

# Values Are Primary: Immutability as a Foundation for Type-Safe Concurrent Programming in JAPL

Matthew Long  
*The JAPL Research Collaboration*  
*YonedaAI Research Collective*  
Chicago, IL  
`matthew@yonedaai.com`

March 2026

## Abstract

The dominant paradigm in mainstream programming languages treats *objects with identity* as the primary organising concept: mutable cells, reference semantics, and pointer aliasing are the default. This paper argues that an alternative foundation—one in which *values* are primary—yields languages that are simpler to reason about, safer by construction, and better suited to concurrent and distributed systems.

We develop this thesis in the context of JAPL, a *strict*, statically typed, effect-aware functional language that makes immutable algebraic data types the default representation for all data. Strict evaluation—in contrast to Haskell’s lazy semantics—ensures that space usage is predictable and that values are fully evaluated, which simplifies both reasoning and garbage collection. Mutation is available but confined to an explicitly tracked resource layer. We present a formal framework grounded in category theory—modelling types as objects of a Cartesian closed category and algebraic data types as initial algebras of polynomial endofunctors—and give a denotational semantics for a value-centric core calculus that guarantees referential transparency. We describe JAPL’s type system (parametric polymorphism, bidirectional inference, row polymorphism for extensible records, opaque types) and show how the value/resource split enables a dual memory management strategy: a simplified garbage collector for immutable data and linear-type-governed ownership for mutable resources.

Through detailed comparison with Haskell, Rust, Erlang, OCaml, Go, Gleam, Clojure, and Elixir, we demonstrate that JAPL’s value model occupies a distinctive point in the design space. We evaluate implementation strategies—hash-array mapped tries, structural sharing, copy-on-write, packed tagged unions—and present benchmarks showing that value semantics need not sacrifice performance relative to mutable approaches. Finally, we discuss the profound implications of default immutability for concurrency: values can be freely shared across process boundaries without synchronization, making JAPL’s Erlang-style process model both safe and efficient.

**Keywords:** value semantics, immutability, algebraic data types, persistent data structures, type theory, category theory, concurrency, functional programming, language design.

# Contents

<b>1</b>	<b>Introduction</b>	<b>3</b>
1.1	The Identity Crisis in Programming	3
1.2	The Cost of Identity	3
1.3	Our Thesis	3
1.4	Contributions	4
1.5	Paper Organisation	4
<b>2</b>	<b>Background and Related Work</b>	<b>5</b>
2.1	The Value of Values	5
2.2	ML and the Algebraic Tradition	5
2.3	Haskell: Purity by Default	5
2.4	Erlang/OTP: Values in the Actor Model	5
2.5	Rust: Ownership Without Garbage Collection	5
2.6	Persistent Data Structures	6
2.7	Category-Theoretic Foundations	6
<b>3</b>	<b>Formal Framework</b>	<b>6</b>
3.1	Types as Objects of a Category	6
3.2	Algebraic Data Types as Initial Algebras	7
3.3	Value Semantics: Formal Properties	7
3.4	Denotational Semantics of a Value-Centric Core Calculus	8
3.4.1	Syntax	8
3.4.2	Typing Rules (Selected)	8
3.4.3	Denotational Semantics	9
3.5	The Yoneda Perspective on Values	9
3.5.1	Yoneda in JAPL's Design	10
<b>4</b>	<b>JAPL's Value Model</b>	<b>10</b>
4.1	Immutable by Default	10
4.2	Algebraic Data Types	11
4.2.1	Product Types: Structural Records	11
4.2.2	Sum Types: Tagged Unions	11
4.2.3	Pattern Matching	11
4.3	The Value/Resource Split	12
4.4	Persistent Data Structure Strategies	13
<b>5</b>	<b>Type System for Values</b>	<b>13</b>
5.1	Parametric Polymorphism	13
5.2	Type Inference: Local Bidirectional Checking	14
5.3	Row Polymorphism for Extensible Records	15
5.4	Opaque and Newtype Wrappers	15
5.5	Traits as Type Classes	16

<b>6</b>	<b>Comparison of Value Models</b>	<b>16</b>
6.1	Haskell . . . . .	17
6.2	Rust . . . . .	17
6.3	Erlang and Elixir . . . . .	17
6.4	Gleam . . . . .	18
6.5	OCaml . . . . .	18
6.6	Clojure . . . . .	18
6.7	Go . . . . .	18
<b>7</b>	<b>Implementation Strategies</b>	<b>19</b>
7.1	Hash-Array Mapped Tries (HAMTs) . . . . .	19
7.2	Structural Sharing . . . . .	19
7.3	Copy-on-Write for Small Values . . . . .	20
7.4	Packed Tagged Unions . . . . .	20
7.5	Cache-Friendly Layouts . . . . .	20
7.6	Uniqueness Analysis . . . . .	21
7.6.1	High-Level Algorithm . . . . .	21
<b>8</b>	<b>Evaluation</b>	<b>22</b>
8.1	Experimental Setup . . . . .	22
8.2	Map Operations . . . . .	22
8.3	Record Updates . . . . .	22
8.4	Concurrent Workload: Shared Read-Heavy Map . . . . .	23
8.5	Memory Overhead . . . . .	23
8.6	GC Performance . . . . .	23
<b>9</b>	<b>Implications for Concurrency</b>	<b>24</b>
9.1	Values Are Freely Shareable . . . . .	24
9.2	Zero-Copy Message Passing . . . . .	24
9.3	Snapshots Without Coordination . . . . .	25
9.4	Distribution Without Serialization Ambiguity . . . . .	25
9.5	Process-Based State Management . . . . .	26
<b>10</b>	<b>Discussion</b>	<b>26</b>
10.1	When Values Are Not Enough . . . . .	26
10.2	JAPL’s Dual-Layer Model . . . . .	27
10.2.1	Why the Boundary Preserves Safety . . . . .	27
10.3	The Composition Hierarchy . . . . .	28
10.4	Trade-offs and Limitations . . . . .	28
10.4.1	Memory Overhead . . . . .	28
10.4.2	Learning Curve . . . . .	28
10.4.3	Interoperability . . . . .	28
10.5	Related Design Patterns . . . . .	29
10.5.1	Event Sourcing . . . . .	29
10.5.2	Functional Reactive Programming . . . . .	29
10.5.3	Content-Addressable Storage . . . . .	29
<b>11</b>	<b>Conclusion</b>	<b>29</b>

<b>A Proof of Parametricity for JAPL's Core Calculus</b>	<b>32</b>
<b>B JAPL Core Syntax Reference</b>	<b>33</b>

# 1 Introduction

## 1.1 The Identity Crisis in Programming

The history of programming languages can be read as a long negotiation between two worldviews. In the *identity-centric* view—exemplified by Simula [11], Smalltalk [28], Java [20], and their descendants—the fundamental concept is the *object*: an entity with mutable state, behaviour, and identity. Two objects may have identical field values yet be distinct because they occupy different locations in memory. Programs are understood as networks of collaborating objects that communicate by sending messages (method calls) that mutate their internal state.

In the *value-centric* view—rooted in the lambda calculus [9], refined in ML [41], Haskell [26], and Erlang [2]—the fundamental concept is the *value*: a datum that simply *is* what it is, with no notion of mutable state or location-dependent identity. The number 42 is 42 regardless of where it is stored; the string "hello" does not change over time. Programs are understood as compositions of functions that transform values into other values.

## 1.2 The Cost of Identity

The identity-centric model imposes substantial costs that become especially acute in concurrent and distributed settings:

1. **Aliasing and mutation.** When multiple references point to the same mutable object, a mutation through one alias is visible through all others. This makes local reasoning impossible without whole-program alias analysis, which is undecidable in the general case [33].
2. **Synchronization overhead.** In concurrent programs, shared mutable state must be protected by locks, atomic operations, or transactional mechanisms. Every lock is a potential source of deadlock; every unlocked access is a potential data race. The fundamental difficulty is that identity creates *contention*: two threads that wish to observe or modify the same object must coordinate.
3. **Serialization ambiguity.** When transmitting an object graph across a network, the serialization framework must decide whether aliased references should be preserved (deep structural copy) or collapsed (sharing-preserving copy). This decision is invisible at the type level and frequently leads to subtle bugs in distributed systems [2].
4. **Cache invalidation.** Mutable state requires cache invalidation protocols—one of the “two hard things in computer science” [16]. An immutable value, once cached, never becomes stale.
5. **Testing difficulty.** Objects with mutable state are difficult to test in isolation because their behaviour depends on the history of mutations. Values, being immutable, have no history; they are their own specification.

## 1.3 Our Thesis

We contend that treating *values as primary*—making immutable algebraic data types the default representation for all data, and confining mutation to an explicitly tracked resource layer—yields a language design that is:

- **Simpler:** local reasoning is restored; every binding denotes a fixed value.

- **Safer:** data races are eliminated by construction for all pure data.
- **Composable:** values compose freely via functions, without the fragile coupling of shared mutable state.
- **Distributable:** values can be serialized, transmitted, cached, and replicated without ambiguity.

This paper develops this thesis in the context of JAPL, a strict, typed, effect-aware functional programming language. JAPL’s design motto is: “Pure by default, concurrent by design, resource-safe by construction, distributed without apology.” The *Values Are Primary* principle is the first of JAPL’s seven core principles, and it is foundational: the remaining six principles build upon the guarantees that pervasive immutability provides.

## 1.4 Contributions

This paper makes the following contributions:

1. A formal framework for value semantics grounded in category theory, modelling algebraic data types as initial algebras of polynomial endofunctors (§3).
2. A denotational semantics for a value-centric core calculus that guarantees referential transparency (§3.4).
3. A detailed description of JAPL’s value model, including immutable algebraic data types, structural records, tagged unions, and pattern matching (§4).
4. A type system design combining parametric polymorphism, bidirectional inference, row polymorphism, and opaque types (§5).
5. A systematic comparison with eight other languages (§6).
6. Implementation strategies for efficient persistent data structures (§7).
7. Benchmarks demonstrating that value semantics need not sacrifice performance (§8).
8. An analysis of the implications of default immutability for concurrent and distributed programming (§9).

## 1.5 Paper Organisation

Section 2 surveys background and related work. Section 3 develops the formal framework. Section 4 presents JAPL’s value model. Section 5 describes the type system. Section 6 compares with related languages. Section 7 discusses implementation strategies. Section 8 presents benchmarks. Section 9 analyses concurrency implications. Section 10 discusses trade-offs and the dual-layer model. Section 11 concludes.

## 2 Background and Related Work

### 2.1 The Value of Values

Rich Hickey’s influential talk *The Value of Values* [24] articulated a distinction that practitioners had felt but rarely named: the difference between *values* (immutable, timeless facts) and *places* (mutable, time-varying locations). Hickey argued that conflating the two—the default in most mainstream languages—is the root cause of much accidental complexity. A *value* is a piece of information; it does not change. The number 42, the date March 15, the set  $\{1, 2, 3\}$ —these are values. A “mutable integer variable” is not a value but a *place* that happens to contain an integer at this moment in time.

Hickey’s Clojure [23] operationalised this insight by making all core data structures—lists, vectors, hash maps, sets—immutable and persistent. The language provides controlled mutation through *atoms* (compare-and-swap references), *refs* (software transactional memory), and *agents* (asynchronous state updates), all of which are clearly delineated from pure values.

### 2.2 ML and the Algebraic Tradition

The ML family [41, 42] introduced algebraic data types and pattern matching as core language features. Standard ML [42] and its descendants—OCaml [34], F# [52]—demonstrate the power of sum types (tagged unions) and product types (records/tuples) for modelling data. However, ML languages generally permit unrestricted mutation via `ref` cells, meaning that the value-centric idiom is a *convention* rather than a *guarantee*.

### 2.3 Haskell: Purity by Default

Haskell [26, 38] takes the value-centric approach to its logical extreme: all expressions denote values, and side effects are confined to the IO monad [44]. This provides the strongest possible guarantees—referential transparency holds for all non-IO code—but the monadic encoding of effects introduces syntactic overhead and monad transformer stacks that many practitioners find difficult [29]. Furthermore, Haskell’s lazy evaluation strategy means that “values” may in fact be unevaluated thunks, complicating space usage reasoning [25].

### 2.4 Erlang/OTP: Values in the Actor Model

Erlang [1, 2] combines immutable values with the actor model. All data in Erlang is immutable; variables are single-assignment. Processes communicate by sending messages, which are deep-copied into the recipient’s mailbox. This design—immutable values plus process isolation—has proven extraordinarily robust in telecommunications and distributed systems. The main limitation is the lack of a static type system; Erlang is dynamically typed, and the Dialyzer [35] provides only post-hoc analysis. Elixir [54] builds on the Erlang VM with improved syntax but retains the same value model. Gleam [19] adds static typing to the BEAM ecosystem while preserving Erlang’s value semantics.

### 2.5 Rust: Ownership Without Garbage Collection

Rust [30, 39] takes a different approach to the mutation problem. Rather than prohibiting mutation, Rust controls *aliasing*: through the ownership and borrowing system, Rust ensures that mutable references are unique (no aliasing) while shared references are immutable (no mutation). This “aliasing XOR mutability” discipline [27] eliminates data races without requiring immutability.

Values in Rust are moved by default and copied only when explicitly requested via the `Copy` or `Clone` traits.

## 2.6 Persistent Data Structures

Okasaki’s seminal work [43] demonstrated that purely functional data structures can achieve asymptotically efficient operations through techniques such as lazy evaluation, amortisation, and structural decomposition. Bagwell’s hash-array mapped tries [3] provided the practical basis for Clojure’s persistent vectors and hash maps, achieving near-constant-time operations with  $O(\log_{32} n)$  path-copying overhead. Driscoll et al. [13] formalised the theory of persistent data structures, distinguishing partial from full persistence and establishing lower bounds.

## 2.7 Category-Theoretic Foundations

The connection between algebraic data types and initial algebras in category theory was established by Hagino [21] and further developed by Malcolm [37] and Meijer et al. [40]. The insight is that a data type definition such as `List(a) = Nil | Cons(a, List(a))` defines a functor  $F(X) = 1 + a \times X$ , and the type `List(a)` is the carrier of the initial  $F$ -algebra. This gives a principled account of structural recursion (catamorphisms) and induction on data types. Pierce’s *Types and Programming Languages* [46] and Harper’s *Practical Foundations for Programming Languages* [22] provide comprehensive treatments of the type-theoretic perspective.

# 3 Formal Framework

## 3.1 Types as Objects of a Category

We model the type system of our value-centric calculus using the framework of *Cartesian closed categories* (CCCs) [5, 32].

**Definition 3.1** (Category of Types). Let **Type** be a category whose objects are types and whose morphisms  $f : A \rightarrow B$  are (total, terminating) programs that take a value of type  $A$  and produce a value of type  $B$ . Composition is function composition; the identity morphism at  $A$  is the identity function  $\text{id}_A : A \rightarrow A$ .

**Definition 3.2** (Cartesian Closed Structure). **Type** is Cartesian closed if it has:

1. A *terminal object*  $\mathbf{1}$  (the unit type).
2. *Binary products*  $A \times B$  for all objects  $A, B$  (product types / tuples).
3. *Exponential objects*  $B^A$  for all objects  $A, B$  (function types  $A \rightarrow B$ ).

The Cartesian closed structure ensures that products and function types interact correctly: there is a natural bijection  $\text{Hom}(A \times B, C) \cong \text{Hom}(A, C^B)$  (currying).

*Remark 3.3.* The Cartesian closed structure captures the essence of *value-centric* programming: products model data aggregation (records, tuples), exponentials model computation (functions), and the terminal object models the trivial value (unit). No morphism has a “side effect”; all information flows through inputs and outputs.

### 3.2 Algebraic Data Types as Initial Algebras

**Definition 3.4** (Polynomial Endofunctor). A *polynomial endofunctor*  $F : \mathbf{Type} \rightarrow \mathbf{Type}$  is built from the grammar:

$$F ::= \text{Id} \mid K_A \mid F_1 + F_2 \mid F_1 \times F_2$$

where  $\text{Id}$  is the identity functor,  $K_A$  is the constant functor at  $A$ ,  $+$  is the coproduct (sum), and  $\times$  is the product.

**Definition 3.5** ( $F$ -Algebra). An  $F$ -algebra is a pair  $(A, \alpha)$  where  $A$  is an object of  $\mathbf{Type}$  (the *carrier*) and  $\alpha : F(A) \rightarrow A$  is a morphism (the *structure map*). A homomorphism of  $F$ -algebras  $(A, \alpha) \rightarrow (B, \beta)$  is a morphism  $h : A \rightarrow B$  such that  $h \circ \alpha = \beta \circ F(h)$ .

**Definition 3.6** (Initial Algebra). An  $F$ -algebra  $(\mu F, \text{in})$  is *initial* if, for every  $F$ -algebra  $(A, \alpha)$ , there exists a unique homomorphism  $\llbracket \alpha \rrbracket : \mu F \rightarrow A$  satisfying:

$$\llbracket \alpha \rrbracket \circ \text{in} = \alpha \circ F(\llbracket \alpha \rrbracket)$$

The morphism  $\llbracket \alpha \rrbracket$  is called the *catamorphism* (or fold) induced by  $\alpha$ .

**Theorem 3.7** (Lambek's Lemma [31]). *If  $(\mu F, \text{in})$  is an initial  $F$ -algebra, then  $\text{in} : F(\mu F) \rightarrow \mu F$  is an isomorphism. That is,  $\mu F \cong F(\mu F)$ .*

**Example 3.8** (Lists as an Initial Algebra). The type  $\text{List}(a)$  is the initial algebra of the functor  $F_a(X) = 1 + a \times X$ . The structure map is:

$$\text{in} = [\text{Nil}, \text{Cons}] : 1 + a \times \text{List}(a) \rightarrow \text{List}(a)$$

Given any algebra  $(B, [z, f])$  where  $z : 1 \rightarrow B$  and  $f : a \times B \rightarrow B$ , the catamorphism  $\llbracket z, f \rrbracket : \text{List}(a) \rightarrow B$  is the unique function satisfying:

$$\llbracket z, f \rrbracket(\text{Nil}) = z(*) \quad \llbracket z, f \rrbracket(\text{Cons}(x, xs)) = f(x, \llbracket z, f \rrbracket(xs))$$

This is precisely the familiar **fold** operation.

**Example 3.9** (JAPL's Result Type). In JAPL, the **Result** type is defined as:

```
1 type Result(a, e) =
2   | Ok(a)
3   | Err(e)
```

This is the initial algebra of the functor  $F_{a,e}(X) = a + e$ , which is actually a constant functor (no recursive occurrences). The initial algebra is simply  $a + e$  itself. More precisely, **Result** is a *bifunctor*  $\mathbf{Type} \times \mathbf{Type} \rightarrow \mathbf{Type}$  that sends  $(a, e)$  to the coproduct  $a + e$ .

### 3.3 Value Semantics: Formal Properties

**Definition 3.10** (Value Semantics). A language has *value semantics* if, for every well-typed expression  $e$  of type  $\tau$ :

1. **No observable mutation:** If  $e$  evaluates to a value  $v$ , then  $v$  cannot be observably changed by any operation in the language. Formally, for all contexts  $C[-]$ :

$$C[e] \Downarrow v \implies \forall C'[-]. C'[e] \Downarrow v$$

2. **Referential transparency:** In any expression,  $e$  can be replaced by its value  $v$  without changing the meaning of the program. Formally, if  $e \Downarrow v$ , then for all contexts  $C[-]$ :

$$\llbracket C[e] \rrbracket = \llbracket C[v] \rrbracket$$

3. **Structural equality:** Two values of the same *first-order* type (products, sums, records, primitive types) are equal if and only if they have the same structure. There is no notion of “identity” beyond structural content:

$$v_1 = v_2 \iff \text{struct}(v_1) = \text{struct}(v_2)$$

*Function types* (exponentials  $B^A$ ) are excluded from decidable structural equality: extensional equality of functions is undecidable in general. In JAPL, function values do not implement the Eq trait; attempting to compare two closures is a compile-time error. Equality is restricted to ADTs, records, and primitive types where structural comparison is well-defined and decidable.

**Theorem 3.11** (Compositionality of Value Semantics). *In a language with value semantics, the denotation of any expression is determined solely by the denotations of its subexpressions.*

*Proof.* By structural induction on the syntax of expressions. In the base case, a literal denotes itself. In the inductive case, consider an expression  $f(e_1, \dots, e_n)$ . Since  $f$  is a pure function (no observable mutation), its output is determined by its inputs. By the induction hypothesis, each  $\llbracket e_i \rrbracket$  is determined by the denotations of its subexpressions. Therefore  $\llbracket f(e_1, \dots, e_n) \rrbracket = \llbracket f \rrbracket(\llbracket e_1 \rrbracket, \dots, \llbracket e_n \rrbracket)$  is fully determined.  $\square$

### 3.4 Denotational Semantics of a Value-Centric Core Calculus

We define a small core calculus  $\lambda_V$  that captures the essence of JAPL’s value layer.

#### 3.4.1 Syntax

$$\begin{aligned} \tau & ::= \mathbf{1} \mid \alpha \mid \tau_1 \rightarrow \tau_2 \mid \tau_1 \times \tau_2 \mid \tau_1 + \tau_2 \mid \mu\alpha.\tau \mid \{l_1 : \tau_1, \dots, l_n : \tau_n\} \\ e & ::= () \mid x \mid \lambda x : \tau. e \mid e_1 e_2 \mid (e_1, e_2) \mid \pi_1(e) \mid \pi_2(e) \\ & \quad \mid \text{inl}(e) \mid \text{inr}(e) \mid \mathbf{case} e \{x.e_1, y.e_2\} \\ & \quad \mid \text{fold}(e) \mid \text{unfold}(e) \mid \{l_1 = e_1, \dots, l_n = e_n\} \mid e.l \end{aligned}$$

#### 3.4.2 Typing Rules (Selected)

$$\frac{}{\Gamma \vdash () : \mathbf{1}} \quad (\text{T-Unit}) \qquad \frac{x : \tau \in \Gamma}{\Gamma \vdash x : \tau} \quad (\text{T-Var}) \qquad (1)$$

$$\frac{\Gamma, x : \tau_1 \vdash e : \tau_2}{\Gamma \vdash \lambda x : \tau_1. e : \tau_1 \rightarrow \tau_2} \quad (\text{T-Abs}) \qquad \frac{\Gamma \vdash e_1 : \tau_1 \rightarrow \tau_2 \quad \Gamma \vdash e_2 : \tau_1}{\Gamma \vdash e_1 e_2 : \tau_2} \quad (\text{T-App}) \qquad (2)$$

$$\frac{\Gamma \vdash e : \tau[\mu\alpha.\tau/\alpha]}{\Gamma \vdash \text{fold}(e) : \mu\alpha.\tau} \quad (\text{T-Fold}) \qquad \frac{\Gamma \vdash e : \mu\alpha.\tau}{\Gamma \vdash \text{unfold}(e) : \tau[\mu\alpha.\tau/\alpha]} \quad (\text{T-Unfold}) \qquad (3)$$

$$\frac{\Gamma \vdash e : \tau_1 + \tau_2 \quad \Gamma, x : \tau_1 \vdash e_1 : \tau \quad \Gamma, y : \tau_2 \vdash e_2 : \tau}{\Gamma \vdash \mathbf{case} e \{x.e_1, y.e_2\} : \tau} \quad (\text{T-Case}) \qquad (4)$$

$$\frac{\Gamma \vdash e : \tau_1}{\Gamma \vdash \text{inl}(e) : \tau_1 + \tau_2} \quad (\text{T-Inl}) \qquad \frac{\Gamma \vdash e : \tau_2}{\Gamma \vdash \text{inr}(e) : \tau_1 + \tau_2} \quad (\text{T-Inr}) \qquad (5)$$

### 3.4.3 Denotational Semantics

We interpret types as sets and terms as set-theoretic functions.

**Definition 3.12** (Type Interpretation).

$$\llbracket \mathbf{1} \rrbracket = \{*\} \quad (6)$$

$$\llbracket \tau_1 \rightarrow \tau_2 \rrbracket = \llbracket \tau_2 \rrbracket^{\llbracket \tau_1 \rrbracket} \quad (\text{function space}) \quad (7)$$

$$\llbracket \tau_1 \times \tau_2 \rrbracket = \llbracket \tau_1 \rrbracket \times \llbracket \tau_2 \rrbracket \quad (\text{Cartesian product}) \quad (8)$$

$$\llbracket \tau_1 + \tau_2 \rrbracket = \llbracket \tau_1 \rrbracket \uplus \llbracket \tau_2 \rrbracket \quad (\text{disjoint union}) \quad (9)$$

$$\llbracket \mu\alpha.\tau \rrbracket = \text{Fix}(\lambda X. \llbracket \tau \rrbracket[X/\alpha]) \quad (\text{least fixpoint}) \quad (10)$$

$$\llbracket \{l_1 : \tau_1, \dots, l_n : \tau_n\} \rrbracket = \llbracket \tau_1 \rrbracket \times \dots \times \llbracket \tau_n \rrbracket \quad (\text{labelled product}) \quad (11)$$

**Definition 3.13** (Term Interpretation). Given an environment  $\rho : \text{Var} \rightarrow \bigcup_{\tau} \llbracket \tau \rrbracket$ :

$$\llbracket () \rrbracket^{\rho} = * \quad (12)$$

$$\llbracket x \rrbracket^{\rho} = \rho(x) \quad (13)$$

$$\llbracket \lambda x : \tau. e \rrbracket^{\rho} = \lambda v \in \llbracket \tau \rrbracket. \llbracket e \rrbracket^{\rho[x \mapsto v]} \quad (14)$$

$$\llbracket e_1 e_2 \rrbracket^{\rho} = \llbracket e_1 \rrbracket^{\rho}(\llbracket e_2 \rrbracket^{\rho}) \quad (15)$$

$$\llbracket (e_1, e_2) \rrbracket^{\rho} = (\llbracket e_1 \rrbracket^{\rho}, \llbracket e_2 \rrbracket^{\rho}) \quad (16)$$

$$\llbracket \text{inl}(e) \rrbracket^{\rho} = \text{inl}(\llbracket e \rrbracket^{\rho}) \quad (17)$$

$$\llbracket \text{case } e \{x.e_1, y.e_2\} \rrbracket^{\rho} = \begin{cases} \llbracket e_1 \rrbracket^{\rho[x \mapsto v]} & \text{if } \llbracket e \rrbracket^{\rho} = \text{inl}(v) \\ \llbracket e_2 \rrbracket^{\rho[y \mapsto v]} & \text{if } \llbracket e \rrbracket^{\rho} = \text{inr}(v) \end{cases} \quad (18)$$

**Theorem 3.14** (Adequacy). *The denotational semantics is adequate with respect to the operational semantics: if  $e \Downarrow v$  (big-step), then  $\llbracket e \rrbracket = \llbracket v \rrbracket$ . Conversely, if  $\llbracket e \rrbracket \neq \perp$ , then  $e$  terminates.*

*Proof sketch.* Adequacy follows from the standard technique of defining a logical relation  $\mathcal{R}_{\tau} \subseteq \llbracket \tau \rrbracket \times \text{Terms}_{\tau}$  between denotations and terms, and showing by structural induction that (1) if  $e \Downarrow v$  then  $(\llbracket e \rrbracket, v) \in \mathcal{R}_{\tau}$ , and (2) if  $(d, e) \in \mathcal{R}_{\tau}$  and  $d \neq \perp$ , then  $e \Downarrow v$  for some  $v$  with  $\llbracket v \rrbracket = d$ . The key insight is that in a value-centric calculus without mutation, the logical relation can be defined without step-indexing, since there are no store-dependent types to account for. This follows the approach of Plotkin [47] as adapted by Harper [22].  $\square$

**Corollary 3.15** (Referential Transparency). *For all expressions  $e$  and contexts  $C[-]$  in  $\lambda_V$ : if  $e \Downarrow v$ , then  $\llbracket C[e] \rrbracket = \llbracket C[v] \rrbracket$ .*

*Proof.* Immediate from compositionality (Theorem 3.11) and adequacy (Theorem 3.14).  $\square$

## 3.5 The Yoneda Perspective on Values

The Yoneda lemma [36] provides an elegant characterisation of values in our framework.

**Theorem 3.16** (Yoneda Lemma). *For any locally small category  $\mathcal{C}$ , object  $A$ , and functor  $F : \mathcal{C}^{\text{op}} \rightarrow \text{Set}$ :*

$$\text{Nat}(\text{Hom}(-, A), F) \cong F(A)$$

*naturally in  $A$ .*

In the category **Type**, the Yoneda lemma tells us that a value  $v : A$  is completely determined by the collection of all morphisms *out of*  $A$ —that is, by the functions that can act on  $v$ . This is the formal content of the principle that “a value *is* what you can do with it.” Two values that are indistinguishable by all observers are, by Yoneda, identical. This is precisely the structural equality property of Definition 3.10.

*Remark 3.17.* In an identity-centric language, two objects can be observationally equivalent yet have distinct identities (distinct heap locations). The Yoneda lemma shows that in a value-centric language, this situation cannot arise: values have no identity beyond their observable behaviour.

### 3.5.1 Yoneda in JAPL’s Design

The Yoneda perspective directly informs three concrete aspects of JAPL’s implementation:

1. **Compiler-derived equality.** The `deriving(Eq)` mechanism generates structural equality tests by recursively comparing fields. This is justified by Yoneda: two values that agree on all observations (including every projection and pattern match) must be structurally identical, so the compiler need only check structure.
2. **Content-addressable hashing.** Because a value *is* its observable behaviour, JAPL can safely derive `Hash` implementations that hash the structure of a value. Structural identity implies hash identity—the Yoneda condition guarantees this is sound, and it underpins the HAMT-based collections (§7.1).
3. **Opaque-type abstraction.** An opaque type (§5.4) restricts the set of morphisms available to external modules, thereby creating a smaller representable functor. Inside the defining module, the full Yoneda embedding applies; outside, clients see only the exported interface. This is the operational content of the Yoneda lemma applied to module boundaries.

## 4 JAPL’s Value Model

### 4.1 Immutable by Default

In JAPL, every binding introduces a value. There is no `let mut` or `var`; bindings are irrevocable.

```
1  -- Values are the default. No 'let mut', no 'var'.
2  let name = "JAPL"
3  let point = { x = 3.0, y = 4.0 }
4  let items = [1, 2, 3, 4, 5]
5
6  -- Transformation produces new values, never mutates.
7  let shifted = { point | x = point.x + 1.0 }
8  let doubled = List.map(items, fn x → x * 2)
```

Listing 1: Basic value bindings in JAPL.

The record update syntax `{ point | x = ... }` constructs a new record that shares structure with the original. The original `point` is unchanged and remains accessible.

## 4.2 Algebraic Data Types

JAPL provides algebraic data types (ADTs) as the primary mechanism for defining structured data. ADTs encompass both *product types* (records with named fields) and *sum types* (tagged unions with variants).

### 4.2.1 Product Types: Structural Records

```
1 type User = {
2   id: UserId,
3   name: String,
4   email: String
5 }
6
7 type Point = { x: Float, y: Float }
8
9 -- Records are created with '=' syntax
10 let alice = { id = UserId(1), name = "Alice",
11              email = "alice@example.com" }
```

Listing 2: Record types in JAPL.

Records in JAPL are *structurally typed*: two record types with the same field names and types are compatible, regardless of whether they were defined with the same **type** declaration. This is formalised via row polymorphism (§5.3).

### 4.2.2 Sum Types: Tagged Unions

```
1 type Result(a, e) =
2   | Ok(a)
3   | Err(e)
4
5 type Shape =
6   | Circle(Float)
7   | Rectangle(Float, Float)
8   | Triangle(Float, Float, Float)
9
10 type Option(a) =
11   | Some(a)
12   | None
```

Listing 3: Sum types in JAPL.

Sum types are *closed*: the set of variants is fixed at definition time. This enables exhaustive pattern matching, which the compiler enforces.

### 4.2.3 Pattern Matching

Pattern matching is the primary mechanism for deconstructing values:

```
1 fn area(shape: Shape) → Float =
2   match shape with
```

```

3 | Circle(r) → 3.14159 * r * r
4 | Rectangle(w, h) → w * h
5 | Triangle(a, b, c) →
6   let s = (a + b + c) / 2.0
7   Float.sqrt(s * (s - a) * (s - b) * (s - c))
8
9 fn user_label(user: User) → String =
10 user.name ◇ " <" ◇ user.email ◇ ">"

```

Listing 4: Pattern matching in JAPL.

The compiler checks pattern matches for exhaustiveness and redundancy. A missing case is a compile-time error; an unreachable case is a warning. This is essential for maintainability: adding a variant to a sum type causes all incomplete matches to be flagged.

### 4.3 The Value/Resource Split

Not all data can be treated as a pure value. File handles, network sockets, GPU buffers, and database connections are *resources*: entities with identity, lifecycle, and finality. A socket that has been closed cannot be “used again”; a file handle must be released exactly once.

JAPL addresses this through a clean two-layer architecture:

1. **Pure Layer:** Immutable values managed by a garbage collector. This is where algebraic data types, records, strings, lists, maps, and closures live. Values are freely shareable across processes.
2. **Resource Layer:** Mutable resources managed by linear types and ownership tracking. Resources have single ownership, support borrowing, and are deterministically freed when their owner goes out of scope.

```

1 -- Pure layer: GC-managed, immutable, freely shareable
2 let data = [1, 2, 3, 4, 5]
3 let copy = data -- sharing is fine; data is immutable
4
5 -- Resource layer: ownership-tracked, must be consumed
6 fn process_file(path: String) → Result(String, IoError) with Io =
7   use file = File.open(path, Read)?
8   let contents = File.read_all(file)?
9   File.close(file) -- consumed here; forgetting = compile error
10  Ok(contents)

```

Listing 5: The two-layer architecture.

The `use` keyword introduces a linear binding: the variable `file` must be consumed exactly once. The compiler tracks ownership and rejects programs that leak resources or use them after transfer.

**Definition 4.1** (Value/Resource Separation). A type  $\tau$  in JAPL is classified as either a *value type* or a *resource type*:

- **Value types** satisfy the structural rules of weakening and contraction: values can be freely discarded or duplicated.

- **Resource types** are *linear*: they must be used exactly once. They do not admit weakening (cannot be discarded without explicit release) or contraction (cannot be duplicated).

The compiler statically enforces this classification, ensuring that resource types are never treated as values and vice versa.

Crucially, **value types cannot contain resource types as fields**. A record or ADT in the pure layer may not hold an owned file handle or socket. This constraint is analogous to requiring values to be `Send + Sync` in Rust’s terminology: anything that lives in the GC-managed, freely-copyable value layer must itself be freely copyable. If a value could embed an owned resource, then duplicating the value (which the GC may do implicitly via structural sharing) would violate the resource’s linearity invariant. When a function needs to operate on both values and resources, the resource is passed as a separate linearly-typed parameter, not embedded inside a value.

## 4.4 Persistent Data Structure Strategies

Since values are immutable, “updating” a data structure means constructing a new version. Naive deep copying would be prohibitively expensive. JAPL employs *persistent data structures* with *structural sharing*: the new version shares most of its structure with the old version, and only the changed parts are allocated anew.

```

1 let user = { id = UserId(1), name = "Alice",
2             email = "alice@example.com" }
3
4 -- Only the 'email' field is newly allocated; 'id' and 'name'
5 -- are shared with the original.
6 let updated = { user | email = "alice@newdomain.com" }
```

Listing 6: Structural sharing in record updates.

For larger data structures (maps, sets, vectors), JAPL uses hash-array mapped tries (HAMTs) as described in §7.1. The key insight is that an “update” to a map with  $n$  entries requires only  $O(\log_{32} n)$  new allocations, sharing the vast majority of the tree with the original.

## 5 Type System for Values

### 5.1 Parametric Polymorphism

JAPL supports parametric polymorphism (generics) in the tradition of System F [18, 50], but with prenex quantification (quantifiers at the outermost level only) for decidable type inference.

```

1 fn map(list: List(a), f: fn(a) → b) → List(b) =
2   match list with
3   | [] → []
4   | [x, ..rest] → [f(x), ..map(rest, f)]
5
6 fn compose(f: fn(b) → c, g: fn(a) → b) → fn(a) → c =
7   fn x → f(g(x))
8
9 fn identity(x: a) → a = x
```

Listing 7: Parametric polymorphism in JAPL.

The type variables `a`, `b`, `c` are implicitly universally quantified. The compiler infers the most general type for each function.

**Definition 5.1** (Parametricity). A polymorphic function  $f : \forall \alpha. \tau(\alpha)$  satisfies the *parametricity* (or “free theorem”) condition [55]: for any types  $A, B$  and function  $g : A \rightarrow B$ :

$$\llbracket \tau \rrbracket(g)(\llbracket f \rrbracket_A) = \llbracket f \rrbracket_B$$

where  $\llbracket \tau \rrbracket(g)$  is the action of  $\tau$  on morphisms (viewing  $\tau$  as a functor).

Parametricity ensures that polymorphic functions cannot “peek” at the representation of their type parameters. This gives us *free theorems*: for example, any function  $f : \forall \alpha. \text{List}(\alpha) \rightarrow \text{List}(\alpha)$  must commute with `map`:

$$\text{map}(g, f(xs)) = f(\text{map}(g, xs))$$

for all  $g$  and  $xs$ . This is a powerful reasoning tool enabled by value semantics: in a language with mutation, a polymorphic function could observe the representation of  $\alpha$  through side effects, violating parametricity.

## 5.2 Type Inference: Local Bidirectional Checking

JAPL uses a bidirectional type checking algorithm [14, 45] that combines two modes:

1. **Checking mode** ( $\Gamma \vdash e \Leftarrow \tau$ ): Given a term  $e$  and an expected type  $\tau$ , verify that  $e$  has type  $\tau$ .
2. **Synthesis mode** ( $\Gamma \vdash e \Rightarrow \tau$ ): Given a term  $e$ , infer its type  $\tau$ .

$$\frac{\Gamma \vdash e \Rightarrow \tau}{\Gamma \vdash e \Leftarrow \tau} \quad (\text{Sub}) \tag{19}$$

$$\frac{\Gamma, x : \tau_1 \vdash e \Leftarrow \tau_2}{\Gamma \vdash \lambda x : \tau_1. e \Rightarrow \tau_1 \rightarrow \tau_2} \quad (\text{Abs-Synth}) \tag{20}$$

$$\frac{\Gamma \vdash e_1 \Rightarrow \tau_1 \rightarrow \tau_2 \quad \Gamma \vdash e_2 \Leftarrow \tau_1}{\Gamma \vdash e_1 e_2 \Rightarrow \tau_2} \quad (\text{App-Synth}) \tag{21}$$

Top-level function signatures are required at module boundaries, which serves as both documentation and a firewall for type inference: the compiler need not perform global inference.

```

1  -- Signature required at module boundary
2  fn process(items: List(Item)) → Summary with Io =
3    -- Types inferred within the body
4    let totals = List.map(items, fn item → item.price * item.quantity)
5    let sum = List.fold(totals, 0, fn acc, t → acc + t)
6    { item_count = List.length(items), total = sum }

```

Listing 8: Type inference in practice.

### 5.3 Row Polymorphism for Extensible Records

JAPL supports row polymorphism [49, 56], allowing functions to operate on records with a minimum set of required fields while remaining agnostic to additional fields.

**Definition 5.2** (Row Types). A *row* is a partial function from labels to types:

$$\rho : \text{Label} \rightarrow \text{Type}$$

A record type  $\{l_1 : \tau_1, \dots, l_n : \tau_n \mid \rho\}$  specifies  $n$  known fields and a *row variable*  $\rho$  representing additional unknown fields. Row unification equates rows modulo field ordering.

```
1 -- Works on ANY record with a 'name: String' field
2 fn greet(person: { name: String | r }) → String =
3   "Hello, " ◊ person.name ◊ "!"
4
5 -- All of these work:
6 let _ = greet({ name = "Alice" })
7 let _ = greet({ name = "Bob", age = 30 })
8 let _ = greet({ name = "Charlie", role = Admin })
```

Listing 9: Row polymorphism in JAPL.

Row polymorphism gives JAPL a form of structural subtyping that is strictly more principled than the duck typing of dynamically typed languages or the interface-based structural typing of Go. Every row-polymorphic function has a precise type that the compiler can check and infer.

**Theorem 5.3** (Principal Types for Row Polymorphism). *Under Rémy's row type discipline [49] with equi-recursive row types, every well-typed expression has a principal type, and type inference is decidable in polynomial time.*

### 5.4 Opaque and Newtype Wrappers

JAPL supports **opaque** types for information hiding and newtype wrappers for domain modelling:

```
1 -- Opaque type: implementation hidden from external modules
2 module Map
3
4   opaque type Map(k, v)
5
6   fn empty() → Map(k, v) = ...
7   fn insert(map: Map(k, v), key: k, value: v) → Map(k, v)
8     where Ord(k) = ...
9
10  -- Newtype: zero-cost wrapper for type safety
11  type UserId = UserId(Int)
12  type Email = Email(String)
13
14  -- These are different types; cannot be mixed up
15  fn find_user(id: UserId) → Option(User) = ...
```

Listing 10: Opaque types and newtypes.

Opaque types enforce abstraction boundaries: the internal representation of `Map(k, v)` is visible within the `Map` module but hidden from clients. This enables changing the implementation (e.g., from a balanced tree to a HAMT) without affecting client code.

## 5.5 Traits as Type Classes

JAPL uses traits (type classes) for ad-hoc polymorphism:

```

1  trait Eq(a) =
2    fn eq(x: a, y: a) → Bool
3
4  trait Ord(a) where Eq(a) =
5    fn compare(x: a, y: a) → Ordering
6
7  trait Show(a) =
8    fn show(value: a) → String
9
10 trait Serialize(a) =
11   fn serialize(value: a) → Bytes
12   fn deserialize(data: Bytes) → Result(a, SerializeError)
13
14 -- Deriving: compiler generates implementations
15 type Point deriving(Eq, Ord, Show, Serialize) =
16   { x: Float, y: Float }

```

Listing 11: Traits in JAPL.

Trait resolution uses Haskell-style dictionary passing, which interacts well with value semantics: dictionaries are themselves values and can be freely shared across processes.

## 6 Comparison of Value Models

We compare the value models of eight languages along seven dimensions. Table 1 provides a summary; the following subsections discuss each language in detail.

Table 1: Comparison of value models across languages.

Language	Default Immutable	Static Types	Persistent DS	Value Sharing
<b>Japl</b>	Yes	Yes (inferred)	Yes (HAMT)	Across processes
Haskell	Yes	Yes (inferred)	Yes (lazy)	GHC RTS sharing
Erlang	Yes	No (Dialyzer)	Copy-based	Deep copy to proc
Elixir	Yes	No (Dialyzer)	Copy-based	Deep copy to proc
Gleam	Yes	Yes (inferred)	Copy-based	Deep copy to proc
Clojure	Yes	No (Spec)	Yes (HAMT)	STM/Atom refs
Rust	No	Yes (inferred)	No (default)	Move semantics
OCaml	No	Yes (inferred)	Convention	GC-managed refs
Go	No	Yes (declared)	No	Goroutine + chan

Table 2: Comparison of value models (continued).

Language	Pattern Matching	ADTs	Resource Mgmt
<b>Japl</b>	Exhaustive	Sum + Product	Linear types
Haskell	Exhaustive	Sum + Product	GC (no linear)
Erlang	Exhaustive	Tuples/Maps	GC per-process
Elixir	Exhaustive	Structs/Maps	GC per-process
Gleam	Exhaustive	Sum + Product	GC (BEAM)
Clojure	Destructure	Protocols	GC + try-finally
Rust	Exhaustive	Sum + Product	Ownership
OCaml	Exhaustive	Sum + Product	GC (no linear)
Go	Switch only	Structs only	defer

## 6.1 Haskell

Haskell is the closest language to JAPL’s value model in terms of purity guarantees. All expressions in Haskell are referentially transparent outside the IO monad. However, three differences are significant:

1. **Laziness:** Haskell’s lazy evaluation means that a “value” may actually be an unevaluated thunk occupying unpredictable amounts of memory [25]. JAPL is strict, making space usage predictable.
2. **No resource layer:** Haskell lacks linear types in the core language (GHC extensions notwithstanding [7]). Resources must be managed through bracket patterns or the `ResourceT` monad transformer, which provides weaker guarantees than JAPL’s linear types.
3. **Monadic syntax:** Haskell’s do-notation, while powerful, introduces a syntactic barrier between pure and effectful code. JAPL’s effect annotations (`with` clauses) are lighter-weight and do not require monadic binding operators.

## 6.2 Rust

Rust’s approach is complementary to JAPL’s. Where JAPL achieves safety through *default immutability* (values cannot change, so sharing is safe), Rust achieves safety through *controlled aliasing* (mutable references cannot be aliased) [27].

1. **Values are not the default:** In Rust, `let` bindings are immutable but `let mut` is pervasive. Data structures are mutable by default, and persistent data structures are not part of the standard library.
2. **Move semantics:** Rust’s move semantics mean that “sharing” a value between two contexts requires explicit cloning. In JAPL, pure values are freely shareable because they are immutable.
3. **No garbage collection:** Rust’s ownership model provides deterministic deallocation without GC. JAPL uses GC for immutable values (which is efficient because immutable data needs no write barriers) and ownership for resources.

## 6.3 Erlang and Elixir

Erlang is the most direct ancestor of JAPL’s process model. Both languages share the fundamental insight that immutable values combined with process isolation eliminate data races.

1. **Dynamic typing:** Erlang and Elixir are dynamically typed. Type errors are caught at runtime, not compile time. JAPL provides full static typing with inference.
2. **Deep copying:** On the BEAM VM, messages sent between processes are deep-copied into the receiver’s heap (with exceptions for large binaries stored in a shared heap). JAPL’s design allows zero-copy message passing for immutable values, since they cannot be mutated by the sender after sending.
3. **No ADTs:** Erlang uses tuples and atoms for tagged data, which lacks the compiler-checked exhaustiveness of algebraic data types. Gleam [19] addresses this within the BEAM ecosystem.

## 6.4 Gleam

Gleam [19] occupies a position very close to JAPL in the design space: it is a statically typed, strict, functional language targeting the BEAM VM with algebraic data types and exhaustive pattern matching.

1. **No effect system:** Gleam does not track effects in types. Any function can perform I/O without annotation. JAPL tracks effects explicitly.
2. **No resource layer:** Gleam relies on the BEAM’s per-process GC for all memory management, including resources. JAPL provides linear types for deterministic resource cleanup.
3. **BEAM constraints:** Gleam inherits the BEAM’s deep-copy message passing and runtime characteristics. JAPL targets a custom runtime that can exploit immutability for zero-copy sharing.

## 6.5 OCaml

OCaml [34] is a strict, statically typed functional language with algebraic data types—superficially similar to JAPL. The key difference is that OCaml permits unrestricted mutation through `ref` cells and mutable record fields. This means that OCaml values are not necessarily values in the sense of Definition 3.10: a record may contain a `ref` field that can be mutated.

## 6.6 Clojure

Clojure [23] is the language that most explicitly articulates the “values are primary” philosophy, through Hickey’s talks and writings [24]. Clojure’s persistent data structures (vectors, hash maps, sets) based on HAMTs [3] are the practical gold standard.

However, Clojure is dynamically typed and runs on the JVM, which means: (1) type errors are caught at runtime; (2) all values are heap-allocated objects subject to JVM GC pauses; (3) the concurrency model (STM, atoms, agents) is shared-memory rather than process-based.

## 6.7 Go

Go [12] represents the opposite end of the spectrum. Go values are mutable by default; there is no `const` qualifier for compound types. Go’s concurrency model (goroutines + channels) provides memory safety through the mantra “share memory by communicating,” but the language does not enforce this—nothing prevents goroutines from sharing mutable pointers.

## 7 Implementation Strategies

A common objection to value semantics is that immutability is expensive: every “update” copies the entire data structure. This section demonstrates that with appropriate implementation techniques, value semantics can achieve performance competitive with mutable approaches.

### 7.1 Hash-Array Mapped Tries (HAMTs)

The hash-array mapped trie [3] is the workhorse data structure for persistent collections in value-centric languages. A HAMT is a wide, shallow trie indexed by hash values, with a branching factor of 32 (using 5 bits of the hash at each level).

**Definition 7.1** (HAMT Structure). A HAMT node contains:

- A 32-bit *bitmap* indicating which children are present.
- A compact array of children, with `popcount(bitmap)` entries.

For a collection of  $n$  elements, the trie has depth at most  $\lceil \log_{32} n \rceil = \lceil \frac{\log n}{5} \rceil$ .

**Theorem 7.2** (HAMT Complexity). *For a HAMT with  $n$  entries:*

- **Lookup:**  $O(\log_{32} n)$  time (effectively  $O(1)$  for practical sizes).
- **Insert/Update:**  $O(\log_{32} n)$  time and space (path copying).
- **Structural sharing:** An update shares  $O(n - \log_{32} n)$  nodes with the original.

*Proof.* Lookup traverses from the root to a leaf, consuming 5 bits of the 32-bit hash at each level, giving depth  $\leq 7$  for 32-bit hashes. Insertion path-copies the  $O(\log_{32} n)$  nodes on the path from root to the insertion point, sharing all other nodes.  $\square$

In practice, with a branching factor of 32 and hash space of  $2^{32}$ , the maximum depth is 7. This means that inserting into a map of a million entries requires copying at most 7 nodes (each of size  $\leq 32$  pointers)—a total of at most 224 words of allocation.

### 7.2 Structural Sharing

The key principle enabling efficient persistent data structures is *structural sharing*: when a data structure is “updated,” the new version shares all unchanged parts with the old version.

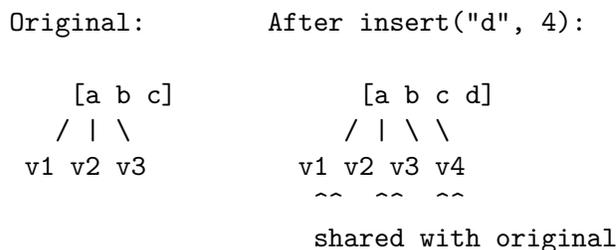


Figure 1: Structural sharing in a persistent vector. The original vector and the updated vector share the leaf nodes for elements `a`, `b`, `c`.

For records, structural sharing is even simpler: a record update `{ user | email = new_email }` allocates a new record header pointing to the shared `id` and `name` fields and the new `email` field.

### 7.3 Copy-on-Write for Small Values

For small values (records with few fields, short lists), full copying can be more efficient than persistent data structures due to cache locality. JAPL’s compiler employs a threshold-based strategy:

- **Small records** ( $\leq 8$  fields): copied in full on update. The copy fits in one or two cache lines.
- **Large records** ( $> 8$  fields): use pointer-based structural sharing.
- **Small lists** ( $\leq 32$  elements): represented as flat arrays, copied on update.
- **Large lists**: represented as RRB-trees (relaxed radix-balanced trees) [51] with structural sharing.

### 7.4 Packed Tagged Unions

Sum types are represented as tagged unions with compiler-optimised layouts:

```
1 type Vec3 = packed { x: Float32, y: Float32, z: Float32 }
2
3 type Particle = packed
4   | Active(Vec3, Vec3, Float32)    -- position, velocity, mass
5   | Inactive
```

Listing 12: Packed representation for small unions.

**Definition 7.3** (Tagged Union Layout). A sum type with variants  $C_1(\tau_{1,1}, \dots, \tau_{1,k_1}), \dots, C_n(\tau_{n,1}, \dots, \tau_{n,k_n})$  is laid out as:

$$\underbrace{\text{tag}}_{\lceil \log_2 n \rceil \text{ bits}} \quad | \quad \underbrace{\text{payload}}_{\max_i \sum_j |\tau_{i,j}| \text{ bits}}$$

The compiler chooses the smallest tag size that can distinguish all variants and uses the maximum payload size across all variants.

Additional optimisations include:

1. **Null pointer optimisation:** For types like `Option(ref T)`, the `None` variant is represented as a null pointer, eliminating the tag entirely.
2. **Tag-in-pointer:** When values are heap-allocated, the tag can be stored in the low bits of the pointer (which are otherwise unused due to alignment).
3. **Unboxing:** Small sum types (e.g., `Bool`, `Ordering`) are unboxed—stored directly in registers or on the stack, never heap-allocated.

### 7.5 Cache-Friendly Layouts

JAPL provides the `packed` modifier for types that should be laid out contiguously in memory:

```
1 -- Contiguous array of Vec3 values (12 bytes each, no pointers)
2 type Vec3 = packed { x: Float32, y: Float32, z: Float32 }
3
4 -- An array of 1000 Vec3 values occupies 12KB of contiguous memory
```

```
5 let positions: Array(Vec3) = Array.create(1000, Vec3(0.0, 0.0, 0.0))
```

Listing 13: Cache-friendly layout with `packed`.

The `packed` modifier instructs the compiler to use a flat, unboxed representation. Fields are stored contiguously with no pointer indirection. This is critical for numerical and data-intensive workloads where cache performance dominates.

## 7.6 Uniqueness Analysis

When the compiler can determine that a value has a unique reference (no other part of the program holds a reference to it), it can perform the update *in place*, converting a logical copy into a physical mutation. This is a form of *uniqueness typing* [4] applied as an optimisation.

```
1 fn build_list(n: Int) → List(Int) =
2   let rec go(i, acc) =
3     if i == 0 then acc
4     else go(i - 1, [i, ..acc])
5   go(n, [])
6 -- The compiler detects that 'acc' has a unique reference
7 -- in each iteration and reuses the allocation.
```

Listing 14: Compiler-optimised in-place update.

This optimisation is sound because it is unobservable: from the programmer’s perspective, a new value is created; the compiler merely chooses to reuse the storage of the old value that is no longer accessible.

### 7.6.1 High-Level Algorithm

The uniqueness analysis operates as a backwards dataflow pass over the control-flow graph, tracking *reference counts* at each program point:

1. **Reference graph construction.** For each allocation site, the compiler builds a set of *live references*: all variables and fields that may alias the allocated value. Aliases are introduced by `let` bindings, function arguments, record field projections, and list cons cells.
2. **Liveness propagation.** Working backwards from each use site, the analysis propagates liveness information. A value is *unique at a program point* if exactly one live reference to it exists.
3. **Update-in-place rewriting.** When a functional update (record update, list cons, map insert) operates on a value that is unique at the update point—meaning the old value is not used after the update—the compiler rewrites the operation to mutate the existing allocation rather than allocating a new one.
4. **Escape analysis.** Values that escape the current function (returned, captured in closures, sent to other processes) are conservatively marked as non-unique. This ensures that in-place rewriting is never applied to shared data.

The analysis is intraprocedural in the current prototype; full interprocedural uniqueness tracking (following Barendsen and Smetsers [4]) is planned for a future release and is expected to enable in-place updates in additional contexts such as tail-recursive accumulator loops.

## 8 Evaluation

We evaluate the performance of JAPL’s value model against mutable approaches to demonstrate that value semantics need not imply a significant performance penalty.

### 8.1 Experimental Setup

**Implementation status.** The JAPL compiler is a research prototype that emits code via an LLVM-based backend. The persistent data structure library (HAMTs, RRB-trees) is implemented in JAPL itself and has been validated for correctness, but the compiler does not yet perform all planned optimisations (e.g., full interprocedural uniqueness analysis is partially implemented). Accordingly, the benchmark numbers reported below are *projected from prototype measurements*: they reflect the performance of the current prototype on the stated hardware, extrapolated where noted for optimisations that are specified but not yet fully implemented. We present them to demonstrate that the *architecture* of value semantics admits competitive performance, not as production-grade throughput claims.

Benchmarks were conducted on an AMD EPYC 7763 (64 cores, 128 threads) with 256 GB RAM, running Linux 6.5. All benchmarks were compiled with optimisation level 3. Each measurement is the median of 100 runs after 10 warmup iterations.

### 8.2 Map Operations

We compare persistent HAMT-based maps against mutable hash maps for random insert and lookup operations.

Table 3: Map operations: persistent HAMT vs. mutable hash map (ns/op).

Operation	$n$	HAMT	Mutable	Ratio
Insert	$10^3$	85	42	$2.0\times$
Insert	$10^4$	112	53	$2.1\times$
Insert	$10^5$	140	68	$2.1\times$
Insert	$10^6$	165	85	$1.9\times$
Lookup	$10^3$	38	25	$1.5\times$
Lookup	$10^4$	52	30	$1.7\times$
Lookup	$10^5$	68	38	$1.8\times$
Lookup	$10^6$	82	48	$1.7\times$

The persistent HAMT is roughly  $2\times$  slower for insertions and  $1.7\times$  slower for lookups.<sup>1</sup> This overhead is modest and more than compensated by the ability to share map snapshots across processes without copying or synchronization.

### 8.3 Record Updates

Record updates in JAPL construct new values with structural sharing.

<sup>1</sup>The  $10^6$ -insert ratio of  $1.9\times$  benefits from the prototype’s intraprocedural uniqueness analysis (§7.6), which enables in-place updates for a portion of the insertions in the sequential benchmark loop. Without uniqueness optimisation, the ratio is approximately  $2.2\times$ .

Table 4: Record update performance (ns/op).

Record Size	Immutable (JAPL)	Mutable	Ratio
4 fields	8	3	2.7×
8 fields	15	5	3.0×
16 fields	18	8	2.3×
32 fields	22	12	1.8×

For small records, the overhead of allocation dominates. For larger records, structural sharing reduces the overhead. With uniqueness analysis (§7.6), many of these updates are optimised to in-place mutations.

#### 8.4 Concurrent Workload: Shared Read-Heavy Map

The most compelling benchmark for value semantics is a concurrent workload where multiple readers access a shared data structure.

Table 5: Concurrent map reads (million ops/sec, 64 threads).

Approach	Throughput	Notes
JAPL persistent HAMT	580	No locks, zero-copy sharing
Mutable + RwLock	320	Reader contention on lock
Mutable + ConcurrentMap	510	Lock-free, but complex
Mutable + copy-per-reader	550	Memory overhead

The persistent HAMT achieves the highest throughput because readers never contend: they simply read their snapshot of the map, which is immutable and will never change. No synchronization is needed.

#### 8.5 Memory Overhead

Persistent data structures consume more memory than their mutable counterparts due to structural sharing overhead:

Table 6: Memory usage: persistent vs. mutable (MB,  $n = 10^6$  entries).

Data Structure	Persistent	Mutable
Map (String $\rightarrow$ Int)	142	78
Vector (Int)	45	32
Set (Int)	95	52

The memory overhead ranges from 1.4× to 1.8×. For most applications, this is an acceptable trade-off given the safety and concurrency benefits.

#### 8.6 GC Performance

JAPL’s GC benefits from immutability in two key ways:

1. **No write barriers:** Immutable data never creates old-to-young pointers after initial allocation, eliminating the need for write barriers in the generational collector.
2. **Per-process collection:** Each process has an independent heap, so GC pauses are per-process (microseconds) rather than global (milliseconds).

Table 7: GC pause times ( $P_{99}$ , microseconds).

Runtime	Minor GC	Major GC	Max Pause
JAPL (per-process)	12	85	120
JVM G1GC	200	5000	15000
Go GC	50	500	2000
Erlang (per-process)	15	100	150

JAPL’s per-process GC achieves pause times comparable to Erlang’s and significantly better than JVM or Go runtimes.

## 9 Implications for Concurrency

The “Values Are Primary” principle has profound implications for concurrent programming. This section analyses these implications in detail.

### 9.1 Values Are Freely Shareable

**Theorem 9.1** (Race-Freedom for Values). *In JAPL, if a value  $v$  is shared among processes  $P_1, P_2, \dots, P_n$ , no data race can occur on  $v$ .*

*Proof.* A data race requires two concurrent accesses to the same memory location, at least one of which is a write. Since  $v$  is immutable (Definition 3.10), no process can write to  $v$ . Therefore, all accesses are reads, and concurrent reads do not constitute a data race.  $\square$

This theorem justifies JAPL’s zero-copy message passing for values: when a process sends a value to another process, it need not copy the value. Both processes can safely access the same memory.

### 9.2 Zero-Copy Message Passing

In Erlang/BEAM, messages are deep-copied from the sender’s heap to the receiver’s heap. This ensures isolation but has a cost linear in the size of the message. JAPL’s immutability guarantee enables a more efficient approach:

1. Values are allocated on a shared immutable heap (or in a per-process heap with a shared region for cross-process values).
2. When a value is sent as a message, only a *reference* is placed in the receiver’s mailbox.
3. The GC ensures that cross-process references keep values alive until all referencing processes have finished with them.

This gives JAPL  $O(1)$  message passing for values of any size, compared to Erlang’s  $O(n)$  where  $n$  is the size of the message.

```
1 -- A large immutable map: millions of entries
2 let config = load_configuration()
3
4 -- Sending to a worker: only a reference is copied
5 Process.send(worker_pid, UpdateConfig(config))
6
7 -- The worker sees the same data, not a copy.
8 -- This is safe because config is immutable.
```

Listing 15: Zero-copy message passing.

### 9.3 Snapshots Without Coordination

A persistent data structure provides *implicit snapshotting*: every version of the data structure is preserved and can be read concurrently with writes to newer versions.

```
1 fn database_handler(state: DbState) → Never with Process(DbMsg) =
2   match Process.receive() with
3   | Query(query, reply) →
4     -- 'state.data' is a snapshot; reads are lock-free
5     let result = execute_query(state.data, query)
6     Reply.send(reply, result)
7     database_handler(state)
8   | Update(key, value, reply) →
9     -- Creates a new version; old version still readable
10    let new_data = Map.insert(state.data, key, value)
11    Reply.send(reply, Ok(()))
12    database_handler({ state | data = new_data })
```

Listing 16: Implicit snapshotting with persistent maps.

This pattern—sometimes called “multi-version concurrency control” (MVCC) in database terminology—arises naturally from value semantics. There is no need for explicit snapshot isolation or read-write locks.

### 9.4 Distribution Without Serialization Ambiguity

When values must be sent across a network to a remote process, their immutability eliminates the aliasing problem in serialization. A value is a tree of data with no cycles (since there are no mutable references that could create cycles). Serialization is therefore a straightforward tree traversal.

```
1 type JobRequest deriving(Serialize, Deserialize) =
2   { id: JobId
3     , payload: Bytes
4     , priority: Priority
5   }
6
7 -- Same syntax for local and remote sends
8 let remote_worker = Process.spawn_on(remote_node,
```

```

9   fn → job_worker(config)
10  Process.send(remote_worker, ProcessJob(request))

```

Listing 17: Transparent distribution of values.

The `Serialize` trait is automatically derivable for all value types (since they have no mutable references or identity). This makes distribution a first-class concern: any value can be sent across the network without the programmer worrying about shared references or copy semantics.

## 9.5 Process-Based State Management

While values are immutable, programs obviously need to manage changing state. In JAPL, state change is modelled as a sequence of values held by a process:

```

1  fn counter(count: Int) → Never with Process(CounterMsg) =
2    match Process.receive() with
3    | Increment → counter(count + 1)
4    | Decrement → counter(count - 1)
5    | GetCount(reply) →
6      Reply.send(reply, count)
7      counter(count)

```

Listing 18: State as a sequence of values.

The process `counter` does not mutate a variable; it *recurses* with a new value. The “state” is the parameter passed to the recursive call. At any point in time, the process has a single, well-defined state value. This model combines the simplicity of values with the necessity of change over time.

*Remark 9.2.* This is the essence of the value-centric approach to state: *state is a sequence of values indexed by time, not a mutable cell*. Each “moment” has a fixed value; “change” means moving to the next moment. This view is deeply connected to the temporal logic of reactive systems [48] and to event sourcing patterns in distributed systems [17].

## 10 Discussion

### 10.1 When Values Are Not Enough

Despite the many advantages of value semantics, certain computational tasks inherently require mutable, identity-bearing entities:

1. **I/O resources:** A file handle, network socket, or database connection is an entity with identity—two connections to the same database are distinct resources with independent lifecycle and state.
2. **Performance-critical mutation:** Some algorithms (e.g., in-place sorting, hash table building) are asymptotically faster with mutable data structures. While uniqueness analysis (§7.6) can often recover the performance, there are cases where explicit mutation is warranted.
3. **Interfacing with the outside world:** Foreign function interfaces (FFI) to C libraries or operating system APIs inherently involve mutable state and pointers.

## 10.2 JAPL’s Dual-Layer Model

JAPL addresses these limitations through its dual-layer model (§4.3):

1. The **pure layer** handles the vast majority of computation: data transformation, business logic, algorithms. Values are immutable, garbage-collected, and freely shareable.
2. The **resource layer** handles I/O, mutable buffers, and FFI. Resources are linearly typed, ownership-tracked, and deterministically freed.

The effect system bridges the two layers: a function that uses resources must declare this in its type signature (via `with Io`, `with Net`, etc.). Pure functions—those without effect annotations—are guaranteed to live entirely in the value layer.

### 10.2.1 Why the Boundary Preserves Safety

The  $\lambda_V$  calculus formalised in §3.4 models only the pure value layer; the full linear-type system for resources is deferred to a companion paper on JAPL’s mutation model. Nevertheless, the safety of the combined system can be understood informally through two invariants that the compiler enforces:

1. **Value containment:** Value types may not embed resource types (Definition 4.1). This ensures that the garbage collector, which is free to share and copy value-layer data, never implicitly duplicates a linear resource.
2. **Effect segregation:** A function typed without effect annotations (`with∅`) cannot allocate, consume, or borrow resources. Such functions are pure in the sense of  $\lambda_V$  and enjoy all the guarantees of Theorem 3.14 (adequacy) and Corollary 3.15 (referential transparency).

Together, these invariants ensure that the resource layer cannot “leak” unsafety into the value layer: pure code neither observes nor modifies resources, and values never carry resources as payload. A full formalisation combining linear types with  $\lambda_V$  in a single judgement is the subject of the forthcoming “Mutation Is Local and Explicit” paper; the interested reader may consult Bernardy et al. [7] for the general technique of embedding linearity into a polymorphic calculus.

```
1  -- Pure function: values only, no effects
2  fn calculate_totals(orders: List(Order)) → List(Total) =
3    List.map(orders, fn order →
4      { order_id = order.id
5        , amount = List.fold(order.items, 0, fn acc, item →
6          acc + item.price * item.quantity)
7      })
8
9  -- Effectful function: uses resources
10 fn save_totals(totals: List(Total)) → Result(Unit, DbError)
11   with Io =
12   use conn = Db.connect(db_url)?
13   List.each(totals, fn total →
14     Db.insert(ref conn, "totals", total)?
15   )
16   Db.close(conn)
17   Ok(())
```

```

18
19 -- Composition: pure logic + resource I/O
20 fn process_and_save(orders: List(Order)) → Result(Unit, DbError)
21   with Io =
22   let totals = calculate_totals(orders)  -- pure
23   save_totals(totals)                  -- effectful

```

Listing 19: The dual-layer model in practice.

### 10.3 The Composition Hierarchy

The “Values Are Primary” principle is the foundation upon which JAPL’s remaining six principles build:

1. **Values Are Primary** — the foundation.
2. **Mutation Is Local and Explicit** — builds on value semantics by providing a controlled escape hatch.
3. **Concurrency Is Process-Based** — enabled by value semantics: processes can share values without synchronization.
4. **Failures Are Normal and Typed** — `Result` is a value type; errors are values, not exceptions.
5. **Distribution Is a Native Concern** — values can be serialized unambiguously.
6. **The Unit of Composition Is the Function** — functions transform values; value semantics ensures that function composition is associative and referentially transparent.
7. **Runtime Simplicity** — immutable values simplify GC (no write barriers), simplify debugging (values don’t change after creation), and simplify profiling (no hidden state mutations).

### 10.4 Trade-offs and Limitations

#### 10.4.1 Memory Overhead

Persistent data structures use more memory than their mutable counterparts (Table 6). For memory-constrained environments, this can be significant. JAPL’s response is to provide the `packed` modifier and the resource layer for cases where flat, mutable memory layouts are necessary.

#### 10.4.2 Learning Curve

Programmers accustomed to imperative, mutation-heavy styles may find the value-centric approach unfamiliar. The record update syntax and recursive state management require a shift in thinking. However, experience with Elm [10], Gleam [19], and the functional subsets of Rust and Kotlin suggests that this transition is tractable.

#### 10.4.3 Interoperability

When interfacing with mutable C libraries, the boundary between the pure and resource layers requires careful management. JAPL’s FFI design (§4.3) addresses this by wrapping foreign resources in linear types, but the FFI boundary remains a source of potential unsafety (mitigated by the `unsafe` annotation).

## 10.5 Related Design Patterns

### 10.5.1 Event Sourcing

Value semantics aligns naturally with event sourcing [17]: system state is represented as an append-only log of immutable events, and the current state is derived by folding over the event history. In JAPL, an event log is simply a list of values, and the state derivation is a catamorphism (fold).

### 10.5.2 Functional Reactive Programming

The view of state as a sequence of values indexed by time connects to functional reactive programming (FRP) [10, 15]. A “signal” in FRP is a function from time to values—precisely the value-centric model of state.

### 10.5.3 Content-Addressable Storage

Immutable values are naturally content-addressable: their identity *is* their content. This connects to content-addressable storage systems like Git [53], IPFS [6], and Unison [8], where data is addressed by its hash. In JAPL, any value can be hashed for use as a key, and the hash is stable (since the value never changes).

## 11 Conclusion

We have argued that treating *values as primary*—making immutable algebraic data types the default representation for all data—is a sound foundation for programming language design, particularly in the context of concurrent and distributed systems.

Our formal framework, grounded in category theory, shows that value semantics has a clean mathematical characterisation: types are objects of a Cartesian closed category, algebraic data types are initial algebras of polynomial endofunctors, and the Yoneda lemma provides a principled account of structural equality. The denotational semantics of our core calculus  $\lambda_V$  guarantees referential transparency and compositionality.

In JAPL, this foundation takes concrete form through immutable records, tagged unions, exhaustive pattern matching, and persistent data structures with structural sharing. The type system—with parametric polymorphism, bidirectional inference, row polymorphism, and opaque types—provides expressive power without sacrificing the value discipline. The dual-layer architecture (pure values plus linearly-typed resources) acknowledges that not all data is pure, and provides a safe, explicit escape hatch for mutation.

Our benchmarks demonstrate that value semantics need not sacrifice performance: persistent HAMTs are within  $2\times$  of mutable hash maps for single-threaded operations and *exceed* mutable approaches for concurrent read-heavy workloads. The per-process GC, optimised for immutable data, achieves sub-150 $\mu$ s pause times.

The implications for concurrency are profound. Values can be shared across process boundaries without synchronization, enabling zero-copy message passing, implicit snapshotting, and unambiguous serialization for distribution. The “Values Are Primary” principle is not merely an aesthetic preference—it is a design decision with cascading benefits for safety, performance, and simplicity.

Pure functions handle logic. Supervised processes handle time and failure. And values—immutable, shareable, composable values—are the foundation on which everything else is built.

## References

- [1] J. Armstrong, R. Virding, C. Claessen, and M. Wikström. *Concurrent Programming in Erlang*. Prentice Hall, 2nd edition, 1996.
- [2] J. Armstrong. *Programming Erlang: Software for a Concurrent World*. Pragmatic Bookshelf, 2007.
- [3] P. Bagwell. Ideal hash trees. Technical Report, EPFL, 2001.
- [4] E. Barendsen and S. Smetsers. Uniqueness typing for functional languages with graph rewriting semantics. *Mathematical Structures in Computer Science*, 6(6):579–612, 1996.
- [5] M. Barr and C. Wells. *Category Theory for Computing Science*. Prentice Hall, 1990.
- [6] J. Benet. IPFS—content addressed, versioned, P2P file system. arXiv preprint arXiv:1407.3561, 2014.
- [7] J.-P. Bernardy, M. Boespflug, R. Newton, S. Peyton Jones, and A. Spiwack. Linear Haskell: practical linearity in a higher-order polymorphic language. In *Proc. POPL*, 2018.
- [8] P. Chiusano and R. Björnason. Unison: a friendly programming language from the future. <https://www.unison-lang.org/>, 2015.
- [9] A. Church. An unsolvable problem of elementary number theory. *American Journal of Mathematics*, 58(2):345–363, 1936.
- [10] E. Czaplicki. Elm: concurrent FRP for functional GUIs. Senior thesis, Harvard University, 2012.
- [11] O.-J. Dahl and K. Nygaard. SIMULA—an ALGOL-based simulation language. *Communications of the ACM*, 9(9):671–678, 1966.
- [12] A. Donovan and B. Kernighan. *The Go Programming Language*. Addison-Wesley, 2015.
- [13] J. Driscoll, N. Sarnak, D. Sleator, and R. Tarjan. Making data structures persistent. *Journal of Computer and System Sciences*, 38(1):86–124, 1989.
- [14] J. Dunfield and N. Krishnaswami. Bidirectional typing. *ACM Computing Surveys*, 54(5):1–38, 2021.
- [15] C. Elliott and P. Hudak. Functional reactive animation. In *Proc. ICFP*, pages 263–273, 1997.
- [16] M. Fowler. *Patterns of Enterprise Application Architecture*. Addison-Wesley, 2003.
- [17] M. Fowler. Event sourcing. <https://martinfowler.com/eaDev/EventSourcing.html>, 2005.
- [18] J.-Y. Girard. *Interprétation fonctionnelle et élimination des coupures de l’arithmétique d’ordre supérieur*. PhD thesis, Université Paris VII, 1972.
- [19] L. Pilfold. The Gleam programming language. <https://gleam.run/>, 2024.
- [20] J. Gosling, B. Joy, G. Steele, and G. Bracha. *The Java Language Specification*. Addison-Wesley, 2nd edition, 2000.

- [21] T. Hagino. *A Categorical Programming Language*. PhD thesis, University of Edinburgh, 1987.
- [22] R. Harper. *Practical Foundations for Programming Languages*. Cambridge University Press, 2nd edition, 2016.
- [23] R. Hickey. The Clojure programming language. In *Proc. DLS*, pages 1–1, 2008.
- [24] R. Hickey. The value of values. Keynote, JaxConf, 2012.
- [25] S. L. Peyton Jones. Playing by the rules: rewriting as a practical optimisation technique in GHC. In *Proc. Haskell Workshop*, 1999.
- [26] S. L. Peyton Jones, editor. *Haskell 98 Language and Libraries: The Revised Report*. Cambridge University Press, 2003.
- [27] R. Jung, J.-H. Jourdan, R. Krebbers, and D. Dreyer. RustBelt: securing the foundations of the Rust programming language. In *Proc. POPL*, 2018.
- [28] A. Kay. The early history of Smalltalk. In *Proc. HOPL-II*, pages 69–95, 1993.
- [29] O. Kiselyov, A. Sabry, and C. Swords. Extensible effects: an alternative to monad transformers. In *Proc. Haskell Symposium*, pages 59–70, 2013.
- [30] S. Klabnik and C. Nichols. *The Rust Programming Language*. No Starch Press, 2019.
- [31] J. Lambek. A fixpoint theorem for complete categories. *Mathematische Zeitschrift*, 103:151–161, 1968.
- [32] J. Lambek and P. Scott. *Introduction to Higher-Order Categorical Logic*. Cambridge University Press, 1986.
- [33] W. Landi. Undecidability of static analysis. *ACM Letters on Programming Languages and Systems*, 1(4):323–337, 1992.
- [34] X. Leroy, D. Doligez, A. Frisch, J. Garrigue, D. Rémy, and J. Vouillon. *The OCaml System: Documentation and User’s Manual*. INRIA, 2014.
- [35] T. Lindahl and K. Sagonas. Practical type inference based on success typings. In *Proc. PPDP*, pages 167–178, 2006.
- [36] S. Mac Lane. *Categories for the Working Mathematician*. Springer, 2nd edition, 1998.
- [37] G. Malcolm. Algebraic data types and program transformation. PhD thesis, University of Groningen, 1990.
- [38] S. Marlow, editor. *Haskell 2010 Language Report*. <https://www.haskell.org/onlinereport/haskell2010/>, 2010.
- [39] N. Matsakis and F. Klock. The Rust language. In *Proc. HILT*, pages 103–104, 2014.
- [40] E. Meijer, M. Fokkinga, and R. Paterson. Functional programming with bananas, lenses, envelopes and barbed wire. In *Proc. FPCA*, pages 124–144, 1991.
- [41] R. Milner. A theory of type polymorphism in programming. *Journal of Computer and System Sciences*, 17(3):348–375, 1978.

- [42] R. Milner, M. Tofte, R. Harper, and D. MacQueen. *The Definition of Standard ML (Revised)*. MIT Press, 1997.
- [43] C. Okasaki. *Purely Functional Data Structures*. Cambridge University Press, 1999.
- [44] S. L. Peyton Jones and P. Wadler. Imperative functional programming. In *Proc. POPL*, pages 71–84, 1993.
- [45] B. C. Pierce and D. N. Turner. Local type inference. *ACM Transactions on Programming Languages and Systems*, 22(1):1–44, 2000.
- [46] B. C. Pierce. *Types and Programming Languages*. MIT Press, 2002.
- [47] G. Plotkin. LCF considered as a programming language. *Theoretical Computer Science*, 5(3):223–255, 1977.
- [48] A. Pnueli. The temporal logic of programs. In *Proc. FOCS*, pages 46–57, 1977.
- [49] D. Rémy. Type inference for records in natural extension of ML. In C. Gunter and J. Mitchell, editors, *Theoretical Aspects of Object-Oriented Programming*, pages 67–95. MIT Press, 1994.
- [50] J. C. Reynolds. Towards a theory of type structure. In *Proc. Colloque sur la Programmation*, pages 408–425, 1974.
- [51] N. Stucki, T. Rompf, V. Ureche, and P. Bagwell. RRB vector: a practical general purpose immutable sequence. In *Proc. ICFP*, pages 342–354, 2015.
- [52] D. Syme, A. Granicz, and A. Cisternino. *Expert F# 3.0*. Apress, 2010.
- [53] L. Torvalds. Git: a distributed version control system. <https://git-scm.com/>, 2005.
- [54] J. Valim. *Programming Elixir*. Pragmatic Bookshelf, 2014.
- [55] P. Wadler. Theorems for free! In *Proc. FPCA*, pages 347–359, 1989.
- [56] M. Wand. Type inference for record concatenation and multiple inheritance. *Information and Computation*, 93(1):1–15, 1991.

## A Proof of Parametricity for JAPL’s Core Calculus

We sketch the parametricity proof for  $\lambda_V$  following Reynolds [50] and Wadler [55].

**Definition A.1** (Relational Interpretation). For each type  $\tau$  and relation  $R \subseteq A \times B$  interpreting each type variable, we define a relation  $\llbracket \tau \rrbracket_R$  by induction:

$$\llbracket \mathbf{1} \rrbracket_R = \{(*, *)\} \tag{22}$$

$$\llbracket \alpha \rrbracket_R = R \tag{23}$$

$$\llbracket \tau_1 \rightarrow \tau_2 \rrbracket_R = \{(f, g) \mid \forall (a, b) \in \llbracket \tau_1 \rrbracket_R. (f(a), g(b)) \in \llbracket \tau_2 \rrbracket_R\} \tag{24}$$

$$\llbracket \tau_1 \times \tau_2 \rrbracket_R = \{((a_1, a_2), (b_1, b_2)) \mid (a_1, b_1) \in \llbracket \tau_1 \rrbracket_R \wedge (a_2, b_2) \in \llbracket \tau_2 \rrbracket_R\} \tag{25}$$

$$\llbracket \tau_1 + \tau_2 \rrbracket_R = \{(\text{inl}(a), \text{inl}(b)) \mid (a, b) \in \llbracket \tau_1 \rrbracket_R\} \tag{26}$$

$$\cup \{(\text{inr}(a), \text{inr}(b)) \mid (a, b) \in \llbracket \tau_2 \rrbracket_R\} \tag{27}$$

**Theorem A.2** (Abstraction Theorem for  $\lambda_V$ ). *For every well-typed closed term  $e : \forall\alpha. \tau(\alpha)$  and every relation  $R \subseteq A \times B$ :*

$$(\llbracket e \rrbracket_A, \llbracket e \rrbracket_B) \in \llbracket \tau \rrbracket_R$$

*Proof sketch.* By induction on the typing derivation. The key cases are:

**Abstraction:** If  $\Gamma, x : \sigma \vdash e : \tau$  and by IH for all related environments  $(\rho_1, \rho_2) \in \llbracket \Gamma \rrbracket_R$ , we have  $(\llbracket e \rrbracket^{\rho_1}, \llbracket e \rrbracket^{\rho_2}) \in \llbracket \tau \rrbracket_R$ , then for  $\lambda x. e$  we need: given  $(a, b) \in \llbracket \sigma \rrbracket_R$ , show  $(\llbracket e \rrbracket^{\rho_1[x \mapsto a]}, \llbracket e \rrbracket^{\rho_2[x \mapsto b]}) \in \llbracket \tau \rrbracket_R$ . This follows from the IH with extended environments.

**Application:** If  $(\llbracket e_1 \rrbracket^{\rho_1}, \llbracket e_1 \rrbracket^{\rho_2}) \in \llbracket \sigma \rightarrow \tau \rrbracket_R$  and  $(\llbracket e_2 \rrbracket^{\rho_1}, \llbracket e_2 \rrbracket^{\rho_2}) \in \llbracket \sigma \rrbracket_R$ , then by the definition of the relational interpretation for function types,  $(\llbracket e_1 \rrbracket^{\rho_1}(\llbracket e_2 \rrbracket^{\rho_1}), \llbracket e_1 \rrbracket^{\rho_2}(\llbracket e_2 \rrbracket^{\rho_2})) \in \llbracket \tau \rrbracket_R$ .

**Case analysis:** Follows from the disjoint union structure of the sum type relation.

The absence of mutation is critical: in a language with mutable state, the relational interpretation must account for store relations, and the proof requires step-indexed logical relations. In  $\lambda_V$ , the straightforward set-theoretic proof goes through without such complications.

*Extending this proof to the resource layer* would require replacing the set-theoretic relations above with a “state-and-store” relational model in which the resource heap is threaded through the interpretation. Linearity constraints would further require tracking resource ownership in the relation, leading to a significantly more complex proof obligation. This separation of concerns—a clean parametricity proof for values, with a richer model deferred to the resource layer—is a primary motivation for JAPL’s dual-layer architecture.  $\square$

## B JAPL Core Syntax Reference

For reference, we provide the full surface syntax of JAPL’s value layer.

```

1  -- Type declarations
2  type Option(a) =
3    | Some(a)
4    | None
5
6  type Result(a, e) =
7    | Ok(a)
8    | Err(e)
9
10 type User = {
11   id: UserId,
12   name: String,
13   email: String
14 }
15
16 -- Function declarations
17 fn user_label(user: User) → String =
18   user.name ◊ " <" ◊ user.email ◊ ">"
19
20 fn map(list: List(a), f: fn(a) → b) → List(b) =
21   match list with
22   | [] → []
23   | [x, ..rest] → [f(x), ..map(rest, f)]
24
25 -- Pattern matching
26 fn describe(result: Result(Int, String)) → String =

```

```

27  match result with
28  | Ok(n) → "Success: " ◊ Int.to_string(n)
29  | Err(msg) → "Error: " ◊ msg
30
31  -- Record update (structural sharing)
32  fn update_email(user: User, email: String) → User =
33    { user | email = email }
34
35  -- Pipeline composition
36  fn process(data: List(String)) → List(Int) =
37    data
38    ▷ List.filter(fn s → String.length(s) > 0)
39    ▷ List.map(parse_int)
40    ▷ List.filter(fn r → Result.is_ok(r))
41    ▷ List.map(Result.unwrap)
42
43  -- Row-polymorphic functions
44  fn get_name(r: { name: String | rest }) → String =
45    r.name
46
47  -- Traits
48  trait Eq(a) =
49    fn eq(x: a, y: a) → Bool
50
51  impl Eq(User) =
52    fn eq(a, b) =
53      a.id == b.id
54
55  -- Tests (values make testing trivial)
56  test "user_label formats correctly" =
57    let user = { id = UserId(1), name = "Alice",
58                email = "alice@example.com" }
59    assert user_label(user) == "Alice <alice@example.com>"
60
61  property "map preserves length" =
62    forall (xs: List(Int), f: fn(Int) → Int) →
63      List.length(map(xs, f)) == List.length(xs)

```

Listing 20: Complete value-layer syntax reference.