

# Mutation Is Local and Explicit:

Controlled Effects and Linear Ownership  
in the JAPL Programming Language

Matthew Long

*The JAPL Research Collaboration*

*YonedaAI Research Collective*

Chicago, IL

`matthew@yonedaai.com`

March 2026

## Abstract

Shared mutable state is the single most prolific source of defects in concurrent software. We present JAPL’s approach to this problem: a *dual-layer* type system that separates *pure values*—immutable, garbage-collected, and freely shared—from *linear resources*—mutable, ownership-tracked, and deterministically released. The pure layer handles the overwhelming majority of application logic with no annotation burden, while the resource layer gives systems-level control over files, sockets, buffers, and foreign memory through an ownership discipline rooted in linear type theory and Girard’s linear logic. We formalize the dual-layer system, prove type safety (progress and preservation), and show that well-typed JAPL programs satisfy *resource safety*: every acquired resource is released exactly once. A detailed comparison with Rust’s borrow checker, Haskell’s IO monad, Clean’s uniqueness types, and several other approaches demonstrates that JAPL’s design achieves a favorable trade-off between annotation burden, safety guarantees, and runtime performance.

**Keywords:** linear types, ownership, mutation, resource safety, effect systems, substructural type systems, dual-layer runtime

# Contents

<b>1</b>	<b>Introduction</b>	<b>3</b>
1.1	The Mutation Problem . . . . .	3
1.2	Contributions . . . . .	4
1.3	Paper Organization . . . . .	4
<b>2</b>	<b>Background and Related Work</b>	<b>4</b>
2.1	Substructural Type Systems . . . . .	4
2.2	Linear Logic . . . . .	5
2.3	Rust’s Ownership and Borrowing . . . . .	5
2.4	Clean’s Uniqueness Types . . . . .	5
2.5	Linear Haskell . . . . .	6
2.6	Region-Based Memory Management . . . . .	6
2.7	Capability-Based Security . . . . .	6
2.8	Austral, Roc, and Koka . . . . .	6
2.9	Algebraic Effects and Handlers . . . . .	7
<b>3</b>	<b>Formal Framework</b>	<b>7</b>
3.1	Linear Type Theory . . . . .	7
3.2	Connection to Girard’s Linear Logic . . . . .	8
3.3	Categorical Semantics . . . . .	8
3.4	Compact Closed Categories and Duality . . . . .	9
<b>4</b>	<b>JAPL’s Dual-Layer Model</b>	<b>9</b>
4.1	Design Philosophy . . . . .	9
4.2	The Pure Layer . . . . .	9
4.3	The Resource Layer . . . . .	10
4.4	The Split in Practice . . . . .	11
4.5	Formal Typing Judgment . . . . .	12
<b>5</b>	<b>Ownership Semantics</b>	<b>12</b>
5.1	Owned $\langle T \rangle$ . . . . .	12
5.2	Consumption and Return Patterns . . . . .	12
5.3	Borrowing . . . . .	13
5.4	Region-Based Inference . . . . .	14
5.5	Uniqueness Inference . . . . .	14
<b>6</b>	<b>Resource Types</b>	<b>15</b>
6.1	Declaring Resource Types . . . . .	15
6.2	Resource Lifecycle . . . . .	15
6.3	The <code>use</code> Binding . . . . .	16
6.4	Parametric Resource Types . . . . .	16
6.5	Freezing: Resource to Value Conversion . . . . .	16

<b>7</b>	<b>Interaction Between Layers</b>	<b>17</b>
7.1	Pure Functions Using Resources . . . . .	17
7.2	Ownership Transfer Between Processes . . . . .	17
7.3	The Escape Hatch . . . . .	18
7.4	Effect System Interaction . . . . .	18
<b>8</b>	<b>Comparison with Related Languages</b>	<b>19</b>
8.1	Overview . . . . .	19
8.2	Rust: Full Borrow Checker . . . . .	19
8.3	Go: GC Everything . . . . .	19
8.4	Haskell: The IO Monad . . . . .	20
8.5	Erlang: Copy Everything . . . . .	20
8.6	OCaml: Mutable Refs . . . . .	20
8.7	Clean: Uniqueness Types . . . . .	20
8.8	Summary of Trade-offs . . . . .	21
<b>9</b>	<b>Implementation</b>	<b>21</b>
9.1	Architecture Overview . . . . .	21
9.2	Type Checker Implementation . . . . .	22
9.3	Linearity Checking Algorithm . . . . .	22
9.4	Runtime Memory Model . . . . .	22
9.5	Compile-Time Ownership Verification . . . . .	23
9.6	Performance Characteristics . . . . .	23
<b>10</b>	<b>Formal Proofs</b>	<b>24</b>
10.1	Syntax of $\lambda^{\text{JR}}$ . . . . .	24
10.2	Operational Semantics . . . . .	24
10.3	Type Safety . . . . .	25
10.4	Resource Safety . . . . .	26
<b>11</b>	<b>Discussion</b>	<b>27</b>
11.1	When Ownership Pressure Leaks . . . . .	27
11.2	Mitigations . . . . .	27
11.3	The 80/20 of Resource Management . . . . .	28
11.4	Effect System Synergy . . . . .	28
11.5	Comparison with Monadic IO . . . . .	28
11.6	Limitations . . . . .	29
11.7	Future Work . . . . .	29
<b>12</b>	<b>Conclusion</b>	<b>29</b>
<b>A</b>	<b>Complete Typing Rules for <math>\lambda^{\text{JR}}</math></b>	<b>32</b>
<b>B</b>	<b>Extended Examples</b>	<b>33</b>
B.1	Database Connection Pool . . . . .	33
B.2	State Effect with Linear Resources . . . . .	34
B.3	Safe FFI Wrapping . . . . .	34

# 1 Introduction

The history of programming language design can be read, in large part, as a long negotiation with mutable state. Assembly gives the programmer direct access to every register and memory cell. C organizes that access into functions and structs but leaves aliasing entirely to the programmer’s discipline. Java introduces garbage collection and access modifiers but permits unrestricted sharing of mutable objects across threads. Haskell confines all side effects to the IO monad, achieving referential transparency at the cost of pervasive monadic plumbing. Rust enforces ownership and borrowing rules that prevent data races at compile time, but its borrow checker pervades *every* layer of the program, even pure data transformations that could never cause a data race.

JAPL takes a different position: **mutation is local and explicit**. The language is organized around two sharply delineated layers:

- (i) A **pure layer** where values are immutable, garbage-collected, and freely shared. Most application code—business logic, data transformations, protocol handling—lives here.
- (ii) A **resource layer** where mutable resources (file handles, network sockets, GPU buffers, FFI pointers) are tracked by a linear ownership discipline. Resources must be consumed exactly once: used and then released.

This separation is the second of JAPL’s seven core design principles<sup>1</sup> and is the subject of this paper.

## 1.1 The Mutation Problem

Shared mutable state creates three interrelated difficulties:

**Concurrency bugs.** When two threads can both read and write the same memory location, the result depends on scheduling—a data race. Languages without mutation restrictions must rely on locks, atomic operations, or transactional memory, each introducing its own class of bugs (deadlocks, priority inversion, live-lock).

**Aliasing and invalidation.** When multiple references point to the same mutable object, a mutation through one reference can invalidate assumptions made through another. The classic example is iterator invalidation: removing an element from a collection while iterating over it. More subtle is the *borrowed reference* problem in systems code: a pointer into a buffer becomes dangling after the buffer is reallocated.

**Reasoning difficulty.** Equational reasoning—substituting equals for equals—fails in the presence of mutation. If  $f(x)$  mutates  $x$ , then  $\text{let } y = f(x); g(y, x)$  is not equivalent to  $g(f(x), x)$ , because the second occurrence of  $x$  observes different state. This makes refactoring hazardous and formal verification prohibitively expensive.

JAPL’s dual-layer design addresses all three problems. Pure values cannot be mutated, so they are immune to data races, aliasing bugs, and reasoning failures. Resources *can* be mutated, but ownership tracking ensures that at most one reference exists at any time, restoring the ability to reason locally about mutation effects.

---

<sup>1</sup>The seven principles are: (1) Values are primary, (2) Mutation is local and explicit, (3) Concurrency is process-based, (4) Failures are normal and typed, (5) Distribution is a native concern, (6) Functions are the unit of composition, (7) Runtime simplicity matters as much as type power.

## 1.2 Contributions

This paper makes the following contributions:

1. We present a formal account of JAPL’s dual-layer type system, connecting it to Girard’s linear logic and the theory of substructural type systems (Sections 3 and 4).
2. We define the ownership semantics for the resource layer, including consumption, borrowing, region-based inference, and uniqueness inference (Section 5).
3. We formalize JAPL’s resource types and prove that well-typed programs satisfy *resource safety*: every acquired resource is eventually released exactly once (Sections 6 and 10).
4. We describe the interaction between the two layers, including the borrowing protocol and escape hatches (Section 7).
5. We provide a detailed comparison with six other language designs (Section 8) and describe an implementation strategy for the dual-layer runtime (Section 9).

## 1.3 Paper Organization

Section 2 surveys related work. Section 3 develops the formal framework: linear logic, substructural type systems, and categorical semantics. Section 4 presents JAPL’s dual-layer model. Section 5 details ownership semantics. Section 6 defines resource types. Section 7 describes layer interaction. Section 8 compares with related languages. Section 9 discusses implementation. Section 10 contains formal proofs. Section 11 reflects on practical trade-offs. Section 12 concludes.

# 2 Background and Related Work

## 2.1 Substructural Type Systems

The structural rules of a type system determine how hypotheses in a typing context may be used. Classical logic (and the simply-typed lambda calculus) admits three structural rules: *weakening* (a hypothesis may be ignored), *contraction* (a hypothesis may be duplicated), and *exchange* (hypotheses may be reordered). Dropping one or more of these rules yields a *substructural* type system [5]:

**Definition 2.1** (Substructural type systems). *Let  $\Gamma$  denote a typing context. The four principal substructural disciplines are:*

- (a) **Linear**: *No weakening, no contraction. Every variable must be used exactly once.*
- (b) **Affine**: *No contraction, but weakening is allowed. Every variable may be used at most once.*
- (c) **Relevant**: *No weakening, but contraction is allowed. Every variable must be used at least once.*
- (d) **Ordered**: *No weakening, no contraction, no exchange. Variables must be used exactly once, in order of introduction.*

JAPL’s resource layer employs *linear* types: each resource must be consumed exactly once. The pure layer uses standard structural types (with all three rules), since immutable values can be freely shared, duplicated, or ignored.

## 2.2 Linear Logic

Girard’s linear logic [1] provides the proof-theoretic foundation for linear type systems. The key connectives are:

- $A \multimap B$  (linear implication): consuming one  $A$  produces one  $B$ .
- $A \otimes B$  (tensor / multiplicative conjunction): both  $A$  and  $B$  are available simultaneously, each used exactly once.
- $!A$  (of course / exponential): an unlimited supply of  $A$ , recoverable as many times as needed.
- $A \& B$  (with / additive conjunction): a choice between  $A$  and  $B$ , where only one is consumed.
- $A \oplus B$  (plus / additive disjunction): exactly one of  $A$  or  $B$  is provided.

The exponential  $!A$  is critical: it bridges linear and unrestricted worlds. In JAPL, the pure layer corresponds to the  $!$ -fragment of linear logic: every pure value implicitly carries an exponential modality, meaning it can be freely duplicated and discarded. The resource layer corresponds to the linear fragment proper.

## 2.3 Rust’s Ownership and Borrowing

Rust [11] pioneered the use of ownership and borrowing as a practical mechanism for memory safety without garbage collection. Rust’s key rules are:

1. Each value has a single owner.
2. Ownership can be transferred (moved).
3. References can be borrowed: either one mutable reference *or* any number of immutable references, but not both simultaneously.
4. Values are dropped (destructors run) when the owner goes out of scope.

The RustBelt project [12] formalized Rust’s type system using the Iris separation logic framework, proving that well-typed Rust programs (even those using `unsafe`) satisfy memory safety when the unsafe blocks are locally sound.

JAPL differs from Rust in a fundamental architectural decision: Rust applies ownership to *all* values, requiring lifetime annotations even for purely functional data transformations. JAPL restricts ownership to the resource layer, leaving the pure layer free of ownership concerns.

## 2.4 Clean’s Uniqueness Types

Clean [13, 14] uses *uniqueness types* to ensure that a value has at most one reference. If a value is unique, it can be destructively updated in place without violating referential transparency, because no other reference can observe the mutation.

Clean’s uniqueness attribute `*` is propagated by the type checker. For example, `*File` denotes a unique file handle that can be read or written. Uniqueness types are closely related to linear types: a unique value must be used exactly once (or threaded through a computation that returns a new unique value).

JAPL’s resource layer is inspired by Clean’s approach but extends it with explicit `own` and `ref` qualifiers, region inference, and integration with an algebraic effect system.

## 2.5 Linear Haskell

Linear Haskell [10] extends Haskell with linear arrow types  $a \multimap b$ , allowing functions to specify that an argument must be consumed exactly once. The key insight is *retrofitting*: linear types are added to an existing language without breaking backward compatibility, by making linearity a property of function arrows rather than types.

JAPL takes a complementary approach. Rather than adding linearity to arrows, JAPL adds linearity to a distinguished class of *types* (resource types). This makes the boundary between linear and unrestricted code syntactically visible.

## 2.6 Region-Based Memory Management

Tofte and Talpin [15] introduced region-based memory management, where objects are allocated into lexically scoped regions and deallocated en masse when the region goes out of scope. The Cyclone language [16, 17] brought regions to a C-like setting with static safety guarantees.

Region inference automatically determines which region each allocation belongs to, minimizing annotation burden. JAPL uses region-like reasoning in its ownership inference pass: the compiler determines the lifetime of each resource and inserts release operations at the appropriate scope boundaries.

## 2.7 Capability-Based Security

The capability model of security [18, 19] organizes access control around unforgeable tokens (capabilities) rather than identity-based access control lists. A process can only perform an operation if it holds the appropriate capability.

JAPL applies this principle to resource access: a file can only be read if the function holds an `Owned<File>` or `ref File` capability. The effect system (`Io`, `Net`) further restricts which capabilities can be exercised in a given context.

## 2.8 Austral, Roc, and Koka

Several recent languages explore design points closely related to JAPL’s.

**Austral.** Austral [37] is a systems language with linear types in which every type is either *free* (unrestricted, may be copied and discarded) or *linear* (must be consumed exactly once). The dichotomy mirrors JAPL’s pure/resource split, but Austral applies it within a single layer: there is no garbage-collected pure heap, and all memory management is explicit. JAPL’s dual-layer design adds a GC-managed pure layer for ergonomics, and borrowing via `ref T` for non-consuming reads, which Austral lacks.

**Roc.** Roc is a functional language that targets application-level programming and uses a capability-based effect system in which all side effects flow through *platforms*. Like JAPL, Roc keeps the vast majority of code pure and pushes effects to the boundary. Unlike JAPL, Roc does not expose linear or ownership types to the programmer; resource management is delegated entirely to the platform layer.

**Koka.** Koka [30] pioneered row-typed algebraic effects with automatic effect inference. JAPL’s effect system draws on Koka’s design, but Koka does not provide linear types or an ownership

discipline for resources. The combination of row-typed effects *and* linear resource tracking is a distinctive feature of JAPL’s design.

## 2.9 Algebraic Effects and Handlers

Algebraic effects [20, 21] provide a structured alternative to monads for describing side effects. An effect signature declares a set of operations; an effect handler provides implementations for those operations.

JAPL’s effect system (`Io`, `State[s]`, `Fail[e]`, etc.) draws on algebraic effects. Effects form a commutative monoid under composition, and effect handlers serve as natural transformations from effectful computations to pure results [22]. The interaction between effects and linear types is a key contribution of this paper: resource operations are effectful (`with Io`), and the linear discipline ensures that effect handlers cannot accidentally duplicate or discard resources.

## 3 Formal Framework

This section develops the type-theoretic and logical foundations of JAPL’s dual-layer system.

### 3.1 Linear Type Theory

We begin with a standard presentation of the linear lambda calculus [2, 3].

**Definition 3.1** (Linear types). *The types of the linear lambda calculus are given by the grammar:*

$$A, B ::= \alpha \mid A \multimap B \mid A \otimes B \mid !A \mid A \& B \mid A \oplus B \mid \mathbf{1} \mid \mathbf{0}$$

where  $\alpha$  ranges over type variables,  $\multimap$  is linear function space,  $\otimes$  is multiplicative conjunction,  $!$  is the exponential,  $\&$  is additive conjunction,  $\oplus$  is additive disjunction,  $\mathbf{1}$  is the unit of  $\otimes$ , and  $\mathbf{0}$  is the empty type.

**Definition 3.2** (Linear typing context). *A linear context  $\Delta$  is a multiset of typing assumptions  $x : A$  where each variable  $x$  appears at most once. An unrestricted context  $\Gamma$  contains assumptions  $x : A$  where  $x$  may be used arbitrarily many times. A mixed context is a pair  $\Gamma; \Delta$ .*

The typing judgment  $\Gamma; \Delta \vdash e : A$  asserts that expression  $e$  has type  $A$  under unrestricted context  $\Gamma$  and linear context  $\Delta$ , where *every* assumption in  $\Delta$  is consumed exactly once.

**Definition 3.3** (Core typing rules). *The following rules govern the linear lambda calculus with exponentials:*

$$\begin{array}{c} \frac{}{\Gamma; x : A \vdash x : A} \text{VAR-LIN} \qquad \frac{x : A \in \Gamma}{\Gamma; \cdot \vdash x : A} \text{VAR-UN} \qquad \frac{\Gamma; \Delta, x : A \vdash e : B}{\Gamma; \Delta \vdash \lambda x. e : A \multimap B} \multimap\text{-I} \\ \\ \frac{\Gamma; \Delta_1 \vdash e_1 : A \multimap B \quad \Gamma; \Delta_2 \vdash e_2 : A}{\Gamma; \Delta_1, \Delta_2 \vdash e_1 e_2 : B} \multimap\text{-E} \qquad \frac{\Gamma; \Delta_1 \vdash e_1 : A \quad \Gamma; \Delta_2 \vdash e_2 : B}{\Gamma; \Delta_1, \Delta_2 \vdash (e_1, e_2) : A \otimes B} \otimes\text{-I} \\ \\ \frac{\Gamma; \Delta_1 \vdash e : A \otimes B \quad \Gamma; \Delta_2, x : A, y : B \vdash e' : C}{\Gamma; \Delta_1, \Delta_2 \vdash \mathbf{let} (x, y) = e \mathbf{in} e' : C} \otimes\text{-E} \qquad \frac{\Gamma; \cdot \vdash e : A}{\Gamma; \cdot \vdash !e : !A} !\text{-I} \\ \\ \frac{\Gamma; \Delta_1 \vdash e_1 : !A \quad \Gamma, x : A; \Delta_2 \vdash e_2 : B}{\Gamma; \Delta_1, \Delta_2 \vdash \mathbf{let} !x = e_1 \mathbf{in} e_2 : B} !\text{-E} \end{array}$$

The !-I rule requires that the expression  $e$  use *no* linear resources: only unrestricted values can be promoted to the exponential. This is the formal counterpart of JAPL’s rule that pure values cannot capture resources.

### 3.2 Connection to Girard’s Linear Logic

The Curry-Howard correspondence extends to linear logic: proofs correspond to programs, propositions to types, and cut-elimination to evaluation [1, 4].

**Proposition 3.4** (Curry-Howard for linear logic). *Under the linear Curry-Howard correspondence:*

- (i)  $A \multimap B$  corresponds to a function consuming one  $A$  to produce one  $B$ .
- (ii)  $A \otimes B$  corresponds to a pair where both components must be used.
- (iii)  $!A$  corresponds to a value that may be duplicated or discarded.
- (iv)  $A \& B$  corresponds to a lazy pair (choose one projection).
- (v)  $A \oplus B$  corresponds to a tagged union (one alternative is provided).

In JAPL’s dual-layer system, the Curry-Howard correspondence manifests directly: pure-layer types correspond to the !-fragment (all types implicitly carry !), while resource-layer types correspond to the linear fragment (no implicit !).

### 3.3 Categorical Semantics

The categorical semantics of linear types provides a rigorous foundation for the dual-layer design.

**Definition 3.5** (Symmetric monoidal closed category). *A symmetric monoidal closed category (SMCC)  $(\mathcal{C}, \otimes, I, \multimap)$  consists of a category  $\mathcal{C}$  equipped with a bifunctor  $\otimes : \mathcal{C} \times \mathcal{C} \rightarrow \mathcal{C}$ , a unit object  $I$ , and an internal hom  $\multimap$  such that  $- \otimes A \dashv A \multimap -$  for every object  $A$ .*

**Definition 3.6** (Linear-non-linear model). *A linear-non-linear (LNL) model [8] consists of:*

- (i) *A Cartesian closed category  $\mathcal{C}$  (the “non-linear” or “intuitionistic” part).*
- (ii) *A symmetric monoidal closed category  $\mathcal{L}$  (the “linear” part).*
- (iii) *A symmetric monoidal adjunction  $F \dashv G : \mathcal{L} \rightarrow \mathcal{C}$  where  $F : \mathcal{C} \rightarrow \mathcal{L}$  is a symmetric monoidal functor.*

*The exponential ! is recovered as the comonad  $FG$  on  $\mathcal{L}$ .*

**Theorem 3.7** (Benton’s LNL correspondence [8]). *The category of LNL models is equivalent to the category of models of intuitionistic linear logic (ILL) with an exponential ! satisfying the standard structural rules (weakening, contraction, dereliction, promotion).*

This result is directly relevant to JAPL:

**Remark 3.8** (JAPL as an LNL model). *JAPL’s dual-layer system instantiates the LNL framework:*

- $\mathcal{C}$  is the category of pure JAPL types, which is Cartesian closed (it has product types, function types, and a terminal object *Unit*).

- $\mathcal{L}$  is the category of resource JAPL types, which is symmetric monoidal closed (tensor is pair-of-resources, linear function is ownership-transferring function).
- The functor  $F : \mathcal{C} \rightarrow \mathcal{L}$  embeds pure values into the resource layer (a pure value can always be used where a resource is expected, since it carries an implicit!).
- The functor  $G : \mathcal{L} \rightarrow \mathcal{C}$  “forgets” linearity—this corresponds to the **freeze** or **snapshot** operations that convert a resource to an immutable value.

### 3.4 Compact Closed Categories and Duality

In the full linear logic setting, the multiplicative fragment admits a notion of *duality*: every type  $A$  has a dual  $A^\perp$ , and the category of types forms a \*-autonomous category [9].

**Definition 3.9** (\*-Autonomous category). *A \*-autonomous category is a symmetric monoidal closed category  $(\mathcal{C}, \otimes, I, \multimap)$  equipped with a dualizing object  $\perp$  such that the canonical morphism  $A \rightarrow (A \multimap \perp) \multimap \perp$  is an isomorphism for every object  $A$ .*

While JAPL does not expose full classical linear logic to the programmer, the duality principle manifests in the resource layer’s *consumption pattern*: acquiring a resource of type  $R$  creates an obligation (dual) to release it. The linear type system ensures that this obligation is always discharged.

## 4 JAPL’s Dual-Layer Model

This section presents the core design of JAPL’s two-layer type and memory system.

### 4.1 Design Philosophy

The dual-layer model is motivated by an empirical observation: **most code does not need mutation**. Business logic, data transformations, protocol parsing, serialization, routing—all of these are naturally expressed as pure functions over immutable values. Mutation is needed only at the boundaries: reading files, writing sockets, managing GPU memory, interacting with foreign libraries.

By restricting ownership tracking to the resource layer, JAPL achieves two goals simultaneously:

1. **Ergonomics**: The vast majority of code is free of ownership annotations, lifetime parameters, and borrow-checker errors.
2. **Safety**: The code that *does* interact with mutable resources is subject to a rigorous linear discipline that prevents resource leaks, use-after-free, and double-free.

### 4.2 The Pure Layer

The pure layer encompasses all immutable values: algebraic data types, records, tuples, lists, maps, strings, closures, and numeric types.

**Definition 4.1** (Pure types). *The pure types of JAPL are given by:*

$$\tau ::= \alpha \mid \mathit{Int} \mid \mathit{Float} \mid \mathit{Bool} \mid \mathit{String} \mid \mathit{Unit} \mid \tau_1 \times \tau_2 \mid \tau_1 + \tau_2 \mid \tau_1 \rightarrow \tau_2 \mid \mathit{List}[\tau] \mid \{\ell_1 : \tau_1, \dots, \ell_n : \tau_n \mid \rho\}$$

where  $\alpha$  ranges over type variables and  $\rho$  ranges over row variables.

Pure values have the following properties:

- **Immutable:** Once created, a pure value never changes.
- **GC-managed:** Memory is reclaimed by a generational, per-process garbage collector.
- **Freely shareable:** Pure values can be passed to any number of functions, stored in any number of data structures, and sent to any number of processes.
- **Structurally typed:** Two values with the same structure have the same type (row polymorphism for records).

```
1 -- Pure values: immutable, GC-managed, freely shareable
2 let data = [1, 2, 3, 4, 5]
3 let copy = data -- sharing is fine; data is immutable
4
5 -- Pure function: transforms data, no side effects
6 fn transform(data: List[Int]) →List[Int] =
7   List.map(data, fn x →x * 2)
8
9 -- Composition of pure functions
10 fn process_order(order: Order) →Invoice =
11   order
12   ▷ validate
13   ▷ calculate_totals
14   ▷ apply_discounts
15   ▷ generate_invoice
```

Listing 1: Pure-layer code: no ownership concerns

### 4.3 The Resource Layer

The resource layer encompasses all mutable, owned values: file handles, network sockets, database connections, GPU buffers, FFI pointers, and mutable arrays.

**Definition 4.2** (Resource types). *A resource type is introduced by the `resource` keyword:*

$$R ::= \text{resource } T$$

*Resource types are linear: each value of a resource type must be consumed exactly once in every execution path.*

**Definition 4.3** (Ownership qualifiers). *JAPL provides two ownership qualifiers:*

- own  $T$ : Exclusive ownership. The holder can read, write, and release the resource. Corresponds to the linear type  $T$ .*
- ref  $T$ : Borrowed reference. The holder can read the resource but cannot write, release, or transfer it. Corresponds to  $!T$  restricted to a scope.*

```

1 resource File
2 resource Socket
3
4 -- Opening a file acquires ownership (manual management via let)
5 fn read_config(path: Path) →Result[Config, ReadError] with Io =
6   let file = File.open(path, Read)?
7   let contents = File.read_all(file)?
8   File.close(file) -- ownership consumed; compile error if omitted
9   parse_config(contents)
10
11 -- Ownership transfer: send consumes the socket, returns it
12 fn send(socket: Owned[Socket], bytes: Bytes)
13   → (Owned[Socket], Int) with Io =
14   let n = Socket.write(socket, bytes)
15   (socket, n)

```

Listing 2: Resource-layer code: explicit ownership

## 4.4 The Split in Practice

The boundary between pure and resource layers is clean and syntactically visible. Consider a web server:

```

1 -- Pure layer: all business logic
2 fn route(req: Request) →Response =
3   match req.method, req.path with
4   | Get, "/users/" ++id →get_user(id)
5   | Post, "/users"      → create_user(req.body)
6   | _                   → not_found()
7
8 fn get_user(id: String) →Response =
9   match UserRepo.find(id) with
10  | Some(user) →Response.json(200, User.to_json(user))
11  | None       → Response.json(404, error_json("not found"))
12
13 -- Resource layer: IO boundary
14 fn serve(listener: Owned[TcpListener])
15   → Never with Io, Process[ServerMsg] =
16   let (listener, conn) = Tcp.accept(listener)
17   let _ = Process.spawn(fn →handle_conn(conn))
18   serve(listener)
19
20 fn handle_conn(conn: Owned[TcpSocket]) →Unit with Io =
21   let (conn, req_bytes) = Tcp.read(conn)
22   let req = parse_request(req_bytes) -- pure
23   let resp = route(req)             -- pure
24   let conn = Tcp.write(conn, serialize_response(resp))
25   Tcp.close(conn)

```

Listing 3: The dual-layer split in a web server

In this example, the pure functions `route`, `get_user`, `parse_request`, and `serialize_response`

make up the majority of the code and require no ownership annotations. Only `serve` and `handle_conn`, which interact with TCP sockets, operate in the resource layer.

## 4.5 Formal Typing Judgment

The dual-layer system uses a combined typing judgment:

**Definition 4.4** (Dual-layer typing judgment). *The judgment*

$$\Gamma; \Delta \vdash_{\varepsilon} e : \tau$$

*asserts that expression  $e$  has type  $\tau$  under unrestricted (pure) context  $\Gamma$ , linear (resource) context  $\Delta$ , and with effect  $\varepsilon \in \{\text{Pure}, \text{lo}, \dots\}$ .*

*The rules ensure:*

- (i) *Variables in  $\Gamma$  may be used any number of times.*
- (ii) *Variables in  $\Delta$  must each be used exactly once.*
- (iii) *Resource-producing operations (`File.open`, etc.) add bindings to  $\Delta$ .*
- (iv) *Resource-consuming operations (`File.close`, etc.) remove bindings from  $\Delta$ .*
- (v) *At every join point (if-then-else, match), the remaining  $\Delta$  must be identical in all branches.*

## 5 Ownership Semantics

### 5.1 Owned $\langle T \rangle$

The type `Owned $\langle T \rangle$`  represents exclusive ownership of a resource of type  $T$ . The owner has three rights:

1. **Use:** Perform operations on the resource (read, write).
2. **Transfer:** Pass ownership to another function or process.
3. **Release:** Deallocate the resource, releasing its underlying system resource.

These rights are mutually exclusive over time: once ownership is transferred or the resource is released, the original holder loses all rights.

**Definition 5.1** (Ownership transfer). *A function that accepts an `Owned $\langle T \rangle$`  parameter consumes the caller's ownership. If the function returns an `Owned $\langle T \rangle$` , ownership is transferred back to the caller. The typing rule is:*

$$\frac{\Gamma; \Delta, x : \text{Owned}\langle T \rangle \vdash_{\varepsilon} e : \tau}{\Gamma; \Delta \vdash_{\varepsilon} \lambda(x : \text{Owned}\langle T \rangle). e : \text{Owned}\langle T \rangle \rightarrow_{\varepsilon} \tau} \text{OWN-TRANSFER}$$

### 5.2 Consumption and Return Patterns

JAPL supports several patterns for resource consumption:

**Terminal consumption.** The resource is released and not returned:

```

1 fn close_file(file: Owned[File]) →Unit with Io =
2   File.close(file)
3   -- file is consumed; cannot be used after this point

```

Listing 4: Terminal consumption: File.close

**Threaded consumption.** The resource is used and then returned with updated state:

```

1 fn read_line(file: Owned[File])
2   → (Owned[File], String) with Io =
3   let (file, line) = File.read_line(file)
4   (file, line)

```

Listing 5: Threaded consumption: read then return

**Fork-join consumption.** Ownership is temporarily split (borrowing) and then rejoined:

```

1 fn copy_file(src: Owned[File], dst: Owned[File])
2   → (Owned[File], Owned[File]) with Io =
3   let bytes = File.read_all(src)
4   let dst = File.write_all(dst, bytes)
5   (src, dst)

```

Listing 6: Fork-join via borrowing

### 5.3 Borrowing

Borrowing allows temporary read-only access to a resource without transferring ownership.

**Definition 5.2** (Borrow). *A borrow of type  $\text{ref } T$  grants read-only access to a resource of type  $T$ . The borrow is valid only within the lexical scope of the borrowing expression. The typing rule is:*

$$\frac{\Gamma; \Delta, x : \text{Owned}\langle T \rangle \vdash_{\varepsilon} e : \tau \quad x \notin \text{FV}(e) \text{ as owned}}{\Gamma; \Delta, x : \text{Owned}\langle T \rangle \vdash_{\varepsilon} \mathbf{borrow } x \text{ as } y \mathbf{ in } [y/x]e : \tau} \text{BORROW}$$

where  $y : \text{ref } T$  is the borrowed reference, and  $x$  remains available as  $\text{Owned}\langle T \rangle$  after the borrow scope exits.

**Remark 5.3** (Borrow desugaring). *The surface syntax  $\mathbf{borrow } x \text{ as } y \mathbf{ in } e$  desugars to the core calculus form  $\mathbf{borrow}(\text{res}(a), \lambda y. e)$ , where  $x$  is bound to  $\text{res}(a)$  and the body  $e$  is wrapped in a lambda that receives the borrowed value  $y$ . The core reduction rule (Definition 10.4, rule BORROW) then reduces  $\mathbf{borrow}(\text{res}(a), f)$  to  $(f v, \text{res}(a))$ , applying the closure to the live value  $v$  while preserving the resource handle.*

```

1 fn peek(buf: ref Buffer) →Byte =
2   Buffer.get(buf, 0)
3
4 fn process(buf: Owned[Buffer]) →(Owned[Buffer], Byte) with Io =

```

```

5 let byte = peek(ref buf) -- borrow for read-only access
6 (buf, byte)             -- buf is still owned here

```

Listing 7: Borrowing: peek at a buffer without consuming it

The key constraint is that borrows cannot *outlive* the owner: the borrowed reference `ref T` is valid only within the scope where it is created. This is enforced statically by the compiler.

## 5.4 Region-Based Inference

JAPL uses region-based inference [15] to automatically determine resource lifetimes and insert release operations.

**Definition 5.4** (Region). *A region  $r$  is an abstract lifetime associated with a lexical scope. Every resource allocation is assigned a region, and the resource is released when the region ends. Regions form a partial order: if scope  $r_1$  is nested inside scope  $r_2$ , then  $r_1 \sqsubseteq r_2$  (i.e.,  $r_1$  ends before  $r_2$ ).*

The compiler’s region inference pass operates in three phases:

1. **Constraint generation:** For each resource allocation, generate a region variable. For each use, generate a constraint that the resource’s region must contain the use site. For each ownership transfer, propagate the region.
2. **Constraint solving:** Solve the region constraints to find the minimal region for each resource.
3. **Insertion:** Insert release operations at region boundaries where resources go out of scope.

**Example 5.5** (Automatic release). *In the following code, the compiler infers that `conn` must live at least as long as the scope of `process_query`:*

```

1 fn process_query(url: String, sql: String)
2   → Result[Rows, DbError] with Io =
3   use conn = Db.connect(url)?
4   let result = Db.query(ref conn, sql)?
5   -- compiler inserts: Db.close(conn) here
6   Ok(result)

```

Listing 8: Region inference inserts `Db.close` automatically

*The `use` binding signals that `conn` is a linear resource; the compiler ensures it is released at the end of its region.*

## 5.5 Uniqueness Inference

Beyond region inference, JAPL performs *uniqueness inference* to determine when a value has a single reference, enabling in-place mutation optimizations.

**Definition 5.6** (Uniqueness). *A value is unique if it has exactly one reference in the program state. A unique immutable value can be safely updated in-place (destructive update), because no other reference can observe the change.*

**Proposition 5.7** (Uniqueness optimization). *If a pure value  $v$  is unique at a record-update site  $\{v \mid \ell = e\}$ , the compiler may perform the update in-place, avoiding allocation. This is semantically equivalent to creating a fresh copy with the field changed.*

*Proof.* Since  $v$  has a single reference, no other expression can observe the difference between in-place update and copy-on-write. By parametricity, the program cannot distinguish the two implementations.  $\square$

This optimization is particularly important for the common pattern of threading state through recursive functions:

```

1 fn server_loop(state: ServerState) ->Never with Process[Msg] =
2   let msg = Process.receive()
3   -- If state is unique (no other references), the update
4   -- { state | count = state.count + 1 } is performed in-place.
5   let new_state = { state | count = state.count + 1 }
6   server_loop(new_state)

```

Listing 9: Uniqueness inference enables in-place update

## 6 Resource Types

### 6.1 Declaring Resource Types

Resource types are declared with the `resource` keyword:

```

1 resource File
2 resource Socket
3 resource DbConnection
4 resource GpuBuffer
5 resource MutableArray[a]

```

Listing 10: Resource type declarations

The `resource` keyword instructs the compiler to apply linear typing rules to values of this type. A resource cannot be silently dropped, duplicated, or aliased.

### 6.2 Resource Lifecycle

Every resource follows a lifecycle:

1. **Acquisition:** A resource is created via a constructor (e.g., `File.open`, `Tcp.connect`). The constructor returns `Owned<R>`.
2. **Use:** The resource is operated on via functions that take `Owned<R>` or `ref R`. Functions that take `Owned<R>` must return it (threaded pattern) or consume it (terminal pattern).
3. **Release:** The resource is released via a destructor (e.g., `File.close`, `Tcp.disconnect`). The destructor consumes `Owned<R>` and returns no resource.

**Definition 6.1** (Resource safety). *A program is resource-safe if, for every resource acquired during execution, the resource is eventually released exactly once.*

**Theorem 6.2** (Well-typed programs are resource-safe). *If  $\Gamma; \cdot \vdash_{\epsilon} e : \tau$  (i.e., the program type-checks with an empty linear context), then execution of  $e$  is resource-safe.*

We defer the proof to Section 10.

### 6.3 The use Binding

The use keyword introduces a linear binding with automatic release semantics:

```
1 fn read_entire_file(path: String) →Result[String, IoError] with Io =
2   use file = File.open(path, Read)?
3   let contents = File.read_all(file)?
4   -- File.close(file) is automatically inserted by the compiler
5   -- at the end of the 'use' scope
6   Ok(contents)
```

Listing 11: The use binding

The use binding desugars to a try-finally pattern: even if the body raises an error (via the ? operator or a crash), the resource is released.

**Definition 6.3** (use desugaring). *The binding `use x = e1 in e2` desugars to:*

$$\text{let } x = e_1 \text{ in try } e_2 \text{ finally } \text{release}(x)$$

where *release* is the destructor associated with the resource type of *x*.

### 6.4 Parametric Resource Types

Resources can be parametric:

```
1 resource MutableArray[a]
2
3 fn sort_in_place(arr: own MutableArray[Int])
4   → own MutableArray[Int] =
5   let len = MutableArray.length(arr)
6   loop i = 0, arr = arr while i < len do
7     loop j = i + 1, arr = arr while j < len do
8       if MutableArray.get(arr, j) < MutableArray.get(arr, i) then
9         let arr = MutableArray.swap(arr, i, j)
10        continue(j + 1, arr)
11      else
12        continue(j + 1, arr)
13    continue(i + 1, arr)
14  arr
```

Listing 12: Parametric resource types

The linearity constraint applies to the outer resource container, not to its type parameter. The elements of a `MutableArray[Int]` are pure `Int` values and can be freely copied.

### 6.5 Freezing: Resource to Value Conversion

A resource can be *frozen* to produce an immutable value, transferring data from the resource layer to the pure layer:

```
1 fn build_message() →Bytes with Io =
2   let buf = Buffer.alloc(1024)      -- resource layer
3   let buf = Buffer.write(buf, 0, header)
```

```

4 let buf = Buffer.write(buf, 64, payload)
5 Buffer.freeze(buf)           -- pure layer: immutable Bytes
6 -- buf is consumed; only the frozen Bytes value remains

```

Listing 13: Freezing a buffer into immutable bytes

**Definition 6.4** (Freeze). A freeze operation has type  $\text{Owned}\langle R \rangle \multimap \tau$ , consuming the resource and producing a pure value. After freezing, the resource is released and the returned value lives in the pure layer (managed by GC).

Freezing corresponds to the functor  $G : \mathcal{L} \rightarrow \mathcal{C}$  in the LNL model (Theorem 3.6): it maps a linear resource to an unrestricted pure value.

## 7 Interaction Between Layers

### 7.1 Pure Functions Using Resources

Pure functions cannot *own* resources, but they can *borrow* them. A function that borrows a resource via  $\text{ref } T$  does not participate in the ownership protocol; it simply reads the resource’s current state.

```

1 -- This function is pure except for the borrow
2 fn buffer_checksum(buf: ref Buffer) →Int =
3   let bytes = Buffer.to_bytes_view(buf) -- read-only view
4   Bytes.fold(bytes, 0, fn acc, b →acc + b)
5
6 fn validate_and_send(buf: Owned[Buffer], sock: Owned[Socket])
7   → (Owned[Buffer], Owned[Socket]) with Io =
8   let sum = buffer_checksum(ref buf) -- borrow for pure computation
9   if sum > 0 then
10    let (sock, _) = Socket.send(sock, Buffer.to_bytes(ref buf))
11    (buf, sock)
12  else
13    (buf, sock)

```

Listing 14: Pure function borrowing a resource

### 7.2 Ownership Transfer Between Processes

When a resource is sent to another process, ownership transfers:

```

1 fn producer() →Unit with Io, Process[ProducerMsg] =
2   let buf = Buffer.alloc(4096)
3   let buf = Buffer.write(buf, 0, generate_data())
4   -- Ownership transfers to the consumer process
5   Process.send(consumer_pid, ProcessBuffer(buf))
6   -- buf is consumed; using it here is a compile error
7
8 fn consumer() →Never with Io, Process[ConsumerMsg] =
9   match Process.receive() with
10  | ProcessBuffer(buf) →
11    -- We now own buf
12    let data = Buffer.freeze(buf)

```

```

13 process_data(data)
14 consumer()

```

Listing 15: Ownership transfer between processes

### 7.3 The Escape Hatch

For interoperability with C libraries and low-level systems code, JAPL provides an `unsafe` escape hatch that allows bypassing linearity checks:

```

1 -- unsafe block: linearity is not checked
2 fn raw_mmap(fd: Owned[File], size: Int) →Owned[MappedRegion] with Io =
3   unsafe {
4     let ptr = foreign_mmap(File.raw_fd(ref fd), size)
5     MappedRegion.from_raw(ptr, size)
6   }
7 -- Outside unsafe: normal linearity rules apply

```

Listing 16: The unsafe escape hatch

The `unsafe` keyword has several properties:

1. It must appear explicitly in the source code, making it searchable and auditable.
2. It does not propagate: a function containing `unsafe` can still present a safe interface to its callers.
3. It is tracked by the effect system: `unsafe` blocks require the `Unsafe` effect, which cannot be silently introduced.

### 7.4 Effect System Interaction

Resource operations are effectful: they carry the `Io` effect. The effect system and the ownership system cooperate:

**Definition 7.1** (Effect-ownership interaction). *The following invariants hold:*

- (i) *Resource acquisition requires an effect (typically `Io`).*
- (ii) *Resource release requires the same effect as acquisition.*
- (iii) *Pure functions ( $\varepsilon = \text{Pure}$ ) cannot acquire or release resources; they can only borrow.*
- (iv) *Effect handlers cannot duplicate or discard resources: the handler's continuation must preserve the linear context.*

**Proposition 7.2** (Effect handlers preserve linearity). *If an effect handler  $h$  transforms computation  $\Gamma; \Delta \vdash_{\varepsilon} e : \tau$  into  $\Gamma; \Delta' \vdash_{\varepsilon'} e' : \tau'$ , then  $\Delta' = \Delta$  (the linear context is unchanged).*

*Proof.* By induction on the structure of the effect handler. The handler provides implementations for effect operations, but each operation implementation must preserve the linear context: it receives a continuation that expects the same linear resources, and it must pass exactly those resources to the continuation. Since the handler cannot forge or destroy linear resources (this is enforced by the `!-I` rule, which prevents linear resources from appearing in unrestricted positions), the linear context is preserved.  $\square$

## 8 Comparison with Related Languages

### 8.1 Overview

Table 1: Comparison of mutation control strategies across languages

Language	Mutation model	Resource mgmt	Annotation	GC
JAPL	Dual-layer	Linear ownership	Low (pure layer)	Hybrid
Rust	Ownership+borrow	Ownership+RAII	High (everywhere)	None
Go	Unrestricted	GC + <code>defer</code>	None	Full GC
Haskell	IO monad	GC + finalizers	Medium (monadic)	Full GC
Erlang	Copy everything	Per-process GC	None	Per-proc GC
OCaml	Mutable refs	GC	Low	Full GC
Clean	Uniqueness types	Uniqueness	Medium	GC

### 8.2 Rust: Full Borrow Checker

Rust applies its ownership and borrowing discipline to *every* value, including purely functional data. This provides strong guarantees—no GC is needed, and memory safety is ensured at compile time—but imposes a significant annotation burden.

Consider a simple data transformation in Rust:

```
fn transform(data: Vec<i32>) -> Vec<i32> {
    data.into_iter().map(|x| x * 2).collect()
}
```

The programmer must choose between `Vec<i32>` (ownership transfer), `&Vec<i32>` (immutable borrow), and `&mut Vec<i32>` (mutable borrow). For a pure transformation, this choice is an unnecessary cognitive burden.

In JAPL, the same function requires no ownership annotations:

```
1 fn transform(data: List[Int]) ->List[Int] =
2   List.map(data, fn x ->x * 2)
```

Listing 17: JAPL vs. Rust: pure data transformation

The trade-off is that JAPL’s pure layer requires a garbage collector, while Rust avoids GC entirely. We argue that this trade-off is worthwhile for application-level code, where GC pauses are negligible relative to I/O latency.

### 8.3 Go: GC Everything

Go takes the opposite extreme: all memory is garbage-collected, and there is no ownership system. Resources are managed by convention (`defer` and the `io.Closer` interface) rather than by the type system.

The consequence is that resource leaks are possible:

```
func processFile(path string) error {
    f, err := os.Open(path)
    if err != nil { return err }
```

```

    // Easy to forget: defer f.Close()
    data, err := io.ReadAll(f)
    // If we return here, f is leaked
    ...
}

```

JAPL’s linear types make this class of bug impossible: the compiler rejects any program that fails to consume a resource.

## 8.4 Haskell: The IO Monad

Haskell confines side effects to the IO monad, ensuring that pure functions are referentially transparent. However, the monadic style has several drawbacks:

1. Monadic binding (`>=>`) is syntactically heavier than direct-style code.
2. Composing different effects (state, error, IO) requires monad transformers, which are notoriously difficult to use.
3. Resources within IO are not linearly tracked; finalizers and `bracket` provide safety but not compile-time guarantees.

JAPL’s effect system provides effect tracking without monadic syntax: effects are listed after `with` in function signatures, and composition is automatic (effects form a commutative monoid). The resource layer adds linear tracking that Haskell’s IO lacks.

## 8.5 Erlang: Copy Everything

Erlang achieves concurrency safety by making all values immutable and copying data between processes. This eliminates shared mutable state entirely but has performance implications for large data transfers.

JAPL inherits Erlang’s process isolation model for pure values but extends it with ownership transfer for resources: a resource can be moved between processes without copying, while maintaining the invariant that only one process owns it at a time.

## 8.6 OCaml: Mutable Refs

OCaml allows mutation through mutable record fields and `ref` cells. The type system does not distinguish pure and impure code; any function may have side effects.

This pragmatic approach simplifies the type system but sacrifices the ability to reason about purity. JAPL’s effect annotations (`with Io`, `with State[s]`) restore this reasoning while keeping the annotation burden low (effects are inferred within function bodies).

## 8.7 Clean: Uniqueness Types

Clean’s uniqueness types [13] are the closest relative of JAPL’s resource layer. Clean uses a `*` attribute to mark unique values; the compiler ensures that unique values are not aliased.

JAPL extends Clean’s approach in several ways:

1. **Explicit layer separation:** JAPL does not apply uniqueness to all types; only declared `resource` types are subject to linear discipline.

2. **Borrowing:** Clean does not have a borrowing mechanism; in Clean, any use of a unique value consumes it. JAPL’s `ref` qualifier allows non-consuming reads.
3. **Effect integration:** JAPL’s resource operations are annotated with effects, providing an additional layer of safety.
4. **Process integration:** JAPL supports ownership transfer between processes, which Clean (not being an actor-based language) does not address.

## 8.8 Summary of Trade-offs

Table 2: Trade-off analysis: safety guarantees vs. programmer burden

	<i>Data race free</i>	<i>Resource safe</i>	<i>No GC needed</i>	<i>Low annotation</i>	<i>Effect tracking</i>
JAPL	✓	✓		✓	✓
Rust	✓	✓	✓		
Go				✓	
Haskell	✓				✓
Erlang	✓			✓	
OCaml				✓	
Clean	✓	✓			

JAPL is the only language in this comparison that simultaneously provides data-race freedom, compile-time resource safety, low annotation burden (for the majority of code), and effect tracking. The cost is a garbage collector for the pure layer; we argue in Section 11 that this is an acceptable trade-off.

## 9 Implementation

### 9.1 Architecture Overview

The JAPL compiler and runtime implement the dual-layer model as follows:

1. **Parser:** Produces an AST with explicit `use`, `own`, and `ref` annotations.
2. **Type checker:** Performs bidirectional type checking with local inference. Maintains separate unrestricted ( $\Gamma$ ) and linear ( $\Delta$ ) contexts.
3. **Linearity checker:** Verifies that all linear variables are consumed exactly once on every execution path.
4. **Region inference:** Assigns regions to resource allocations and inserts release operations.
5. **Uniqueness analysis:** Identifies pure values with a single reference for in-place update optimization.

6. **Code generation:** Compiles to native code via LLVM or Cranelift.
7. **Runtime:** Manages the dual-layer memory model at execution time.

## 9.2 Type Checker Implementation

The type checker uses a *splitting* judgment to partition the linear context among subexpressions [5]:

**Definition 9.1** (Context splitting). *A split of linear context  $\Delta$  into  $\Delta_1$  and  $\Delta_2$  (written  $\Delta = \Delta_1 \uplus \Delta_2$ ) assigns each variable in  $\Delta$  to exactly one of  $\Delta_1$  or  $\Delta_2$ .*

For function application  $e_1 e_2$ , the type checker:

1. Generates a fresh split of the current linear context.
2. Checks  $e_1$  against  $\Delta_1$ .
3. Checks  $e_2$  against  $\Delta_2$ .
4. Verifies that  $\Delta_1$  and  $\Delta_2$  are both fully consumed.

For conditionals and pattern matching, the type checker verifies that all branches consume the *same* set of linear variables:

$$\frac{\Gamma; \Delta_0 \vdash e : \tau \quad \forall i. \Gamma; \Delta_i \vdash p_i \Rightarrow e_i : \sigma \quad \Delta_1 = \Delta_2 = \dots = \Delta_n}{\Gamma; \Delta_0, \Delta_1 \vdash \mathbf{match} \ e \ \mathbf{with} \ \{p_i \Rightarrow e_i\} : \sigma} \text{MATCH}$$

## 9.3 Linearity Checking Algorithm

The linearity checker runs as a separate pass after type checking. It traverses the typed AST and maintains a set of “live” linear variables. The algorithm is:

1. At a `let` binding of a resource type: add the variable to the live set.
2. At a use of a linear variable: remove it from the live set. If already removed, report a “use after move” error.
3. At a branch point (if/match): check each branch independently; verify that all branches have the same live set at exit.
4. At function exit: verify that the live set matches the function’s return type (i.e., any linear variable still live must be returned).
5. At scope exit for `use` bindings: if the variable is still live, insert a release call.

## 9.4 Runtime Memory Model

The runtime maintains two memory regions per process:

**Immutable heap (GC-managed).** All pure values reside here. The garbage collector is generational and per-process:

- **Nursery:** Small, per-process bump allocator. Most values die young and are collected cheaply.
- **Old generation:** Surviving values are promoted to a shared old generation, collected concurrently.
- **No write barriers:** Since data is immutable, there are no pointer mutations to track. This significantly simplifies the GC.
- **Process death:** When a process terminates, its entire nursery is reclaimed instantly.

**Resource arena (ownership-managed).** All linear resources reside here. There is no garbage collection; resources are freed deterministically when consumed:

- **Allocation:** Resources are allocated in a per-process arena.
- **Deallocation:** When a resource is consumed (released), its memory is immediately freed.
- **Transfer:** When ownership is transferred to another process, the resource is moved from one process's arena to another's.
- **No fragmentation:** Arena allocation avoids the fragmentation problems of general-purpose allocators.

## 9.5 Compile-Time Ownership Verification

The compiler performs ownership verification as a fixed-point computation over the control-flow graph (CFG):

1. Build the CFG for each function.
2. At each CFG node, compute the set of live linear variables.
3. At each join point, check that the live sets from all predecessors are identical.
4. At the function exit node, check that the live set matches the expected output (linear variables that appear in the return type).
5. If any check fails, emit a diagnostic with the offending variable and the point where it was consumed or escaped.

This analysis is similar to liveness analysis in traditional compilers but with the additional constraint that each linear variable must be consumed *exactly once*, not “at least once” or “at most once.”

## 9.6 Performance Characteristics

The dual-layer design has favorable performance characteristics:

- **Pure-layer overhead:** GC pauses are bounded by the nursery size and are per-process (no global stop-the-world). Uniqueness inference eliminates unnecessary copies.
- **Resource-layer overhead:** Zero runtime overhead beyond the actual resource operations; ownership is verified entirely at compile time.
- **Combined:** The compiler can optimize across layers. For example, a buffer that is allocated, filled, and frozen within a single function can be stack-allocated if uniqueness analysis proves it does not escape.

## 10 Formal Proofs

This section proves type safety and resource safety for JAPL’s dual-layer system. We work with a core calculus  $\lambda^{\text{JR}}$  (JAPL Resources) that captures the essential features.

### 10.1 Syntax of $\lambda^{\text{JR}}$

**Definition 10.1** (Core calculus  $\lambda^{\text{JR}}$ ).

$$\begin{array}{ll}
 \text{Types} & \tau ::= B \mid \tau_1 \rightarrow \tau_2 \mid \tau_1 \times \tau_2 \mid R \mid \text{Owned}(R) \\
 \text{Expressions} & e ::= x \mid \lambda x. e \mid e_1 e_2 \mid (e_1, e_2) \mid \pi_i(e) \\
 & \quad \mid \text{new}_R \mid \text{use}(e_1, e_2) \mid \text{release}(e) \mid \text{borrow}(e_1, e_2) \\
 & \quad \mid \text{freeze}(e) \mid v \\
 \text{Values} & v ::= \lambda x. e \mid (v_1, v_2) \mid c \mid r \\
 \text{Resources} & r ::= \text{res}(a) \quad (\text{resource at address } a)
 \end{array}$$

where  $B$  ranges over base types and  $R$  ranges over resource type names.

**Remark 10.2** (Effects in  $\lambda^{\text{JR}}$ ). *The core calculus  $\lambda^{\text{JR}}$  deliberately omits effect tracking (`!o`, `State[s]`, etc.) to isolate the contribution of linear ownership. Effect annotations appear in the surface-language typing rules (Theorem 4.4) and in the resource-specific rules (*T-New*, *T-Release*, *T-Freeze*), but a full formalization of the effect system—including effect rows, handler semantics, and the interaction between effect polymorphism and linearity—is the subject of a companion paper on JAPL’s algebraic effect system. The two systems compose orthogonally: linearity governs how many times a resource is used, while effects govern what operations may be performed. Theorem 7.2 states the key compatibility property between the two.*

**Definition 10.3** (Store). *A store  $\sigma$  is a partial function from addresses to resource states:*

$$\sigma : \text{Addr} \rightarrow \text{ResourceState}$$

where  $\text{ResourceState} ::= \text{Live}(v) \mid \text{Released}$ .

## 10.2 Operational Semantics

We write  $\perp_R$  for the *default initial state* of resource type  $R$ : the distinguished value that a freshly allocated resource holds before any operations have been performed on it. Each resource type declaration implicitly defines its  $\perp_R$  (e.g., an empty buffer, an unconnected socket descriptor).

**Definition 10.4** (Small-step reduction). *The reduction relation  $\langle e, \sigma \rangle \longrightarrow \langle e', \sigma' \rangle$  is defined by:*

$$\begin{array}{c} \frac{}{\langle (\lambda x. e) v, \sigma \rangle \longrightarrow \langle [v/x]e, \sigma \rangle} \text{APP} \qquad \frac{a \notin \text{dom}(\sigma)}{\langle \text{new}_R, \sigma \rangle \longrightarrow \langle \text{res}(a), \sigma[a \mapsto \text{Live}(\perp_R)] \rangle} \text{NEW} \\ \\ \frac{\sigma(a) = \text{Live}(v)}{\langle \text{release}(\text{res}(a)), \sigma \rangle \longrightarrow \langle (), \sigma[a \mapsto \text{Released}] \rangle} \text{RELEASE} \\ \\ \frac{\sigma(a) = \text{Live}(v)}{\langle \text{freeze}(\text{res}(a)), \sigma \rangle \longrightarrow \langle v, \sigma[a \mapsto \text{Released}] \rangle} \text{FREEZE} \\ \\ \frac{\sigma(a) = \text{Live}(v)}{\langle \text{borrow}(\text{res}(a), f), \sigma \rangle \longrightarrow \langle (f v, \text{res}(a)), \sigma \rangle} \text{BORROW} \end{array}$$

## 10.3 Type Safety

**Definition 10.5** (Store typing). *A store typing  $\Sigma$  is a partial function from addresses to resource types:*

$$\Sigma : \text{Addr} \rightarrow \{\text{Live}(R) \mid \text{Released}\}$$

We write  $\sigma : \Sigma$  when for every  $a \in \text{dom}(\Sigma)$ :

- If  $\Sigma(a) = \text{Live}(R)$ , then  $\sigma(a) = \text{Live}(v)$  for some  $v$  of type  $R$ .
- If  $\Sigma(a) = \text{Released}$ , then  $\sigma(a) = \text{Released}$ .

**Theorem 10.6** (Progress). *If  $\Gamma; \Delta \vdash_\varepsilon e : \tau$  and  $\sigma : \Sigma$ , then either:*

- $e$  is a value, or
- there exist  $e', \sigma'$  such that  $\langle e, \sigma \rangle \longrightarrow \langle e', \sigma' \rangle$ .

*Proof.* By induction on the typing derivation  $\Gamma; \Delta \vdash_\varepsilon e : \tau$ .

**Case Var-Lin:**  $e = x$  and  $x : \text{Owned}\langle R \rangle \in \Delta$ . Then  $x$  is bound to a value (a resource reference  $\text{res}(a)$ ), so  $e$  is a value.

**Case Var-Un:** Similar.

**Case  $\multimap$ -I:**  $e = \lambda x. e'$ , which is a value.

**Case  $\multimap$ -E:**  $e = e_1 e_2$ . By induction, either  $e_1$  is a value or it can step. If  $e_1$  can step, so can  $e$ . If  $e_1$  is a value, it has type  $A \multimap B$ , so it is a lambda  $\lambda x. e'$ . By induction, either  $e_2$  is a value or it can step. If  $e_2$  can step, so can  $e$ . If  $e_2$  is a value  $v$ , then  $e$  can step by the App rule.

**Case New:**  $e = \text{new}_R$ . We can always find a fresh address  $a \notin \text{dom}(\sigma)$  (addresses are unbounded), so the New rule applies.

**Case Release:**  $e = \text{release}(\text{res}(a))$ . By the typing rule,  $\text{res}(a) : \text{Owned}\langle R \rangle$ , so  $a \in \text{dom}(\Sigma)$  with  $\Sigma(a) = \text{Live}(R)$ . Since  $\sigma : \Sigma$ , we have  $\sigma(a) = \text{Live}(v)$ , so the Release rule applies.

The remaining cases (Freeze, Borrow, pairs, projections) follow similarly.  $\square$

**Theorem 10.7** (Preservation). *If  $\Gamma; \Delta \vdash_\varepsilon e : \tau$  and  $\sigma : \Sigma$  and  $\langle e, \sigma \rangle \longrightarrow \langle e', \sigma' \rangle$ , then there exist  $\Gamma', \Delta', \Sigma'$  such that:*

- (a)  $\Gamma'; \Delta' \vdash_\varepsilon e' : \tau$
- (b)  $\sigma' : \Sigma'$
- (c)  $\Sigma \subseteq \Sigma'$  (the store typing only grows or transitions Live entries to Released)

*Proof.* By induction on the typing derivation, with case analysis on the reduction rule applied.

**Case App:**  $e = (\lambda x. e') v$  steps to  $[v/x]e'$ . By the T-App rule, the linear context is split as  $\Delta = \Delta_1 \uplus \Delta_2$  with  $\Gamma; \Delta_1 \vdash (\lambda x. e') : \tau_1 \multimap \tau$  and  $\Gamma; \Delta_2 \vdash v : \tau_1$ . By the substitution lemma for the dual-layer system (Theorem 10.8),  $\Gamma; \Delta' \vdash [v/x]e' : \tau$  where  $\Delta' = \Delta_1 \uplus (\Delta_2 \setminus \{x : \tau_1\})$ , i.e.,  $\Delta$  with  $x$ 's binding consumed. The store is unchanged.

**Case New:**  $e = \text{new}_R$  steps to  $\text{res}(a)$  with  $\sigma' = \sigma[a \mapsto \text{Live}(\perp_R)]$ . Set  $\Sigma' = \Sigma[a \mapsto \text{Live}(R)]$ . Then  $\text{res}(a) : \text{Owned}\langle R \rangle$  under  $\Delta' = \Delta, a : \text{Owned}\langle R \rangle$ , and  $\sigma' : \Sigma'$ .

**Case Release:**  $e = \text{release}(\text{res}(a))$  steps to  $()$  with  $\sigma' = \sigma[a \mapsto \text{Released}]$ . Set  $\Sigma' = \Sigma[a \mapsto \text{Released}]$ . The result  $()$  has type  $\text{Unit}$ . The linear context  $\Delta' = \Delta \setminus \{a : \text{Owned}\langle R \rangle\}$ : the binding is consumed.

**Case Freeze:** Similar to Release, but the result is the value  $v$  stored in the resource, which has a pure type.

**Case Borrow:**  $e = \text{borrow}(\text{res}(a), f)$  steps to  $(f v, \text{res}(a))$ . The resource remains live; the linear binding is preserved. The result is a pair of the function application result and the resource.

The remaining cases are standard.  $\square$

**Lemma 10.8** (Substitution). *If  $\Gamma; \Delta, x : A \vdash_\varepsilon e : B$  and  $\Gamma; \cdot \vdash v : A$  (if  $A$  is unrestricted) or  $v = \text{res}(a)$  with  $a : \text{Owned}\langle R \rangle$  (if  $A$  is linear), then  $\Gamma; \Delta' \vdash_\varepsilon [v/x]e : B$  where:*

- $\Delta' = \Delta$  if  $A$  is unrestricted.
- $\Delta' = \Delta$  if  $A$  is linear and  $x$  is consumed exactly once in  $e$ .

*Proof.* By induction on the typing derivation, tracking the number of occurrences of  $x$  in  $e$ . The linearity checker ensures exactly one occurrence for linear variables.  $\square$

## 10.4 Resource Safety

**Theorem 10.9** (Resource safety). *If  $\Gamma; \cdot \vdash_\varepsilon e : \tau$  (the program type-checks with an empty linear context) and  $\langle e, \emptyset \rangle \longrightarrow^* \langle v, \sigma \rangle$  (the program terminates with value  $v$  and final store  $\sigma$ ), then for every  $a \in \text{dom}(\sigma)$ :*

$$\sigma(a) = \text{Released}$$

*That is, every resource acquired during execution has been released.*

*Proof.* We establish the invariant: at every step of reduction, the number of Live entries in  $\sigma$  equals the size of the linear context  $\Delta$ .

**Base case:** Initially,  $\sigma = \emptyset$  and  $\Delta = \cdot$ , so the invariant holds trivially.

**Inductive step:** By preservation (Theorem 10.7), each reduction step preserves the correspondence between  $\Delta$  and  $\Sigma$ . Specifically:

- New: adds one entry to both  $\sigma$  (Live) and  $\Delta$ , preserving the count.

- Release/Freeze: transitions one entry in  $\sigma$  from Live to Released and removes one entry from  $\Delta$ , preserving the count.
- All other rules: preserve both  $\sigma$  and the size of  $\Delta$ .

**Conclusion:** When the program terminates with  $\Gamma; \cdot \vdash v : \tau$  (empty linear context), the number of Live entries in  $\sigma$  is zero. Therefore every entry is Released.  $\square$

**Corollary 10.10** (No resource leaks). *Well-typed JAPL programs (without `unsafe`) cannot leak resources: every file is closed, every socket is disconnected, every buffer is freed.*

**Corollary 10.11** (No use-after-free). *Well-typed JAPL programs cannot use a resource after it has been released: the linear type system ensures that the binding is consumed by the release operation and cannot be used again.*

**Corollary 10.12** (No double-free). *Well-typed JAPL programs cannot release a resource twice: the linear type system ensures exactly-once consumption.*

## 11 Discussion

### 11.1 When Ownership Pressure Leaks

The dual-layer design succeeds when the boundary between pure and resource code is clean. In practice, however, there are cases where ownership concerns propagate into otherwise-pure code:

**Resource-carrying data structures.** If a data structure contains a resource (e.g., a list of open file handles), the entire structure becomes linear. This forces the programmer to manage ownership for the container, even if the business logic only cares about the data inside.

**Callbacks with resources.** A higher-order function that accepts a callback may need to thread resources through the callback, complicating the interface.

**Long-lived resources.** A database connection pool manages resources with lifetimes that span many function calls. The pool itself becomes a resource, and every function that uses the pool must participate in the ownership protocol.

### 11.2 Mitigations

JAPL provides several mechanisms to reduce ownership pressure:

**Borrowing (ref).** Most pure functions that need to read a resource can take a borrowed reference, avoiding ownership transfer entirely.

**Scoped resources (use).** The `use` binding automatically manages resource lifetimes within a scope, reducing boilerplate.

**Process encapsulation.** Long-lived resources can be owned by a dedicated process, which exposes a pure message-passing interface. Other processes interact with the resource through messages, never seeing the ownership:

```
1 -- The database process owns the connection
2 fn db_process(conn: Owned[DbConnection])
3   → Never with Io, Process[DbMsg] =
4   match Process.receive() with
5   | Query(sql, reply) →
6     let (conn, result) = Db.query(conn, sql)
7     Reply.send(reply, result)
8     db_process(conn)
9   | Shutdown →
10    Db.close(conn)
11    Process.exit(Normal)
12
13 -- Callers never see Owned[DbConnection]
14 fn get_user(db: Pid[DbMsg], id: UserId)
15   → Result[User, DbError] with Process =
16   let result = Process.call(db, fn reply → Query(user_sql(id), reply))
17   parse_user_row(result)
```

Listing 18: Process encapsulation hides ownership

This pattern is the recommended way to manage long-lived resources in JAPL: the process boundary naturally encapsulates the ownership, and the caller’s interface is pure message-passing.

**Region-based grouping.** Multiple resources acquired in the same scope can be grouped into a single region, reducing the number of individual ownership threads.

### 11.3 The 80/20 of Resource Management

Our experience with prototype JAPL applications suggests the following distribution:

- **80%** of code is pure: data transformations, business logic, protocol handling, serialization. No ownership annotations needed.
- **15%** of code uses scoped resources (`use`): open file, read, close; connect, query, disconnect. Ownership is automatic via the `use` binding.
- **5%** of code requires explicit ownership management: long-lived resources, cross-process transfers, complex resource graphs. This is where the linear type system provides its greatest value.

The dual-layer design ensures that the 80% majority pays no annotation cost, while the 5% minority gets strong compile-time guarantees.

### 11.4 Effect System Synergy

The effect system and the ownership system reinforce each other. The effect system ensures that resource operations are visible in function signatures (`with Io`), preventing hidden side effects. The ownership system ensures that resources are properly managed even when effects are composed.

Together, they provide a stronger guarantee than either alone: a function with signature `fn(x: Int) -> Int` is guaranteed to be pure—it cannot secretly open a file, allocate a buffer, or access the network.

## 11.5 Comparison with Monadic IO

The Haskell approach confines *all* effects to the `IO` monad, including both resource management and general side effects. This conflates two concerns:

1. **Effect tracking:** Knowing that a function performs IO.
2. **Resource management:** Ensuring that resources are properly released.

JAPL separates these concerns: the effect system handles (1), and the ownership system handles (2). This separation allows finer-grained reasoning: two functions may both have effect `IO`, but one manages resources (visible in its type via `Owned`) while the other merely reads a file.

## 11.6 Limitations

The dual-layer approach has several limitations:

1. **GC dependency:** The pure layer requires a garbage collector, making JAPL unsuitable for environments where GC is unacceptable (embedded systems, real-time audio).
2. **Layer crossing cost:** Moving data between layers (freezing a buffer into immutable bytes) may involve a copy if the resource is not uniquely held.
3. **Learning curve:** Programmers must understand the distinction between pure and resource types, even if the pure layer requires no annotations.
4. **Expressiveness:** Some patterns that are natural in Rust (e.g., self-referential structures with lifetimes) are more difficult to express in JAPL’s simpler ownership model.

## 11.7 Future Work

Several directions remain for future investigation:

1. **Graded linearity:** Allowing resources to be used exactly  $n$  times (useful for connection pools with bounded concurrency).
2. **Session types:** Extending resource types with session type protocols, ensuring that resources are used in the correct order (e.g., connect, authenticate, query, disconnect).
3. **Dependent resource types:** Allowing resource types to depend on values, enabling richer invariants (e.g., a file handle that tracks whether it is open for reading, writing, or both).
4. **Formal verification:** Connecting JAPL’s type system to separation logic via the LNL model, enabling machine-checked proofs of resource safety.

## 12 Conclusion

We have presented JAPL’s second core principle: **mutation is local and explicit**. The dual-layer type system—pure values managed by GC, linear resources managed by ownership—achieves a design point that has not been explored by prior languages:

1. **Safety without pervasive annotation:** The majority of code lives in the pure layer and requires no ownership annotations, lifetime parameters, or borrow-checker reasoning.
2. **Resource safety by construction:** Every resource acquired by a well-typed program is released exactly once, as proven by our type safety and resource safety theorems.
3. **Effect visibility:** The effect system makes mutation visible in function signatures, enabling equational reasoning for pure code and local reasoning for effectful code.
4. **Process integration:** Ownership transfer between processes extends the linear discipline to concurrent settings, while process encapsulation hides ownership behind pure message-passing interfaces.

The formal framework—rooted in Girard’s linear logic, Benton’s LNL adjunction model, and the theory of substructural type systems—provides a rigorous foundation for the design. The categorical interpretation (Section 3.3) shows that JAPL’s dual-layer system is not an *ad hoc* engineering compromise but a principled instantiation of the well-studied relationship between linear and Cartesian closed categories.

We believe that the dual-layer approach represents a practical sweet spot for language design: it gives systems programmers the control they need over resources, while giving application programmers the simplicity they want for everyday code. The principle that mutation should be local and explicit—not absent, not pervasive, but *controlled*—is central to JAPL’s mission of being “pure by default, concurrent by design, resource-safe by construction, distributed without apology.”

## References

- [1] Jean-Yves Girard. Linear logic. *Theoretical Computer Science*, 50(1):1–102, 1987.
- [2] Philip Wadler. Linear types can change the world! In *Programming Concepts and Methods*, pages 561–581. North-Holland, 1990.
- [3] Philip Wadler. A taste of linear logic. In *Mathematical Foundations of Computer Science*, volume 711 of LNCS, pages 185–210. Springer, 1993.
- [4] Philip Wadler. Propositions as sessions. In *Proceedings of the 17th ACM SIGPLAN International Conference on Functional Programming (ICFP)*, pages 273–286, 2012.
- [5] David Walker. Substructural type systems. In Benjamin C. Pierce, editor, *Advanced Topics in Types and Programming Languages*, chapter 1, pages 3–43. MIT Press, 2005.
- [6] Benjamin C. Pierce. *Types and Programming Languages*. MIT Press, 2002.
- [7] Benjamin C. Pierce, editor. *Advanced Topics in Types and Programming Languages*. MIT Press, 2004.

- [8] P. Nick Benton. A mixed linear and non-linear logic: Proofs, terms and models. In *Computer Science Logic (CSL)*, volume 933 of LNCS, pages 121–135. Springer, 1995.
- [9] Michael Barr. \*-Autonomous categories and linear logic. *Mathematical Structures in Computer Science*, 1(2):159–178, 1991.
- [10] Jean-Philippe Bernardy, Mathieu Boespflug, Ryan R. Newton, Simon Peyton Jones, and Arnaud Spiwack. Linear Haskell: Practical linearity in a higher-order polymorphic language. *Proceedings of the ACM on Programming Languages*, 2(POPL):1–29, 2018.
- [11] The Rust Team. *The Rust Programming Language*. No Starch Press, 2021. <https://doc.rust-lang.org/book/>.
- [12] Ralf Jung, Jacques-Henri Jourdan, Robbert Krebbers, and Derek Dreyer. RustBelt: Securing the foundations of the Rust programming language. *Proceedings of the ACM on Programming Languages*, 2(POPL):1–34, 2017.
- [13] Erik Barendsen and Sjaak Smetsers. Uniqueness typing for functional languages with graph rewriting semantics. *Mathematical Structures in Computer Science*, 6(6):579–612, 1996.
- [14] Rinus Plasmeijer and Marko van Eekelen. *Clean Language Report, version 2.1*. University of Nijmegen, 2001.
- [15] Mads Tofte and Jean-Pierre Talpin. Region-based memory management. *Information and Computation*, 132(2):109–176, 1997.
- [16] Trevor Jim, J. Gregory Morrisett, Dan Grossman, Michael W. Hicks, James Cheney, and Yanling Wang. Cyclone: A safe dialect of C. In *USENIX Annual Technical Conference*, pages 275–288, 2002.
- [17] Dan Grossman, Greg Morrisett, Trevor Jim, Michael Hicks, Yanling Wang, and James Cheney. Region-based memory management in Cyclone. In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, pages 282–293, 2002.
- [18] Jack B. Dennis and Earl C. Van Horn. Programming semantics for multiprogrammed computations. *Communications of the ACM*, 9(3):143–155, 1966.
- [19] Mark S. Miller. *Robust Composition: Towards a Unified Approach to Access Control and Concurrency Control*. PhD thesis, Johns Hopkins University, 2006.
- [20] Gordon Plotkin and Matija Pretnar. Handlers of algebraic effects. In *Programming Languages and Systems (ESOP)*, volume 5502 of LNCS, pages 80–94. Springer, 2009.
- [21] Andrej Bauer and Matija Pretnar. Programming with algebraic effects and handlers. *Journal of Logical and Algebraic Methods in Programming*, 84(1):108–123, 2015.
- [22] Ohad Kammar, Sam Lindley, and Nicolas Oury. Handlers in action. In *Proceedings of the 18th ACM SIGPLAN International Conference on Functional Programming (ICFP)*, pages 145–158, 2013.
- [23] Andrew K. Wright and Matthias Felleisen. A syntactic approach to type soundness. *Information and Computation*, 115(1):38–94, 1994.

- [24] Joe Armstrong. *Programming Erlang: Software for a Concurrent World*. Pragmatic Bookshelf, 2003.
- [25] Simon Peyton Jones. *Haskell 98 Language and Libraries: The Revised Report*. Cambridge University Press, 2003.
- [26] Xavier Leroy. *The OCaml System: Documentation and User’s Manual*. INRIA, 2000.
- [27] Alan A. A. Donovan and Brian W. Kernighan. *The Go Programming Language*. Addison-Wesley, 2015.
- [28] Sylvan Clebsch, Sophia Drossopoulou, Sebastian Blessing, and Andy McNeil. Deny capabilities for safe, fast actors. In *Proceedings of the 5th International Workshop on Programming Based on Actors, Agents, and Decentralized Control*, pages 1–12, 2015.
- [29] Jesse A. Tov and Riccardo Pucella. Practical affine types. In *Proceedings of the 38th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, pages 447–458, 2011.
- [30] Daan Leijen. Type directed compilation of row-typed algebraic effects. In *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages (POPL)*, pages 486–499, 2017.
- [31] Karl Mazurak, Jianzhou Zhao, and Steve Zdancewic. Lightweight linear types in System F<sup>o</sup>. In *Proceedings of the 5th ACM SIGPLAN Workshop on Types in Language Design and Implementation (TLDI)*, pages 77–88, 2010.
- [32] Neel Krishnaswami, Pierre Pradic, and Nick Benton. Integrating linear and dependent types. In *Proceedings of the ACM on Programming Languages*, 3(POPL):1–28, 2019.
- [33] Peter Sewell, Francesco Zappa Nardelli, Scott Owens, Gilles Peskine, Thomas Ridge, Susmit Sarkar, and Rok Strniša. Ott: Effective tool support for the working semanticist. *Journal of Functional Programming*, 20(1):71–122, 2010.
- [34] Robin Milner. A theory of type polymorphism in programming. *Journal of Computer and System Sciences*, 17(3):348–375, 1978.
- [35] John C. Reynolds. Types, abstraction and parametric polymorphism. In *Information Processing 83*, pages 513–523. Elsevier, 1983.
- [36] Eugenio Moggi. Notions of computation and monads. *Information and Computation*, 93(1):55–92, 1991.
- [37] Fernando Borretti. The Austral programming language. <https://austral-lang.org/>, 2023.
- [38] Samson Abramsky. Computational interpretations of linear logic. *Theoretical Computer Science*, 111(1–2):3–57, 1993.
- [39] Marco Gaboardi, Shin-ya Katsumata, Dominic Orchard, Flavien Breuvar, and Tarmo Uustalu. Combining effects and coeffects via grading. In *Proceedings of the 21st ACM SIGPLAN International Conference on Functional Programming (ICFP)*, pages 476–489, 2016.

## A Complete Typing Rules for $\lambda^{\text{JR}}$

For reference, we present the complete set of typing rules for the core calculus.

**Definition A.1** (Full typing rules).

$$\begin{array}{c}
\frac{x : \tau \in \Gamma}{\Gamma; \cdot \vdash x : \tau} \text{ T-VAR-UN} \qquad \frac{}{\Gamma; x : \tau \vdash x : \tau} \text{ T-VAR-LIN} \qquad \frac{\Gamma; \Delta, x : \tau_1 \vdash e : \tau_2}{\Gamma; \Delta \vdash \lambda x. e : \tau_1 \multimap \tau_2} \text{ T-ABS} \\
\\
\frac{\Gamma, x : \tau_1; \cdot \vdash e : \tau_2}{\Gamma; \cdot \vdash \lambda x. e : \tau_1 \rightarrow \tau_2} \text{ T-ABS-UN} \qquad \frac{\Gamma; \Delta_1 \vdash e_1 : \tau_1 \multimap \tau_2 \quad \Gamma; \Delta_2 \vdash e_2 : \tau_1}{\Gamma; \Delta_1, \Delta_2 \vdash e_1 e_2 : \tau_2} \text{ T-APP} \\
\\
\frac{\Gamma; \Delta_1 \vdash e_1 : \tau_1 \rightarrow \tau_2 \quad \Gamma; \Delta_2 \vdash e_2 : \tau_1}{\Gamma; \Delta_1, \Delta_2 \vdash e_1 e_2 : \tau_2} \text{ T-APP-UN} \qquad \frac{\Gamma; \Delta_1 \vdash e_1 : \tau_1 \quad \Gamma; \Delta_2 \vdash e_2 : \tau_2}{\Gamma; \Delta_1, \Delta_2 \vdash (e_1, e_2) : \tau_1 \otimes \tau_2} \text{ T-PAIR} \\
\\
\frac{\Gamma; \Delta_1 \vdash e_1 : \tau_1 \otimes \tau_2 \quad \Gamma; \Delta_2, x : \tau_1, y : \tau_2 \vdash e_2 : \tau}{\Gamma; \Delta_1, \Delta_2 \vdash \mathbf{let} (x, y) = e_1 \mathbf{in} e_2 : \tau} \text{ T-LETPAIR} \\
\\
\frac{}{\Gamma; \cdot \vdash_{\text{lo}} \mathbf{new}_R : \text{Owned}\langle R \rangle} \text{ T-NEW} \qquad \frac{\Gamma; \Delta \vdash e : \text{Owned}\langle R \rangle}{\Gamma; \Delta \vdash_{\text{lo}} \mathbf{release}(e) : \text{Unit}} \text{ T-RELEASE} \\
\\
\frac{\Gamma; \Delta \vdash e : \text{Owned}\langle R \rangle \quad R \hookrightarrow \tau}{\Gamma; \Delta \vdash_{\text{lo}} \mathbf{freeze}(e) : \tau} \text{ T-FREEZE} \\
\\
\frac{\Gamma; \Delta, x : \text{Owned}\langle R \rangle \vdash e_1 : \text{Owned}\langle R \rangle \quad \Gamma; y : \text{ref } R \vdash e_2 : \tau}{\Gamma; \Delta, x : \text{Owned}\langle R \rangle \vdash \mathbf{borrow}(x \text{ as } y, e_2) : \tau \otimes \text{Owned}\langle R \rangle} \text{ T-BORROW} \\
\\
\frac{\Gamma; \Delta_0 \vdash e : \text{Bool} \quad \Gamma; \Delta \vdash e_1 : \tau \quad \Gamma; \Delta \vdash e_2 : \tau}{\Gamma; \Delta_0, \Delta \vdash \mathbf{if } e \mathbf{ then } e_1 \mathbf{ else } e_2 : \tau} \text{ T-IF}
\end{array}$$

Note that *T-Abs-Un* requires the linear context to be empty ( $\cdot$ ): an unrestricted function  $\tau_1 \rightarrow \tau_2$  may be duplicated, so it must not close over linear resources—otherwise, applying the closure twice would duplicate those resources, violating linearity.

Note that *T-If* requires both branches to consume the same linear context  $\Delta$ . This is the key rule that prevents conditional resource leaks.

## B Extended Examples

### B.1 Database Connection Pool

```

1 resource DbConnection
2
3 type PoolMsg =
4   | Checkout(Reply[Owned[DbConnection]])
5   | Return(Owned[DbConnection])
6   | Shutdown

```

```

7
8 fn pool_process(conns: List[Owned[DbConnection]],
9                 waiting: List[Reply[Owned[DbConnection]]])
10  → Never with Io, Process[PoolMsg] =
11 match Process.receive() with
12 | Checkout(reply) →
13   match conns with
14   | [conn, ..rest] →
15     Reply.send(reply, conn)
16     pool_process(rest, waiting)
17   | [] →
18     pool_process(conns, List.append(waiting, [reply]))
19 | Return(conn) →
20   match waiting with
21   | [waiter, ..rest] →
22     Reply.send(waiter, conn)
23     pool_process(conns, rest)
24   | [] →
25     pool_process([conn, ..conns], waiting)
26 | Shutdown →
27   List.each(conns, fn conn → Db.close(conn))
28   Process.exit(Normal)

```

Listing 19: A resource-safe connection pool

## B.2 State Effect with Linear Resources

```

1 -- Explicit mutable state via effect
2 fn accumulate(items: List[Int]) → Int with State[Int] =
3   List.each(items, fn x → State.modify(fn acc → acc + x))
4   State.get()
5
6 -- Resource layer: mutable buffer (ownership tracked)
7 fn fill_buffer(data: Bytes) → Bytes with Io =
8   let buf = Buffer.alloc(1024)
9   let buf = Buffer.write(buf, 0, data)
10  Buffer.freeze(buf)
11
12 -- Combining pure, state, and resource layers
13 fn process_and_store(items: List[Int], path: Path)
14  → Result[Unit, IOError] with Io =
15   let total = State.run(0, fn → accumulate(items))
16   let file = File.open(path, Write)?
17   File.write(file, Int.to_string(total))
18   File.close(file)
19   Ok(())

```

Listing 20: Composing effects with linear resources

## B.3 Safe FFI Wrapping

```

1  -- Foreign resource: a C library handle
2  resource CHandle
3
4  foreign "C" fn c_lib_open(config: CString) →Ptr[CHandle]
5  foreign "C" fn c_lib_process(h: Ptr[CHandle], data: Ptr[Byte], len: CInt) →CInt
6  foreign "C" fn c_lib_close(h: Ptr[CHandle]) →Unit
7
8  -- Safe JAPL wrapper: all ownership is explicit
9  fn open_handle(config: String) →Result[Owned[CHandle], Error] with Io =
10     use cstr = CString.from(config)
11     let ptr = unsafe c_lib_open(cstr)
12     if Ptr.is_null(ptr) then Err(OpenFailed)
13     else Ok(CHandle.from_raw(ptr))
14
15  fn process_data(handle: Owned[CHandle], data: Bytes)
16     → (Owned[CHandle], Result[Int, Error]) with Io =
17     let result = unsafe c_lib_process(
18         CHandle.as_ptr(ref handle),
19         Bytes.as_ptr(data),
20         Bytes.length(data)
21     )
22     if result < 0 then (handle, Err(ProcessFailed(result)))
23     else (handle, Ok(result))
24
25  fn close_handle(handle: Owned[CHandle]) →Unit with Io =
26     unsafe c_lib_close(CHandle.into_raw(handle))

```

Listing 21: Safe wrapper around a C FFI resource