10.1145/1356058.1356061

[97] Sven Woop, Carsten Benthin, Ingo Wald, Gregory S Johnson, and Eric Tabellion. 2014. Exploiting Local Orientation Similarity for Efficient Ray Traversal of Hair and Fur. In *Eurographics/ACM SIGGRAPH Symposium on High Performance Graphics*. The Eurographics Association. doi:10.2312/hpg.20141092

[98] Rohan Yadav, Alex Aiken, and Fredrik Kjolstad. 2022. DISTAL: the distributed tensor algebra compiler. In *Proceedings of the 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation* (San Diego, CA, USA) *(PLDI 2022)*. Association for Computing Machinery, New York, NY, USA, 286–300. doi:10.1145/3519939.3523437

[99] Rohan Yadav, Alex Aiken, and Fredrik Kjolstad. 2022. SpDISTAL: compiling distributed sparse tensor computations. In *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis* (Dallas, Texas) *(SC '22)*. IEEE Press, Article 59, 15 pages. doi:10.1109/SC41404.2022.00064

[100] Qianchuan Ye and Benjamin Delaware. 2019. A verified protocol buffer compiler. In *Proceedings of the 8th ACM SIGPLAN International Conference on Certified Programs and Proofs* (Cascais, Portugal) *(CPP 2019)*. Association for Computing Machinery, New York, NY, USA, 222–233. doi:10.1145/3293880.3294105

[101] Henri Ylitie, Tero Karras, and Samuli Laine. 2017. Efficient incoherent ray traversal on GPUs through compressed wide BVHs. In *Proceedings of High Performance Graphics*. 1–13. doi:10.1145/3105762.3105773

[102] Keren Zhou, Mario Lezcano-Casado, Adam P. Goucher, Akhmed Rakhmati, Jeff Niu, Justin Lebar, Pawel Szczerbuk, Peter Bell, Phil Tillet, Thomas Raoux, and Zahi Moudallal. 2025. Linear Layouts: Robust Code Generation of Efficient Tensor Computation Using F_2. In *Proceedings of the 31st ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 1* (Pittsburgh, PA, USA) *(ASPLOS '26)*. Association for Computing Machinery, New York, NY, USA, 132–146. doi:10.1145/3760250.3762221

A Provenance Lattice

Figure 12 shows the depth-one provenance lattice of the red-black tree's **Node** and **Leaf** constructors. Equality of provenances is decidable (Lean: `Prov.decidableEquiv`); two provenances read finitely many paths, constructors, and literals, and projecting a value onto these reads preserves both.

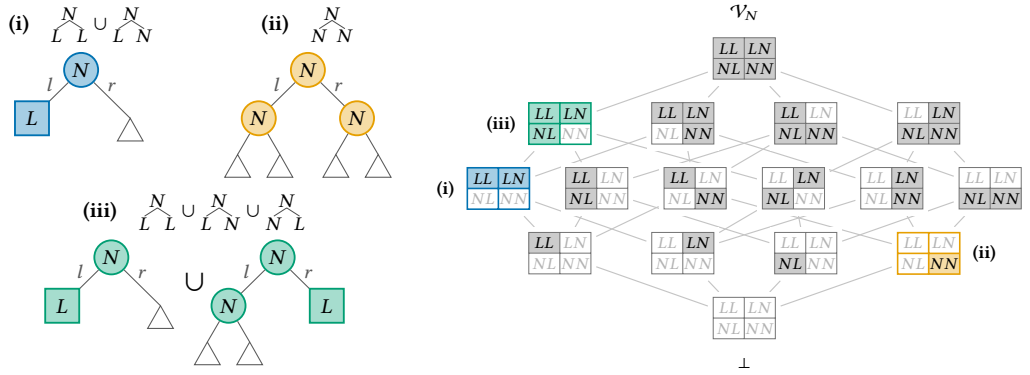


Fig. 12. Constructor provenance of the children l and r of a **Node** of the red-black Tree, abbreviating **Leaf** as L and **Node** as N. Triangles denote arbitrary subtrees. Each matrix shades the admitted child pairs, its cell AB standing for children A and B, and edges order the matrices by inclusion. The examples, which require (i) a left leaf, (ii) two node children, or (iii) at least one leaf child, share the colors of their matrices.

B Well-Formedness

The judgment $A \vdash \ell$ wf checks a layout ℓ on its admitted subjects A, the source values that can reach ℓ (Section 4.2). A layout definition for T is verified on all values of its type, $\mathcal{V}_{\mathsf{T}}$, and each rule checks every immediate sublayout on the subjects that can reach it. The judgment therefore guarantees every admitted subject has at least one representation (*totality*, Lemma 3). The rules of Figure 13 define the well-formedness judgment of Section 5.3, which every rewrite judgment of Appendix C requires of both sides.

$$\frac{\ell \in \{\epsilon, \mathsf{pad}_n, \mathsf{zero}_n\}}{A \vdash \ell \ \mathsf{wf}} \qquad \frac{A \subseteq \mathcal{V}_{\mathsf{u}_n}}{A \vdash \mathsf{u}_n \ \mathsf{wf}} \qquad \frac{A \subseteq \mathrm{dom}[\![e]\!]}{A \vdash \langle e \rangle \ \mathsf{wf}}$$

$$\frac{A \subseteq \mathrm{dom}\,\pi \qquad A.\pi \vdash \ell \ \mathsf{wf}}{A \vdash \pi : \ell \ \mathsf{wf}} \qquad \frac{A \vdash \ell \ \mathsf{wf}}{A \vdash \&_n(\ell) \ \mathsf{wf}} \qquad \frac{A \vdash \ell_0 \ \mathsf{wf} \qquad A \vdash \ell_1 \ \mathsf{wf}}{A \vdash \ell_0 \times \ell_1 \ \mathsf{wf}} \qquad \frac{\mathsf{T} \in \mathrm{dom}\,\Delta \qquad A \subseteq \mathcal{V}_{\mathsf{T}}}{A \vdash \mathsf{ref}\,\mathsf{T} \ \mathsf{wf}}$$

$$\frac{A \subseteq \bigcup_j \phi_j \qquad A \cap \phi_j \vdash \ell_j \ \mathsf{wf}}{A \vdash \bigcup\{\phi_j \Rightarrow \ell_j\} \ \mathsf{wf}} \qquad \frac{A \subseteq \bigcup_i A_{\pi=c_i} \qquad A_{\pi=c_i} \subseteq \phi_i \qquad A_{\pi=c_i} \vdash \ell_i \ \mathsf{wf}}{A \vdash \bigsqcup^{\pi}\{(c_i, \phi_i) \Rightarrow \ell_i\} \ \mathsf{wf}}$$

$$\frac{\forall \mathsf{T} \in \mathrm{dom}\,\Delta,\ \mathcal{V}_{\mathsf{T}} \vdash \Delta(\mathsf{T}) \ \mathsf{wf}}{\Delta \ \mathsf{wf}}$$

Fig. 13. Well-formedness on the admitted subjects A, under the layout environment Δ. We write $\mathrm{dom}[\![e]\!]$ for the values on which e is defined and $\mathrm{dom}\,\pi$ for those on which $v.\pi$ is defined. A named layout passes the selected subjects $A.\pi$ to its payload. We write $A_{\pi=c_i}$ for the admitted subjects whose key π is defined and equals c_i. In the mechanization, a disjoint sum keyed by a tag takes its guards from the tag's declaration, so $A_{\pi=c_i} \subseteq \phi_i$ follows from the selector's soundness. Lean: `WF, WF.{sumMap_iff, taggedMap_iff}, LayoutEnv.WF`

C Rewrite Judgments

Figure 14 states the rules of Figure 2 as judgments $A \vdash_\Delta \ell_0 \longrightarrow \ell_1$ on admitted subjects A under the layout environment Δ, each establishing local refinement (Definition 4). A judgment with $\longleftrightarrow$ abbreviates both directions. Besides the displayed premises, each rule requires both sides to be well formed on A (Appendix B). Here $j, k, m, n \in \mathbb{N}$. Ω and Ψ are arbitrary lists of arms, and an indexed arm $\phi_j \Rightarrow \ell_j$ stands for a list of arms indexed from 0. In tag, the arms may be listed in any order, and $\mathrm{guards}(t)$ lists their guards in that order (Section 4).

$$\frac{}{A \vdash_\Delta \ell \longrightarrow \mathsf{pad}_n \times \ell}\ \mathsf{pad}_n \qquad \frac{}{A \vdash_\Delta \ell \longrightarrow \mathsf{zero}_n \times \ell}\ \mathsf{zero}_n \qquad \frac{}{A \vdash_\Delta \ell \longrightarrow \&_n(\epsilon) \times \ell}\ \mathsf{ptr}_n$$

$$\frac{}{A \vdash_\Delta \ell \longrightarrow \langle e \rangle \times \ell}\ \mathsf{memoize}_e \qquad \frac{}{A \vdash_\Delta \&_n(\ell_0) \times \ell_1 \longrightarrow \&_n(\ell_0 \times \ell_1)}\ \mathsf{box} \qquad \frac{}{A \vdash_\Delta \&_n(\ell) \longrightarrow \ell}\ \mathsf{unbox}$$

$$\frac{}{A \vdash_\Delta \ell_0 \times \ell_1 \longrightarrow \ell_1 \times \ell_0}\ \mathsf{commute} \qquad \frac{}{A \vdash_\Delta (\ell_0 \times \ell_1) \times \ell_2 \longleftrightarrow \ell_0 \times (\ell_1 \times \ell_2)}\ \mathsf{assoc} \qquad \frac{}{A \vdash_\Delta \epsilon \times \ell \longleftrightarrow \ell}\ \mathsf{unit}$$

$$\frac{k \le n}{A \vdash_\Delta \pi : \mathsf{u}_n \longrightarrow \pi[0,k) : \mathsf{u}_k \times \pi[k,n) : \mathsf{u}_{n-k}}\ \mathsf{split}_k \qquad \frac{}{A \vdash_\Delta \bigcup\{\phi \Rightarrow \ell\} \longrightarrow \ell}\ \mathsf{collapse}$$

$$\frac{}{A \vdash_\Delta \mathsf{u}_n \longrightarrow \bigcup\{\mathcal{V}_{\mathsf{u}_m} \Rightarrow \mathsf{u}_m, \top \Rightarrow \mathsf{u}_n\}}\ \mathsf{refine}_{\mathsf{u}_m} \qquad \frac{}{A \vdash_\Delta \ell \longrightarrow \bigcup\{\phi \Rightarrow \ell, \top \Rightarrow \ell\}}\ \mathsf{refine}_\phi$$

$$\frac{}{A \vdash_\Delta \ell \times \bigcup\{\Omega\} \longrightarrow \bigcup\{\ell \times \Omega\}}\ \mathsf{distribute} \qquad \frac{\Omega \neq \varnothing}{A \vdash_\Delta \bigcup\{\ell \times \Omega\} \longrightarrow \ell \times \bigcup\{\Omega\}}\ \mathsf{factor}$$

$$\frac{A \subseteq \bigcup_j \phi_j}{A \vdash_\Delta \bigcup\{\phi_j \Rightarrow \ell_j, \phi \Rightarrow \ell\} \longrightarrow \bigcup\{\phi_j \Rightarrow \ell_j\}}\ \mathsf{prune} \qquad \frac{}{A \vdash_\Delta \bigcup\{\Omega, \phi \Rightarrow \bigcup\{\Psi\}\} \longrightarrow \bigcup\{\Omega, \phi \cap \Psi\}}\ \mathsf{flatten}$$

$$\frac{\Delta \ \mathsf{wf}\ \text{(Section 5.3)}}{A \vdash_\Delta \mathsf{ref}\,\mathsf{T} \longrightarrow \Delta(\mathsf{T})}\ \mathsf{unfold} \qquad \frac{\mathrm{guards}(t) = (\phi_j)}{A \vdash_\Delta \bigcup\{\phi_j \Rightarrow \ell_j\} \longrightarrow t : \mathsf{u}_{|t|} \times \bigsqcup^t\{(j, \phi_j) \Rightarrow \ell_j\}}\ \mathsf{tag}$$

Fig. 14. Rewrite judgments, named as in Figure 2. Lean: `Core.Step`

D Layout Interpretation

Consider $\mathsf{T} = \mathbf{A}(x : \mathsf{u}_1,\ y : \mathsf{u}_1) \mid \mathbf{B}(y : \mathsf{u}_1,\ x : \mathsf{u}_1)$, whose two constructors declare the same one-bit fields but in opposite orders. The derivation below computes the structural representations of $\mathbf{B}(0,1)$ under a rewritten variant of the canonical layout that stores its field x behind a two-bit pointer. Since this layout refers to no type, we interpret it in the empty environment $\varnothing$. We use **blue** to denote an evaluation step.

$$
\begin{aligned}
&\mathcal{W}^{\varnothing}_{\bigcup\{\mathbf{A} \Rightarrow x : \mathsf{u}_1 \times y : \mathsf{u}_1,\ \mathbf{B} \Rightarrow y : \mathsf{u}_1 \times \&_2(x : \mathsf{u}_1)\}}(\mathbf{B}(0,1)) \\
&= \left\{\mathbf{cons}_{\mathbf{B}}(\eta) \mid \eta \in \mathcal{W}^{\varnothing}_{y:\mathsf{u}_1 \times \&_2(x:\mathsf{u}_1)}(\mathbf{B}(0,1))\right\} \\
&\quad \left\{\mathrm{cons}_{\mathbf{B}}(\eta) \mid \eta \in \mathcal{W}^{\varnothing}_{y:\mathsf{u}_1 \times \&_2(x:\mathsf{u}_1)}(\mathbf{B}(0,1))\right\} \\
&= \left\{\mathrm{cons}_{\mathbf{B}}(\eta) \mid \eta \in \mathcal{W}^{\varnothing}_{y:\mathsf{u}_1}(\mathbf{B}(0,1)) \times \mathcal{W}^{\varnothing}_{\&_2(x:\mathsf{u}_1)}(\mathbf{B}(0,1))\right\} \\
&\quad \left\{\mathrm{cons}_{\mathbf{B}}(\eta) \mid \eta \in \mathcal{W}^{\varnothing}_{y:\mathsf{u}_1}(\mathbf{B}(0,1)) \times \mathcal{W}^{\varnothing}_{\&_2(x:\mathsf{u}_1)}(\mathbf{B}(0,1))\right\} \\
&= \left\{\mathrm{cons}_{\mathbf{B}}(\eta) \mid \eta \in \{0\} \times \mathcal{W}^{\varnothing}_{\&_2(x:\mathsf{u}_1)}(\mathbf{B}(0,1))\right\} \\
&\quad \left\{\mathrm{cons}_{\mathbf{B}}(\eta) \mid \eta \in \{0\} \times \mathcal{W}^{\varnothing}_{\&_2(x:\mathsf{u}_1)}(\mathbf{B}(0,1))\right\} \\
&= \left\{\mathrm{cons}_{\mathbf{B}}(\eta) \mid \eta \in \{0\} \times \{\mathbf{ptr}(\alpha, k) \mid \alpha \in \{0,1,2,3\},\ k \in \mathcal{W}^{\varnothing}_{x:\mathsf{u}_1}(\mathbf{B}(0,1))\}\right\} \\
&\quad \left\{\mathrm{cons}_{\mathbf{B}}(\eta) \mid \eta \in \{0\} \times \{\mathbf{ptr}(\alpha, k) \mid \alpha \in \{0,1,2,3\},\ k \in \mathcal{W}^{\varnothing}_{x:\mathsf{u}_1}(\mathbf{B}(0,1))\}\right\} \\
&= \left\{\mathrm{cons}_{\mathbf{B}}(\eta) \mid \eta \in \{0\} \times \{\mathbf{ptr}(0,1),\ \mathbf{ptr}(1,1),\ \mathbf{ptr}(2,1),\ \mathbf{ptr}(3,1)\}\right\} \\
&\quad \left\{\mathbf{cons}_{\mathbf{B}}(\eta) \mid \eta \in \{0\} \times \{\mathbf{ptr}(0,1),\ \mathbf{ptr}(1,1),\ \mathbf{ptr}(2,1),\ \mathbf{ptr}(3,1)\}\right\} \\
&= \begin{Bmatrix} \mathbf{cons}_{\mathbf{B}}(0, \mathbf{ptr}(0,1)), & \mathbf{cons}_{\mathbf{B}}(0, \mathbf{ptr}(1,1)), \\ \mathbf{cons}_{\mathbf{B}}(0, \mathbf{ptr}(2,1)), & \mathbf{cons}_{\mathbf{B}}(0, \mathbf{ptr}(3,1)) \end{Bmatrix}
\end{aligned}
$$

Fig. 15. The structural representations of $\mathbf{B}(0,1)$, computed one evaluation step at a time.

E Producer and Consumer Specialization

Specialization makes matches read and constructions write the chosen layouts before lowering (Section 8.2). A binder is *inline* when its fields lie directly in the cell, and *stored* when the cell holds it as a scalar or pointer. A construction makes an arm *selectable* when its explicit constructors and literals satisfy the arm's guard. We write π for a path and κ for the known provenance (Sections 4.1 and 4.2).

In the running example, the layout unfolds a list's tail into the head's cell, giving arms **Nil**, **Cons**$(h, \textbf{Nil})$, and **Cons**$(h, \textbf{Cons}(h', t'))$ with t' behind a pointer. Specialization rewrites `map` on the left into the function on the right.

```
fn map(xs: List): List =
  match xs {
    Nil => Nil,
    Cons(h, t) => Cons(h + 1, map(t))
  }
```

```
fn map(xs: List): List =
  match xs {
    Nil => Nil,
    Cons(h, Nil) => Cons(h + 1, Nil),
    Cons(h, Cons(h', t')) =>
      Cons(h + 1, Cons(h' + 1, map(t')))
  }
```

Consumer specialization. Algorithm 1 compiles matches into decision trees against the layout [60]. It tests positions demanded by the first live clause in position order, extending κ without a test when only one layout arm remains. Reify groups tests within a cell into sibling arms; tests behind a pointer become a nested match. Resolve recompiles matches on inline binders at their paths, inlining one step of a callee that matches the binder. Thus `map(t)` becomes `Cons(h' + 1, map(t'))`; the call on the stored `t'` remains.

Producer specialization. Algorithm 2 exposes constructors needed to select an arm. Expand inlines a call in a constructor argument and moves the enclosing construction into each resulting match branch. For example, **Cons**$(n, \mathit{range}(n-1))$ cannot select an arm until the tail's constructor is known. One step yields **Cons**$(n, \textbf{Nil})$ and **Cons**$(n, \textbf{Cons}(n-1, \mathit{range}(n-2)))$, each selecting an arm. Expansion may continue toward a more specific arm; otherwise, Dispatch tests the child's constructor and respells it in each branch. Stored pointers pass through under aliases recording their constructors. Dispatch requires inline children to decompose into scalars, pointers, or inline constructors that have no further dispatch.

Bounds. Inlining a recursive call may expose another call without revealing the constructor the layout needs. To ensure compilation terminates, a *budget* bounds this compile-time work, not runtime recursion or input size. Expansion spends one unit per inlined call, starting from $B_\ell = d_\ell(c+1)$. Here d_ℓ is the largest number of positions within a value that any one guard tests, excluding the root. In the list layout, the only test beyond the root asks whether the tail is **Nil** or **Cons**, so $d_\ell = 1$. The parameter c is the size of the largest group of functions connected by tail calls outside constructions. The factor $c+1$ allows c calls that build nothing plus one that exposes a constructor. Let-bound values restart with their own layout's budget. When expansion runs out, dispatch may expose the missing constructors by runtime tests; it starts from $Q_\ell = d_\ell + 1$ and spends one unit per child destruction. If dispatch cannot supply the fields safely, compilation rejects the construction.

Notably, these bounds can reject valid programs. For example, suppose *drop* traverses a list by tail calls before returning **Nil**. Expanding **Cons**$(h, \mathit{drop}(t))$ exhausts the budget while one branch still contains a recursive call. Under the two-element layout, dispatch cannot safely extract the packed child's fields, so compilation rejects the construction.

Algorithm 1 Consumer specialization.

function SpecializeMatch(clauses C, path π) ▷ in the current cell ℓ, under facts κ
  return Reify(the decision tree of C at π, against ℓ under κ)

function Reify(T)
  if T is a leaf (e, σ) **then return** Resolve(e, σ) ▷ σ: binders to paths
  return a match with one reified arm per outcome of T

function Resolve(e, σ)
  if e = match y $\{\ldots\}$ and y is inline **then**
    return SpecializeMatch(the clauses of e, $\sigma(y)$)
  else if $e = g(\ldots, y, \ldots)$, the body of g matches that parameter, and y is inline **then**
    return SpecializeMatch(the clauses of g, $\sigma(y)$) ▷ inline one step of g
  else if $e = y$ and y is stored **then**
    return the field or pointer that stores y
  else if $e = y$ and y is inline as $\mathbf{K}(a_1, \ldots, a_n)$ with every a_i stored or constant **then**
    return $\mathbf{K}(a_1, \ldots, a_n)$ ▷ rebuild one constructor
  else if $e = y$ **then**
    reject ▷ rebuilding y would copy its structure
  end if
  return e with Resolve applied to each subexpression

Algorithm 2 Producer specialization.

function Expand(e, cell layout ℓ, budget b) ▷ initially $b = B_\ell$
  if e = match s $\{p_i \Rightarrow e_i\}_i$ **then**
    return match s $\{p_i \Rightarrow$ Expand$(e_i, \ell, b)\}_i$
  else if e = let $x : \ell_x = v$ in e' **then**
    return let $x : \ell_x =$ Expand(v, ℓ_x, B_{ℓ_x}) in Expand(e', ℓ, b)
  else if e is a construction and some arm of ℓ is selectable **then**
    return e
  else if $e = \mathbf{K}(\ldots, g(\bar{u}), \ldots)$ for a known function g **then**
    if $b > 0$ **then**
      $e' \leftarrow e$ with g inlined and **K** pushed into the arms of its match
      return Expand$(e', \ell, b - 1)$
    else if Dispatch cannot expose the hidden constructor **then**
      reject
    end if
  end if
  return e

function Dispatch(e, cell layout ℓ, budget q) ▷ initially $q = Q_\ell$
  if $q = 0$, e is not a construction, or an arm is selectable **then return** e
  $i \leftarrow$ the first argument whose constructor a guard of ℓ constrains
  for all constructors **C** of e_i's type **do**
    $e_{\mathbf{C}} \leftarrow$ Dispatch$(e[i \leftarrow \mathbf{C}(\bar{y})], \ell, q - 1)$ ▷ $\bar{y}$ fresh
  end for
  return match e_i $\{\mathbf{C}(\bar{y}) \Rightarrow e_{\mathbf{C}}\}_{\mathbf{C}}$

F Physical Encoding and Decoding

The encoder and decoder of Section 7.1 satisfy the following sufficient conditions.

F.1 Discriminant Closure

The judgment $\Gamma \vdash \ell$ disc checks that each tag is stored whole in its scope. Γ records available scalar paths, not run-time values. Write $|\pi|$ for a path's declared scalar width. The paths stored whole unconditionally in ℓ are (Lean: `Stored`):

$$\begin{aligned}
\mathcal{S}(\pi : \mathsf{u}_n) &= \begin{cases} \{\pi\}, & |\pi| \text{ is defined and at most } n, \\ \varnothing, & \text{otherwise,} \end{cases} \\
\mathcal{S}(\&_n(\ell)) &= \mathcal{S}(\ell), \\
\mathcal{S}(\ell_0 \times \ell_1) &= \mathcal{S}(\ell_0) \cup \mathcal{S}(\ell_1), \\
\mathcal{S}(\ell) &= \varnothing \qquad \text{for all other forms.}
\end{aligned}$$

At a product $\ell_0 \times \ell_1$, check both components with $\Gamma \cup \mathcal{S}(\ell_0 \times \ell_1)$. Pointers preserve Γ; names reset it to $\varnothing$. Check every sum arm with Γ, additionally requiring a tagged sum's tag path to belong to Γ. Leaves need no further check; reference bodies are checked separately from $\varnothing$. The forward variant $\Gamma \vdash \ell$ $\mathsf{disc}_{\rightarrow}$ instead checks ℓ_0 with Γ and ℓ_1 with $\Gamma \cup \mathcal{S}(\ell_0)$, so a tag must precede its payload. Importantly, other decoding orders could relax this requirement when tag locations can be determined without first decoding the payloads they select.

F.2 Encoding and Decoding Guarantees

The encoder emits scalars at their widths, most significant bit first, concatenates product components, selects disjoint-sum arms by their tags, and stores each pointer as an address whose heap block contains the pointee. The decoder reads components from left to right. Its sufficient condition is the forward judgment $\varnothing \vdash \ell$ $\mathsf{disc}_{\rightarrow}$ of Appendix F.1, where each tag field must be stored whole before its payload in the same scope. Both functions take a recursion bound n (fuel), written encode_ℓ^n and decode_ℓ^n.

A layout is *concrete* for τ when the representations of $\mathcal{V}_\tau$ contain no markers of ordinary sums. For a minimum pointer width w, a value v satisfies the *address-capacity bound* when some $\eta_0 \in \mathcal{W}_\ell^\Delta(v)$ has $|\eta_0|_{\mathcal{P}} \le 2^w$, where $|\eta|_{\mathcal{P}}$ counts the pointer occurrences in η, so that the encoder can give every occurrence a distinct address.

Theorem 11 (Encoding within address capacity Lean: `Physical.encoding_exists_direct`). *Suppose* $\mathsf{Inv}(\Delta)$, ℓ *is well formed on* $\mathcal{V}_\tau$ *under* Δ *and concrete for* τ, $\varnothing \vdash \ell$ disc, *and* $\varnothing \vdash \Delta(\mathsf{T})$ disc *for every* $\mathsf{T} \in \operatorname{dom}\Delta$. *Suppose every pointer in* ℓ *and the bodies bound by* Δ *has width at least* w. *For every* $v \in \mathcal{V}_\tau$ *satisfying the address-capacity bound, the encoder succeeds at some fuel on some permitted representation,*

$$\exists \eta \in \mathcal{W}_\ell^\Delta(v),\ n \in \mathbb{N},\ m \in \mathcal{M}. \quad \mathsf{encode}_\ell^n(\eta) = m.$$

Lemma 12 (Recovery by forward decoding Lean: `Physical.roundTrip_direct`). *If* Δ *and* ℓ *are well formed,* $\varnothing \vdash \ell$ $\mathsf{disc}_{\rightarrow}$, *and* $\varnothing \vdash \Delta(\mathsf{T})$ $\mathsf{disc}_{\rightarrow}$ *for every* $\mathsf{T} \in \operatorname{dom}\Delta$, *then at every fuel* n *the decoder* decode_ℓ^n *returns* η *for every successful encoding* $\mathsf{encode}_\ell^n(\eta)$ *of a permitted representation* η, *and so does every larger fuel. Each such pair therefore satisfies the storage recovery property Equation* (7).

G Program Preservation

Programs consist of first-order function definitions over source values, with the expression syntax (Lean: `Programs.{Expr, Prog}`)

$$x, f, \mathbf{K} \in \mathsf{Id} \qquad k \in \mathbb{N} \qquad op : \mathbb{N} \times \mathbb{N} \rightharpoonup \mathbb{N}$$
$$e ::= x \mid k \mid op(e, e) \mid \mathbf{K}(\bar{e}) \mid f(\bar{e}) \mid \textbf{let } x = e \textbf{ in } e \mid \textbf{match } e \textbf{ with } \{\mathbf{K}_i(\bar{x}_i) \Rightarrow e_i\} \qquad P ::= \{f_i(\bar{x}_i) = e_i\}.$$

For simplicity, function and match-arm lookup select the first matching name. Both interpretations evaluate arguments from left to right before construction or a call. A call requires the declared arity and evaluates the function body in a fresh environment containing only its parameter bindings; recursive calls use the same P. Variables, literals, scalar operations, and let bindings have their usual call-by-value meaning. An undefined lookup or scalar operation supplies no evaluation derivation.

Source evaluation $P; \rho \vdash e \Downarrow v$ uses the values of Section 4.1. Construction evaluates its arguments to $\bar{v}$ and forms $\mathbf{K}(\bar{v})$, provided this value belongs to the constructor's declared type. Matching selects the arm for the evaluated constructor and binds its fields to the branch variables in declaration order, with equal arities. Structural evaluation uses the same program and scalar operations, but uses represented datatype values, defined below.

The datatype interface. Represented values comprise scalars and datatype handles. For a relation $r \sim v$ from represented values to source values, the interface provides partial operations

$$\mathsf{construct}(\mathbf{K}, \bar{r}), \qquad \mathsf{destruct}(r).$$

Scalars relate only to the same scalar, datatype handles relate only to constructor values, and each represented value relates to at most one source value. For componentwise-related arguments $\bar{r} \sim \bar{v}$, construction succeeds exactly when $\mathbf{K}(\bar{v})$ is admitted by the declaration, including its field types and arity; its result relates to $\mathbf{K}(\bar{v})$. Destruction of a handle related to $\mathbf{K}(\bar{v})$ succeeds with $(\mathbf{K}, \bar{r})$ where $\bar{r} \sim \bar{v}$ in declaration order. These conditions give both directions of the constructor and match cases in evaluation correspondence.

The remainder instantiates this interface with faithful layouts, through a construction policy b and a destruction policy ι satisfying

$$b_{\mathsf{T}}(v) \in \mathcal{W}^{\Delta}_{\mathsf{ref}\,\mathsf{T}}(v), \qquad \iota_{\mathsf{T}}(\eta) = v \quad \text{for } v \in \mathcal{V}_{\mathsf{T}} \text{ and } \eta \in \mathcal{W}^{\Delta}_{\mathsf{ref}\,\mathsf{T}}(v).$$

A valid b meets the left condition on every admitted value and returns nothing for any other value, and a valid ι meets the right one.

Computable policies. The encoder of Lemma 3 gives a computable valid b under $\mathsf{Inv}(\Delta)$ when Δ binds every declared type, and the canonical decoder, composed with the reconstruction functions of further rewrites, gives a valid ι. These are computable functions of the model; implementing them as generated code requires a separate proof of the same interface laws.

Theorem 13 (Derived policies Lean: `Core.UpdateSeq.{decoder_decodes,observations}`). *Let Δ_0 be a canonical environment and $\Delta_0 \longrightarrow^* \Delta_1$ a finite sequence of valid updates that keep their reconstruction functions. Undoing the updates in reverse order, each by applying its function at every reference to the updated type and then applying the canonical decoder gives a valid destruction policy ι_1 for Δ_1. Hence (6) holds with the encoder of Lemma 3 at both endpoints, the canonical decoder at Δ_0, and ι_1 at Δ_1, for all related inputs.*

Views and related inputs. A represented value has the form

$$r ::= k \mid \langle \mathsf{T}, \eta, \pi \rangle.$$

The type T describes the value represented by η; the path π selects a datatype subvalue of that value, possibly of another type. Scalars represent themselves. A view represents a constructor value w exactly when an admitted root has that representation and selects w:

$$\begin{aligned} k \sim_\Delta v &\iff v = k, \\ \langle \mathsf{T}, \eta, \pi \rangle \sim_\Delta w &\iff w \text{ is a constructor value and} \\ &\qquad \exists u \in \mathcal{V}_\mathsf{T}. \quad \eta \in \mathcal{W}^\Delta_{\mathrm{ref}\,\mathsf{T}}(u) \ \wedge \ u.\pi = w. \end{aligned}$$

Environments $\rho' \sim_\Delta \rho$ have the same domain and satisfy $\rho'(x) \sim_\Delta \rho(x)$ at each bound variable. Thus related inputs provide both evaluations with the same source values, while their structural representations may differ.

A destruction policy ι induces the partial reader

$$\alpha_\iota(k) = k, \qquad \alpha_\iota(\langle \mathsf{T}, \eta, \pi \rangle) = \big(\iota_\mathsf{T}(\eta)\big).\pi.$$

Under a valid ι, related values satisfy $r \sim_\Delta v \implies \alpha_\iota(r) = v$. Structural construction first evaluates its arguments to $\bar{r}$ and reads $\bar{v} = \alpha_\iota(\bar{r})$ componentwise. For $v = \mathbf{K}(\bar{v}) \in \mathcal{V}_\mathsf{T}$, the policy b returns a permitted representation and construction produces

$$\langle \mathsf{T}, b_\mathsf{T}(v), \epsilon \rangle.$$

The policies are fixed for an evaluation. We represent each as a partial function on arbitrary type names and values or representations; validity of b requires a permitted result for every value of each declared type and no result for any other value. The functions b_T above are its total restrictions to those values.

Matching evaluates its scrutinee to $\langle \mathsf{T}, \eta, \pi \rangle$ and obtains its constructor through α_ι. For declared fields $f_i : \tau_i$, it binds scalar fields to their recovered scalar values and datatype fields to $\langle \mathsf{T}, \eta, \pi.f_i \rangle$. It then evaluates the selected branch in the extended environment. The extended views keep η and need no separately constructed representation of a field. The branch variables must match the declared field count. These operations and the common rules define $\Delta, b, \iota; P; \rho' \vdash e \Downarrow r$.

Lemma 14 (Evaluation correspondence Lean: `Programs.evaluation_value_iff`). *Suppose Δ is faithful. For valid policies b and ι and related environments $\rho' \sim_\Delta \rho$,*

$$P; \rho \vdash e \Downarrow v \quad \iff \quad \exists r. \quad \Delta, b, \iota; P; \rho' \vdash e \Downarrow r \ \wedge \ r \sim_\Delta v.$$

Applying correspondence to both environments gives Theorem 8. Preserving faithfulness of the environment then gives Theorem 9. A view of a field continues to refer to its enclosing representation, so this argument requires no agreement between source match depth and layout unfolding depth.

H Additional Evaluation Results

H.1 Complete Layout Frontiers

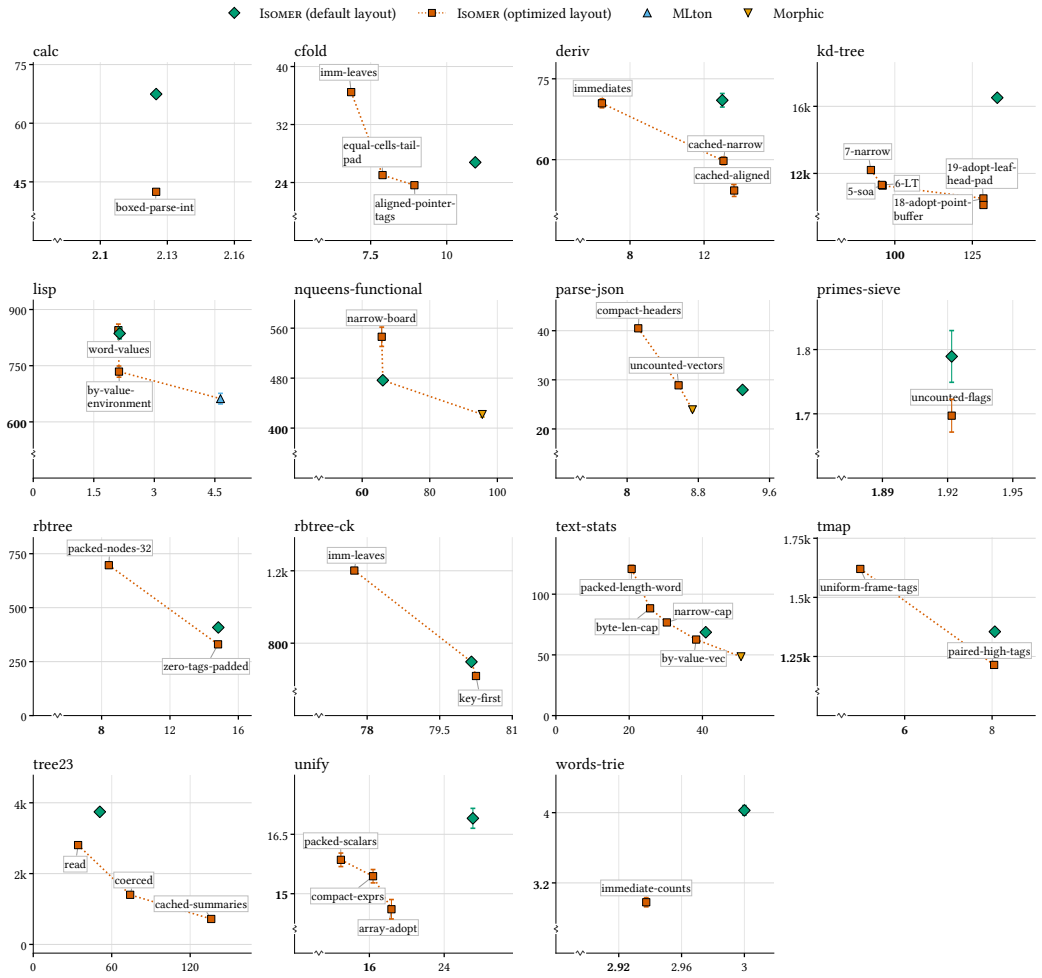


Fig. 16. Lower is better. Mimalloc runtime–RSS frontiers for benchmarks with a non-default Pareto-optimal layout. Each panel plots runtime (ms, vertical) against peak RSS (MiB, horizontal). ◆ marks a default ISOMER layout; ▪···▪ marks an ISOMER layout optimized via rewrites on the frontier (▪) and traces the frontier (·····). Triangles mark MLton (△) and Morphic (▽). Dominated layouts are omitted.

H.2 Peak RSS (Mimalloc)

We report whole-process peak resident set size (RSS), normalized to ISOMER with its default layout. For MLton, this measure includes its GC-managed heap.

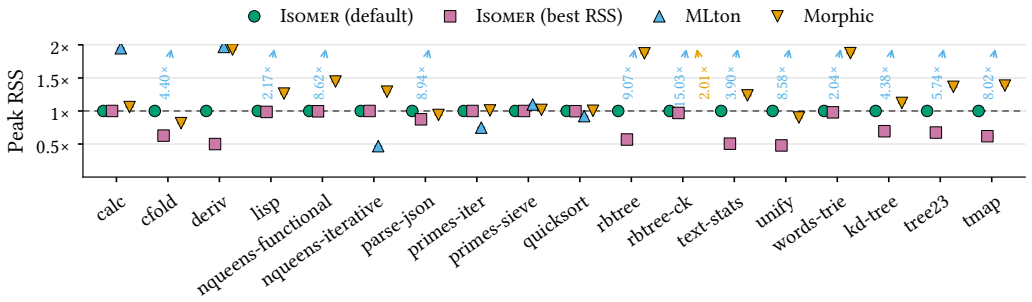


Fig. 17. Lower is better. Mimalloc peak RSS relative to default Isomer. The best-RSS layout is selected independently for each benchmark. Arrows indicate ratios beyond the displayed range.

H.3 Parity Layouts with Morphic

We run Isomer with Morphic's own layouts, derived with our rewrite rules and matched bit for bit, so that only the compilers differ. One program, `primes-iter`, declares no ADTs. Isomer is 1.18× faster (geomean).

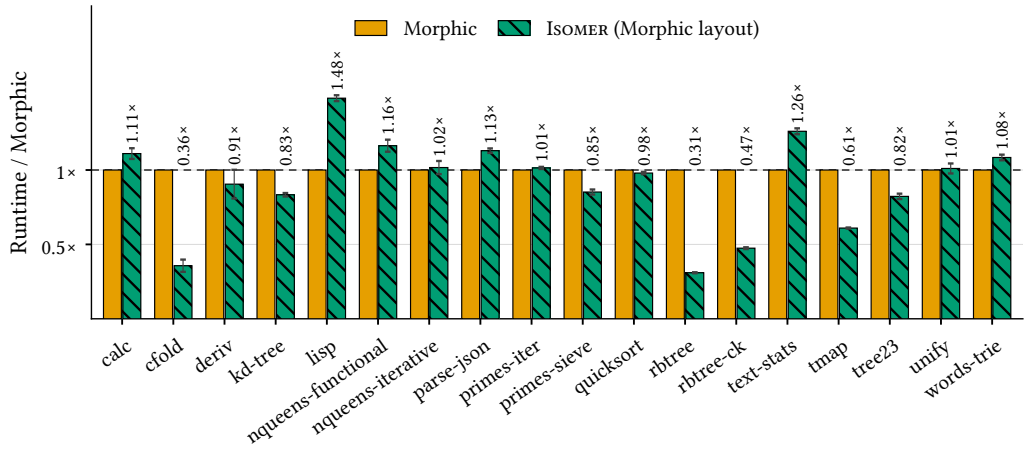


Fig. 18. Lower is better. Mimalloc runtime relative to Morphic, with Isomer at Morphic's layouts.

H.4 System Allocator Results

We repeat the workloads and sampling protocol of Section 9.1 with the system allocator. MLton uses its own tracing heap, so its measurements are shared across allocator comparisons. Isomer with default layout achieves speedups (geomean) of 1.01× versus MLton and 1.41× versus Morphic, using 0.29× and 0.84× their peak RSS, respectively. Morphic-style layouts give a 1.36× speedup over Morphic across all 18 benchmarks. Against Ribbit on `pcf`, `option`, and `rb-tree`, Isomer achieves speedups of 1.20×, 1.45×, and 2.22×, respectively, or 1.57× (geomean).

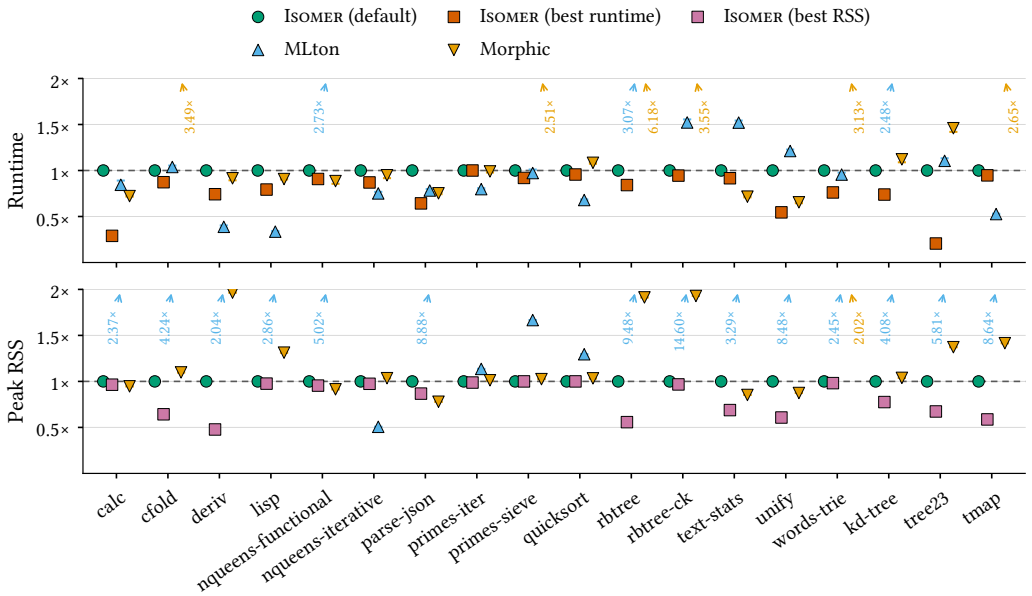


Fig. 19. Lower is better. System-allocator runtime (top) and peak RSS (bottom) relative to default Isomer, using the conventions of Figure 9. The best-runtime and best-RSS layouts are selected independently for each benchmark. Arrows indicate ratios beyond the displayed range.

H.5 MLton Array Representations

For each benchmark, calibration selects the fastest of four MLton array implementations whose output agrees with the reference on the checked workload. The selected implementation is recorded and used for every timed run. `primes-iter` and `tmap` contain no source arrays.

Table 1. MLton array implementations selected for the comparison benchmarks.

Mode	Array Implementation	Benchmarks
`mutate`	mutable	`calc`, `nqueens-iterative`, `parse-json`, `primes-sieve`, `quicksort`, `rbtree-ck`, `text-stats`, `words-trie`, `kd-tree`, `tmap`
`copy`	Flat, copied on update	`lisp`, `rbtree`, `tree23`
`reroot`	Persistent, update chains	`nqueens-functional`, `primes-iter`, `unify`
`trie32`	Persistent, 32-way trie	`cfold`, `deriv`

I Mechanization

Table 2 lists, in paper order, every definition, lemma, and theorem of Sections 4 to 7 and of the appendices, together with the mechanized examples of Sections 2 and 8.1. Files are relative to the root of the Lean development, with the prefix `Ismr/` omitted. We additionally provide `Paper.lean`, which restates each numbered result. Propositional extensionality and quotient soundness are the only axioms used in this formalization.

Table 2. The definitions, theorems, and claims of the paper and its appendices.

Paper	File	Lean
Section 2		
Listings (1), (2), (4) and (5), derived by the rules and faithful	`Examples/Motivation.lean`	`Examples.Motivation.{morphic_core, lowBits_core, shape_core, terminal_core, morphic_stores, morphic_faithful, lowBits_faithful, shape_faithful, terminal_faithful}`
Section 4		
Section 4.1, values, typing, and paths	`Source.lean`	`Value, Values, Value.select, Decls.WF`
Section 4.2, provenances as sets of values, admitted subjects	`Source.lean`	`Prov.holds`
	`Provenance.lean`	`Prov.Equiv`
	`Layout/Basic.lean`	`Domain`
Figure 1, layouts	`Layout/Basic.lean`	`Layout, LayoutEnv`
	`Source.lean`	`Path, Prov`
	`Expr.lean`	`Expr`
Arms and their guards	`Layout/SumArm.lean`	`SumArm, Prov.Unique`
	`Layout/ArmMaps.lean`	`ArmMap`
Tags, selectors, and code widths	`Source.lean`	`Decls.lookupTag, ArmSelector, ArmSelector.{sound, complete}, codeWidth, Decls.tagWidth`
Canonical layout environment	`Layout/Canonical.lean`	`Canonical.env`
Section 4.4, layout contexts	`Layout/Context.lean`	`Ctxt, Ctxt.{fill, holeDom}`
Figure 2, rules and certificates	`Core/Rules.lean`	`Step, Step.{certificate, local}`
	`Core/Local.lean`	`RuleCertificate`
	`Layout/Flatten.lean`	`refineArms, flattenArms, WF.flatten, Rule.flattenRecon, Rule.flattenRecon_rep`
	`Provenance/KeepFirst.lean`	`Prov.keepFirst, Prov.{keepFirst_pairwise, keepFirst_covers, find?_guard_of_mem_keepFirst}`
Definition updates	`Layout/Basic.lean`	`LayoutEnv.set`
Figure 3, the double-ended layout	`Examples/Motivation.lean`	`Examples.Motivation.{doubleEnded_core, doubleEnded_stores, doubleEnded_faithful}`
Section 5		
Section 5.1, structural representations	`Layout/Semantics.lean`	`Enc`
Section 5.1, the canonical representation of the example node	`Examples/Worked.lean`	`Examples.Worked.{canonical_node, canonical_node_only}`
Figure 4, interpretation	`Layout/Semantics.lean`	`Rep`
Definition 1, faithfulness and determination	`Correctness/Faithfulness.lean`	`Faithful, Total, AbsoluteInjective, Determines, AbsoluteInjective.determines`
Section 5.3 and Figure 13, well-formedness	`Layout/WellFormed.lean`	`WF, LayoutEnv.WF`
	`Layout/ArmMaps.lean`	`WF.sumMap_iff`
	`Layout/TaggedArms.lean`	`WF.taggedMap_iff`
Definition 2, environment invariant	`Core/Local.lean`	`Inv`
	`Layout/WellFormed.lean`	`LayoutEnv.WF`
	`Layout/Structural.lean`	`LayoutEnv.Descent`
Lemma 3, totality by a computable encoder	`Correctness/Totality.lean`	`WF.totality, represent_total`
	`Layout/Represent.lean`	`represent, represent_sound`
Section 6		
Section 6.1, sharing, refinement, and its laws	`Core/Local.lean`	`Shares, absoluteInjective_iff_shares, Refines, Refines.{refl, trans, faithful}`
Definition 4, local refinement	`Core/Local.lean`	`LocalRefinement`
Section 6.1, reconstruction functions	`Core/Local.lean`	`Reconstructs, Reconstructs.refines, LocalRefinement.of_reconstructs`
	`Core/Rules.lean`	`Step.certificate`
	`Core/StorageRules.lean`	`Recon.{dropFirst, swap, join}`
Theorem 5, contextual replacement	`Core/Contextual.lean`	`Refines.context, LocalRefinement.context`

Paper	File	Lean
Lemma 6, recursive binding	Core/Recursive.lean	LocalRefinement.bind, Refines.bind
Theorem 7, valid rewrites and soundness	Core/Correctness.lean	Rewrite, RewriteSeq, Rewrite.sound, RewriteSeq.sound
Section 6.2, composition and canonical starts	Core/Correctness.lean	RewriteSeq.faithful
	Layout/Canonical.lean	Canonical.body_wf
	Layout/CanonicalDecode.lean	Canonical.{faithful, decode_rep}
	Core/Canonical.lean	canonical_rewrites_faithful
Section 6.3, interval encoding	Core/Interval.lean	Interval.{Bounds, certificate, local_refinement, rewrite}
Section 6.4, the two evaluations, valid policies, related values	Programs/SourceSemantics.lean	Programs.Eval
	Programs/StructuralSemantics.lean	Programs.StructuralEval
	Programs/Views.lean	Programs.{Builder, Decoder, Builds, Decodes, Represents, FaithfulRoots}
	Correctness/ProgramPreservation.lean	Programs.EnvRel
Theorem 8, representation independence	Correctness/ProgramPreservation.lean	Programs.representation_independence
	Programs/ViewInterface.lean	Programs.view_laws
	Programs/Interface.lean	Programs.Interface.independence
Theorem 9, program preservation	Core/Programs.lean	BodyRewriteSeq.{observations, roots}
	Layout/Structural.lean	LayoutEnv.Binds
Section 6.4, computable policies	Programs/Construct.lean	Programs.Builds.construct
	Programs/CanonicalDecoder.lean	Programs.Decodes.canonical
	Core/Policies.lean	UpdateSeq.{decoder, decoder_decodes}
Section 7		
Equation (7), storage recovery	Physical/Semantic.lean	Physical.{SemanticRoundTrip, RoundTrip}, Physical.RoundTrip.semantic
Theorem 10, recovery, loaded inputs	Core/Storage.lean	storage_recovery
	Physical/Semantic.lean	Physical.SemanticRoundTrip.recovery
	Programs/SemanticStorage.lean	Programs.loaded_evaluation_iff
Section 7.1, the forward codec and a stored node	Physical/Coding.lean	Physical.{encodeLayout, decodeLayout}
	Physical/Direct.lean	Physical.{encoding_exists_direct, roundTrip_direct}
	Examples/Motivation.lean	Examples.Motivation.{lowBits_encode, lowBits_round_trip, lowBits_recovery}
Section 7.2, read order, ambiguity	Physical/Counterexamples.lean	Physical.Limits.{read_order, encoding_ambiguity, no_value_decoder}
	Physical/Semantic.lean	Physical.SemanticRoundTrip.injective
Section 8		
Section 8.1, a finite array as a sum indexed by its length	Examples/Array.lean	Examples.Array.{arms, sized, tagSize, updates, sized_faithful, construct_a79, round_trip}
Appendix A		
Appendix A, deciding provenance equality	Provenance/Decide.lean	Prov.{decidableEquiv, equivCheck_iff, counterexample, counterexample_sound, counterexample_complete}
	Provenance/Finite.lean	Observation.{proj, proj_mem, holds_proj}
Appendix C		
Figure 14, rule judgments	Core/Rules.lean	Step
Appendix D		
Figure 15, the representations of $\mathbf{B}(0, 1)$	Examples/Worked.lean	Examples.Worked.b01_representations
Appendix F		
Appendix F.1, discriminant closure	Layout/Stored.lean	Stored, TagAccess, TagsStored
Fuel-indexed codec	Physical/Coding.lean	Physical.{encodeLayout, decodeLayout, encodeLayout_mono_le, decodeLayout_mono_le}
Theorem 11, encoding within capacity	Physical/Direct.lean	Physical.encoding_exists_direct
	Physical/Existence.lean	Physical.{ConcreteOn, AddressCapacity}
Lemma 12, forward decoding	Physical/Direct.lean	Physical.roundTrip_direct
	Physical/Words.lean	Physical.Direct
Appendix G		
Programs	Programs/Syntax.lean	Programs.{Expr, Prog}
Theorem 13, derived policies	Core/Policies.lean	UpdateSeq.{decoder, decoder_decodes, observations}
	Core/Inversion.lean	BodyUpdate, BodyUpdate.{undo, undo_decodes}, UpdateSeq, UpdateSeq.{undo, undo_decodes}, Ctxt.recon
Lemma 14, evaluation correspondence	Correctness/ProgramPreservation.lean	Programs.evaluation_value_iff