

Lento: An Expressive, Safe, and Efficient Language

[DRAFT 2026-03-01]

William Rågstad
Independent Researcher

Abstract—The result of decades of programming language research can become diluted when retrofitted into mainstream languages. Legacy constraints and compatibility requirements force compromises that lead to complex and unwieldy designs. Consequently, redesigning a new modern language from the ground up is often necessary to fully leverage state-of-the-art advances in typing, verification, and compilation. In this paper we introduce a *high-level and work-in-progress* vision of Lento, a language designed to bridge the gap between high-level abstraction and low-level performance.

Index Terms—Programming Language, Parallelism, Guided Synthesis, Program Specification, Formal Verification, Liquid Types, Dependent Types, Session Types, Effect System, Type Inference

I. INTRODUCTION

Programming languages today often force developers to choose between high-level abstractions that enhance productivity and low-level control that optimizes performance. This dichotomy has led to a fragmented landscape where languages excel in one domain but falter in the other. High-level languages, such as Python and JavaScript, prioritize ease of use and rapid development but often sacrifice performance and compile-time type safety. Conversely, low-level languages like C and Rust offer fine-grained control over system resources but come with steep learning curves, increased complexity, and higher demands on expertise.

Lento aims to bridge this gap by introducing a novel programming language while also incorporating the benefits of the pure functional paradigm. Leveraging strong static typing, lazy evaluation, automated parallelism and concurrency, user-friendly syntax, and more, to offer several advantages over existing languages. It is designed to be safe, performant, and stable while promoting longevity in codebases.

II. VALUES

In the Lento compiler team we define a fundamental set of values people developing the language should adhere to. These extend to the language standard specification and form the basis for all future features and decisions. Our standpoints on these values are outlined below and are goals we strive to achieve in the development of Lento.

A. Longevity

Lento should be a language built to last for decades. We have learned from the mistakes of others and decide to support multiple versions of the language (called language

editions¹) simultaneously. Where libraries in one edition must be backward compatible and seamlessly interoperate with those compiled in other editions. Editions with major breaking changes should include migration guidelines. This allows users to choose when to upgrade and ensures that older codebases remain functional with long-term support (LTS).

B. Stability

Lento should prioritize stability in its features, standard library APIs, and language design.

Once a feature is released, it should not change in a way that breaks existing code. To achieve this, rigorous reviewing of new features include both community feedback and thorough testing. Most importantly, features are designed with future compatibility in mind from the start. Defining a small “external” feature interface to the language, allowing internal changes with semantically equivalent behavior.

C. Performance

Lento strive to be a high-performance language in terms of execution speed and resource efficiency. However, our goal is not to compete with other low-level languages, but rather to provide predictable performance characteristics.

D. Bootstrapping

One controversial standpoint and conscious design choice is that Lento should not be bootstrapped in itself. Instead, we believe that an implementation written in the low-level systems language Rust is a better choice for achieving our goals for the near future. This enables faster and more flexible development between major language editions while still allowing for future self-hosting when Lento is matured. Rust offers the Lento compiler a solid foundation for performance, safety, and cross-platform compatibility.

E. Safety

Lento is designed with safety as a core principle. This includes memory safety, type safety, concurrency safety and behavioral safety through formal program specifications. By leveraging a strong static type system with dependent liquid refinements and a static effect-system, Lento aims to catch errors at compile-time, reducing runtime failures and ensuring that programs behave as intended. Lento builds on top of the principle of immutability as a cornerstone of functional programming, making programs easier to reason

¹This is similar to Rust’s approach with new editions every 3 years, see <https://blog.rust-lang.org/2014/10/30/Stability>.

about and debug. Mutability and side effects are handled safely while maintaining referential transparency.

F. Interoperability

All editions of Lento must be backwards compatible with each other at the binary level through a stable Application Binary Interface (ABI). Two programs compiled with different editions of Lento must be able to interoperate without issues.

Lento should also prioritize interoperability with existing languages and ecosystems. This includes providing a robust Foreign Function Interface (FFI) and ABI marshalling for seamless integration with languages such as C/C++, Rust, and Python.

G. Tooling and Ecosystem

Lento should provide a comprehensive tooling ecosystem to enhance developer productivity. This includes a package manager for easy dependency management, IDE integrations for code completion and debugging, and build tools for efficient project management.

H. Community

Our final standpoint is that we should foster a welcoming and inclusive community around Lento that encourages constructive collaboration and knowledge sharing. This involves creating clear contribution guidelines, maintaining open communication channels, and actively engaging with users and contributors to gather feedback and improve the language.

III. FEATURES

Lento aims to provide a robust set of features that enhance developer productivity while maintaining the core values of the language.

A. Strong Static Type System

This is a fundamental feature and core value of the language itself. Lento provides a dependent type system with liquid refinements and effects to catch both *regular type errors* and *logical contradictions* in specifications all at compile-time using bidirectional type inference without sacrificing expressivity. See Section V for a more detailed explanation of Lento’s type system.

B. Units of Measurement

Lento supports user-defined *units of measure* as a lightweight, inference-friendly suffix on numeric literals for algebraic domains. The goal is to prevent accidental mixing of incompatible quantities (e.g. meters vs seconds) while keeping the runtime model unchanged: unit annotations are checked statically and erased (no runtime wrappers). This follows classic work on dimension typing, analysis and inference and models for polymorphism over physical dimension. [1], [2], [3], [4]

Concretely, numeric literals may carry a unit expression suffix such as `10m/s`, `9.81m/s^2`, or `500kcal`. Suffix elabora-

tion runs at compile time during elaboration/typechecking and must be pure and deterministic.

```

1 fn c_to_k(v) = v + 273.15
2 fn f_to_k(v) = (v - 32.0) * 5.0/9.0 + 273.15
3 type Temp(kelvin: f64)
4 unit K = fn v => Temp(v)
5 unit C = fn v => Temp(c_to_k(v))
6 unit F = fn v => Temp(f_to_k(v))
7 let t: Temp = 200K - 5C + 10F

```

Lento lets numeric literals use unit suffixes (e.g., `5C`, `10F`, `3K`), which are mapped at compile time to a single semantic type (like `Temp`) using pure, static functions. Different units are converted to a consistent basis (e.g., kelvin), so arithmetic like `5C + 10F` works seamlessly and type-checks, with all conversions resolved at compile time. For example, a library may register a complex unit such that `2i` elaborates to the imaginary injection into `Complex` (or quaternions), while keeping `i` as an ordinary identifier outside suffix position. APIs may also require explicit units (e.g., `sleep(10s)`; `sleep(10)` is a type error).

C. Automated Parallelism

Lento targets scalable task-level parallelism with automatic management of parallel work granularity and scheduling, informed by recent advances in automatic parallelism management [5] and classic approaches to automatic parallelization [6].

D. Concurrency Model

E. Memory Model

Region-based memory management and memory safety without garbage collection via linear/affine types, ownership, and borrowing.

F. Program Synthesis

Lento incorporates program synthesis techniques to automatically generate code snippets based on type-level program specifications provided by the developer, see Section V.H. This aligns with type-driven synthesis from polymorphic refinement types [7].

G. Multistage Meta-Programming

Lento supports multistage meta-programming for constant evaluation and staging code generation at compile-time using a staged calculus in a type-safe way with phase separation that is more powerful than traditional hygienic macros.

H. Partial Evaluation

Lento supports partial evaluation to optimize programs by precomputing parts of the code at compile-time based on known inputs and contexts, resulting in more efficient runtime performance. This is *similar to multistage programming* but focuses on optimization rather than code generation, see Section III.G.

I. Holes

Lento supports the use of holes in code, allowing developers to leave parts of the program unspecified during rapid development. Instead of employing gradual typing, holes in Lento must be statically typed but can be left unimplemented. This marks a place for the compiler to “fill in” an implementation using program synthesis techniques (see Section III.F) based on the surrounding context, type information, and specifications.

J. Standard Library

Lento should provide two sets of *officially-supported* standard libraries. First a stable cross-platform core, and extended ext libraries with wide range of functionality out of the box². The extended libraries may have platform-specific dependencies and less stability guarantees compared to the core libraries. See Fig. 1 for the standard library lattice in Lento.




Fig. 1. Standard library hierarchy with core core in blue, extended ext libraries in purple, and user-defined packages in green.

a) *Core core*: Essential data structures (lists, maps, sets, graphs), algorithms (sorting, searching), concurrency (async/await), networking (TCP, UDP), and file I/O APIs.

b) *Extended ext*: Additional libraries and frameworks for common tasks and domains, such as OS interaction (Win32, POSIX), web client and server development (HTTP, WebSocket), data processing (CSV, JSON), databases (SQL, NoSQL), GUI toolkits, scientific computing (numerical methods, statistics), computer graphics (2D/3D rendering, game development), compression (gzip, zlib), cryptography (hashing, encryption), machine learning (linear algebra, neural networks), and more.

c) *Third-party libc*: Lento does not aim to “reinvent the wheel”, but use what knowledge and tools are available. One such resource is libc, a powerful and backwards compatible API with highly efficient implementations. The libc API is cross-platform, offered via GNU and MUSL on Linux, and UCRT on Windows. This library is a trusted base to build on top of with the benefits of:

- Less library code in Lento thanks to reusability.
- Dynamic linking results in light-weight artifacts.
- Performant runtime via optimized implementations.

K. Tooling Ecosystem

A language is only as good as its tooling and ecosystem. Lento should therefore provide a comprehensive developer tooling to enhance productivity and streamline processes.

a) *Package Management*: Lento should provide a built-in package manager to facilitate easy dependency management, versioning, and distribution of libraries and applications. This includes support for semantic versioning, dependency resolution, and publishing packages to an official default central repository configurable by the user. Self-hosted registries should also be supported for private proprietary use cases.

b) *IDE Integrations*: LSP support for popular text editors to provide code completion, syntax highlighting, inline error diagnostics, and debugging capabilities.

c) *Build Tools*:

d) *Testing Framework*:

e) *Documentation Generation*: The language should include tools for generating comprehensive documentation from function signatures, but also from code comments and type annotations. Comments should support markdown-like syntax for formatting, code snippets, and linking to other parts of the documentation. The primary source of truth should always be type signatures to ensure accuracy and consistency thanks to Lento’s strong typing with dependent refinements and formal specifications.

IV. LANGUAGE SYNTAX

This section presents the formal syntax of expressions, types, and program structure.

A. Lento by Example

We begin with a brief overview of Lento’s syntax via some examples. Consider the following recursive fact function:

```
1 fn fact(n: int, acc = 1): int = match n
2   | 0 => acc,
3   | _ => fact(n - 1, n * acc),
```

Here, we define a function fact that takes an integer n and returns its fact. The match expression allows us to pattern match on the value of n, providing different cases for 0 and other integers.

We could also define the type of fact using dependent types and liquid refinements to specify that the result is always non-negative and utilize pattern matching in argument bindings [8]:

```
1 type Nat = { v: int | v >= 0 }
2 fact :: Nat -> Nat -> Nat
3 fn fact(n, acc) = fact(n - 1, n * acc)
4 fn fact(0, acc) = acc
```

Note that the function clause order does not matter *here*, due to automatic ordering of case specificity placing the general case last for us. This is usually done automatically by the compiler, but for ambiguous cases, the compiler requires us to specify the order ourselves explicitly.

B. Formal Syntax

Language terms are defined by the following expression grammar:

²This is heavily inspired by both the Rust and Go standard library designs, see <https://doc.rust-lang.org/std> and <https://pkg.go.dev/std>.

$$\begin{aligned}
e &:: \text{cnst} \mid \text{stm} \mid e_1 \ e_2 \mid (p : \tau) \Rightarrow e \\
&\mid \oplus e \mid e_1 \oplus e_2 \mid e \oplus \{e_1, \dots, e_n\} \\
&\mid [e_1, \dots, e_n] \mid \{l_i : e_i\} \mid (e_1, \dots, e_n) \\
&\mid T(e_1, \dots, e_n) \mid (e : \tau) \mid (e) \\
\text{cnst} &:: () \mid n \mid n \ u \mid n.n \mid n.n \ u \mid \text{true} \mid \text{false} \mid \text{"s"} \mid \text{'c'} \\
u &:: 1_u \mid id_u \mid u_1 * u_2 \mid u_1 / u_2 \mid u^n \mid (u) \\
\text{stm} &:: \text{id} :: \tau \\
&\mid \text{let } p = e_1 \text{ in } e_2 \\
&\mid \text{type } \tau_i = \tau_j \text{ in } e \\
&\mid \text{fn } f(p_1, \dots, p_n) : \tau = e_1 \text{ in } e_2 \\
&\mid \text{match } e \text{ with } \{p_1 \Rightarrow e_1, \dots, p_n \Rightarrow e_n\} \\
&\mid \text{if } e_1 \text{ then } e_2 \text{ else } e_3 \\
&\mid \text{import } m \text{ in } e \\
p &:: \text{cnst} \mid \text{id} \mid _ \mid p : \tau \mid \oplus p \mid p_1 \oplus p_2 \mid p \oplus \\
&\mid [p_1, \dots, p_n] \mid \{l_i : p_i\} \mid (p_1, \dots, p_n) \\
&\mid T(p_1, \dots, p_n) \\
m &:: \text{id} \mid \text{id} . m \mid \text{id}_1 \ \text{as} \ \text{id}_2 \mid \{m_1, \dots, m_n\}
\end{aligned}
\tag{1}$$

Both n and i are metavariables ranging over numeric literals and identifiers, respectively. The metavariable id_u ranges over valid unit-of-measure symbols, not to be confused with the cnst constant $()$: **unit** representing a “nothing” value or empty tuple of type `unit`. In the concrete syntax, unit/suffix literals allow optional whitespace (e.g. [2i](#) and [5 m/s](#)). The constants `s` and `c` represent arbitrary string and character literals. The symbol τ ranges over types and T type constructors such as $\text{Some}(e)$ of type `Option` τ where $e : \tau$, defined in Section V.A. The symbol \oplus ranges over operators (arithmetic, logical, comparison, etc.) in different fixity positions (prefix, infix, postfix). Statements in stm use standard keywords such as **let**, **fn**, **match**, and **if** that are entirely syntactic sugar and desugar to core expressions in e . Only **match** is a fundamental construct in the operational semantics, see Section VI. The formal “in” syntax is *implicit* and only included for formality.

V. TYPE SYSTEM

The design unifies several advanced features within its strong static type system and algebraic effect system to catch a wide range of logical errors at compile time. The type syntax is expressive yet easy to use due to our inference-first principle, allowing developers to define complex types and guarantee properties about their program behavior while minimizing boilerplate code. This chapter describes a formal operational semantics defining program execution: surface-language constructs elaborate into a typed core calculus in which typing and refinements are checked. Type aliases are supported as syntactic sugar but do not impact the core type system.

A. Type Syntax

The core type language is defined inductively for any type-level term τ as follows:

$$\begin{aligned}
\tau &:: \text{id} \mid \mathbf{Type}_i \mid A \mid B \mid C \mid D \mid E \mid M \mid L \mid S \mid N \\
A &:: \tau_i \rightarrow \tau_j \mid \tau_i \ \tau_j \mid \tau_i(\tau_j, \dots, \tau_n) \mid \tau \ \mathbf{all} \ \text{id}_1 : \tau_1, \dots, \text{id}_n : \tau_n \\
B &:: \text{unit} \mid \text{bool} \mid \text{char} \mid \text{str} \mid \text{uint}_u \mid \text{int}_w \mid \text{float}_v \mid \text{num} \\
C &:: \text{record}\{\ell_i : \tau_i\} \mid \text{inductive}\{\ell_i : \tau_i\} \\
D &:: \Pi x : \tau_i. \tau_j \mid \Sigma x : \tau_i. \tau_j \\
E &:: x : \tau_i \rightarrow \varepsilon \rightarrow \tau_j \mid \tau_i ! \varepsilon \rightarrow \tau_j \mid \tau_i ? \varepsilon \rightarrow \tau_j \\
M &:: \Box \tau \mid \Diamond \tau \\
L &:: \{\nu : \tau \mid \varphi\} \\
S &:: \mathbf{end} \mid !\tau.S \mid ?\tau.S \mid \mu t : S \mid \oplus \{\ell_i : S_i\} \mid \&\{\ell_i : S_i\} \\
N &:: \text{literal}(\text{cnst}) \\
\varphi &:: \text{cnst} \mid \nu \oplus \text{cnst} \mid \mathbf{f}_{\text{pure}} \nu \mid \varphi_1 \wedge \varphi_2 \mid \varphi_1 \vee \varphi_2 \mid \neg \varphi
\end{aligned}
\tag{2}$$

Note that $i, j \in I$ ranges over indices in finite index sets I , $u \in \{1, \dots, 128, \text{big}\}$ and $w \in \{8, \dots, 128, \text{big}\}$ denotes signed/unsigned integer bit-widths, $v \in \{32, 64, \text{big}\}$ denotes a floating-point bit-width. Literal constants c narrow types of category B . Predicates φ range over an SMT-decidable logic, \oplus abbreviates a comparison operator (e.g., $=, \neq, <, \leq, >, \geq$) against a constant or a pure comptime function. The refinement variable ν ranges over the value denoted by a refined term. Type categories are as follows:

- A : Type constructor application of parameterized types and type-level functions (e.g., `Map` $\tau_i \ \tau_j$).
- B : Primitive base types and literals.
- C : Composite types (records and recursive sum types with type constructors).
- D : Dependent types.
- E : Effect-annotated function types.
- M : Modal types.
- L : Liquid refinement types.
- S : Session types.
- N : Literal singleton types (for constant propagation and partial evaluation), e.g. `literal(42)`.
- \mathbf{type}_i : The universe of types at level $i \in \mathbb{N}$ (to support cumulative type hierarchies and avoid Girard’s paradox).

The standard library provides numerous predefined types and type constructors built on top of this core syntax, see Section III.J. These include common data structures (`Tuple` $\tau_i \dots \tau_j$, `Vec` $\nu \ \tau$, `Map` $\tau_i \ \tau_j$, `Set` τ), parameterized polymorphic algorithms (sorting, searching), concurrency abstractions (`async/await` effects), and networking primitives (`TCP/UDP` channels with session types).

B. Polymorphism and Type Classes

Lento supports universal polymorphism via a single surface binder **all**. In signatures, **all** is written postfix to keep the term-level shape primary:

```
1 id :: a -> a all a: Type Lento
```

Constraints are written with a postfix **all** clause and elaborate to implicit evidence parameters. To keep inference predictable, Lento uses coherent global instance resolution: for any given head (e.g., [Eq T](#)), instance search either finds

a unique best instance or fails; ambiguous constraints are a type error.

```

1 type class Eq = { eq : Self -> Self -> bool }
2 let (Eq int).eq = prim_int_eq
3 fn (Eq List T).eq(xs, ys) all T : Eq = list_eq_by(T.eq, xs, ys)
4 infix left 4 (==) : a -> a -> bool all a = (Eq a).eq
5 member :: a -> List a -> bool all a : Eq
6 fn member(x, xs) = ...

```

Session-typed communication uses the session-indexed channel type $\mathbf{Chan}(S)$ (see Section V.F). Concrete transports (e.g., TCP, in-memory queues) can be abstracted behind type classes, but we omit the backend API here.

Example composite types in C are instances of record types and inductive types (algebraic data types) such as tuples, lists, 2D points, and option types:

```

Tuple ::= record{0 : int, 1 : bool, 2 : float}
List  $\tau$  ::= inductive{Nil : unit,
                    Cons :  $\Sigma$  head :  $\tau$ . List}
Point2D ::= record{x : int, y : int}
Option  $\tau$  ::= sum{Some :  $\tau$ , None : unit}

```

Types in D include dependent function types (Π types) and dependent pair types (Σ types). Algebraic effect-annotations ε denote function side-effects in category E , while refined types are in L (see Section V.D) and session types are in S (see Section V.F).

C. Dependent Types

Types can be parameterized by values to encode rich invariants and relationships. This enables the expression of rich invariants and relationships between data directly in the type system. For example, we can define a vector type parameterized by its length:

$$\text{Vec}(\nu : \text{Int}, \tau : \text{Type}) = \{\nu : \text{List}(\tau) \mid \text{len}(\nu) = \nu\} \quad (4)$$

Here, $\text{Vec}(n, \tau)$ represents a list of elements of type τ with length exactly n . This permits specifications such as length-indexed collections, value-range invariants, and relational postconditions, while still admitting decidable checking once refinements are restricted to an SMT-decidable fragment. Foundational accounts of dependent and advanced types can be found in [9] and [10].

D. Liquid Refinement Types

Liquid refinements qualify base types with logical predicates, yielding logically qualified data types. Refined base types have the form $\{\nu : B \mid \varphi\}$, where B is a base type and φ is an SMT-decidable predicate verifying properties of the value ν at compile-time. Unless otherwise stated, predicates are interpreted under the variable assignments induced by Γ . For example, the type $\{\nu : \text{Int} \mid 0 \leq \nu \wedge \nu < n\}$ denotes integers in the half-open interval $[0, n)$. To preserve decidability, the refinement logic is restricted to an SMT-decidable subset (typically a quantifier-free fragment), including theories such as linear arithmetic, uninterpreted functions, and finite maps. See [8] for the original development and [11] for subsequent extensions.

E. Modal Types

Lento incorporates modal type constructors to model staged and time-dependent computation. We write necessity and possibility modalities as $\Box\tau$ and $\Diamond\tau$ respectively. Intuitively, $\Diamond\tau$ classifies computations that may produce a τ in the future. These modalities interact with staged meta-programming (see Section III.G), effect tracking (see Section V.G), concurrency abstractions (see Section III.D), and parallelism (see Section III.C). For example, asynchronous functions can be assigned modal signatures that make temporal availability explicit:

$$\begin{aligned} \text{async fetch}(\text{url} : \text{String}) &\rightarrow \Diamond \text{Data} \\ \text{await} : \Diamond \text{Data} &\rightarrow \text{Data} \end{aligned} \quad (5)$$

Here $\Diamond \text{Data}$ records that the result is not necessarily available required by the temporal aspect of this particular asynchronous computational effect.

F. Session Types

Lento uses session types to specify and enforce structured communication protocols. Channels $\mathbf{Chan}(S)$ are parameterized by a session type S to statically guarantee correct sequences of message passing between two parties. A session type consists of special type operators modeling sending and receiving messages, branching, selection, and recursion. The underlying implementation of channels can be separated from the session type specification, allowing for various transport mechanisms (e.g., TCP, in-memory queues), optimizations (e.g., zero-copy), or serialization formats (e.g., JSON, Protobuf³). We write output and input actions as $!\tau.S$ and $?\tau.S$, respectively, together with external choice (branching) and internal choice (selection):

$$\begin{aligned} S ::= \mathbf{end} \mid t \mid !\tau : S \mid ?\tau : S \mid \mu\tau : S \\ \mid \&\{\ell_i : S_i\}_{i \in I} \mid \oplus \{\ell_i : S_i\}_{i \in I} \end{aligned} \quad (6)$$

Where \mathbf{end} denotes the end of a session and t a basic type in B , see Section V.A. For real distributed systems with multiparty sessions (MST) we can use the global protocol type G to describe the overall communication protocol among multiple participants, and project it onto local session types $S_p = G \downarrow p$ for each participant p to ensure that each party adheres to the global protocol. This guarantees communication safety and protocol compliance by construction, with automatic duality between send and receive operations (client/server) by inverting sends/receives and select/branch direction, denoted \bar{S} .

$$\begin{aligned} \overline{!\tau.S} = ?\tau.\bar{S} \quad \overline{?\tau.S} = !\tau.\bar{S} \quad \overline{\mu t.S} = \mu t.\bar{S} \\ \overline{\oplus \{\ell_i : S_i\}_{i \in I}} = \&\{\ell_i : \bar{S}_i\}_{i \in I} \\ \overline{\&\{\ell_i : S_i\}_{i \in I}} = \oplus \{\ell_i : \bar{S}_i\}_{i \in I} \end{aligned} \quad (7)$$

³Protocol Buffers are language-neutral, platform-neutral extensible mechanisms for serializing structured data. From <https://protobuf.dev>

G. Effect System

Lento incorporates static algebraic liquid effect system to reason about and handle side effects in programs such as I/O operations, state mutations, exceptions, concurrency, and even deterministic parallelism. [12] This allows developers to reason about the effects their code may have, enabling safer and more predictable programming patterns.

H. Program Specifications

Verifiable type annotations and function signatures express preconditions, postconditions, and invariants are used to reason about program behavior statically and formally to ensure runtime safety. This allows developers to write more reliable code and reduce the likelihood of logical bugs.

I. The Numeric Lattice and Subtyping

Lento supports a subtyping mechanism that allows for hierarchical relationships between types. This enables polymorphism and code reuse, as functions can accept arguments of a supertype while operating on subtypes.

Unlike standard systems where numeric types are disjoint, Lento organizes scalars into a bounded lattice structure $\mathcal{L} = (B, \subseteq)$. Subtyping is defined as the conjunction of structural lattice inclusion and logical implication rather than implicit coercion, simplifying the elaboration phase.

$$\frac{B_1 \subseteq B_2 \quad \text{Valid}_{\text{SMT}}(\varphi_1 \Rightarrow \varphi_2)}{\Gamma \vdash \{\nu : B_1 \mid \varphi_1\} <: \{\nu : B_2 \mid \varphi_2\}} \text{Sub-Refine} \quad (8)$$

Fig. 2. Refinement Subtyping Rule. A type is a subtype if its base is smaller in the lattice AND its logical predicate implies the parent's predicate.

This allows strictly checked narrow types (e.g., **u8**) to be used where wider types (e.g., **u64**) are expected, but not vice versa, without runtime casting.

See Fig. 3 for the number subtyping lattice in Lento illustrating primitive widening type conversions (numeric promotion) and subtyping relationships. [13]

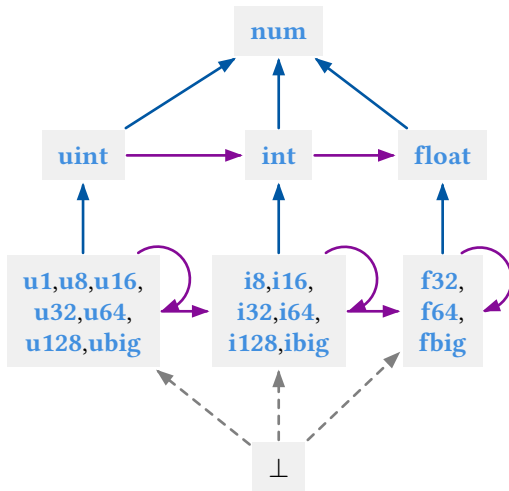


Fig. 3. Numeric type lattice with subtyping in blue and widening in purple.

J. Units of Measure

Units of measure refine numeric types with a phantom unit index and are checked statically, then erased. We write $\text{num}[u]$ for numbers with unit expression u ; the dimensionless case is $u = 1$.

Unit expressions follow the surface grammar from Section IV.B and denote elements of a commutative group of base units with integer exponents. We write $\text{nf}(u)$ for the normalized map $\text{uid} \rightarrow \mathbb{Z}$ obtained by flattening products/divisions and collecting exponents. Two units are equal iff they normalize to the same map:

$$u_1 = u_2 \iff \text{nf}(u_1) = \text{nf}(u_2) \quad (9)$$

The typing rules for arithmetic enforce dimensional consistency:

$$\frac{}{\Gamma \vdash n u : \text{num}[u]} \text{T-UnitLit} \quad (10)$$

$$\frac{\Gamma \vdash e_1 : \text{num}[u] \quad \Gamma \vdash e_2 : \text{num}[u]}{\Gamma \vdash e_1 + e_2 : \text{num}[u]} \text{T-Add} \quad (11)$$

$$\frac{\Gamma \vdash e_1 : \text{num}[u] \quad \Gamma \vdash e_2 : \text{num}[u]}{\Gamma \vdash e_1 - e_2 : \text{num}[u]} \text{T-Sub} \quad (12)$$

$$\frac{\Gamma \vdash e_1 : \text{num}[u_1] \quad \Gamma \vdash e_2 : \text{num}[u_2]}{\Gamma \vdash e_1 * e_2 : \text{num}[u_1 * u_2]} \text{T-Mul} \quad (13)$$

$$\frac{\Gamma \vdash e_1 : \text{num}[u_1] \quad \Gamma \vdash e_2 : \text{num}[u_2]}{\Gamma \vdash \frac{e_1}{e_2} : \text{num}\left[\frac{u_1}{u_2}\right]} \text{T-Div} \quad (14)$$

$$\frac{\Gamma \vdash e : \text{num}[u] \quad z \in \mathbb{Z}}{\Gamma \vdash e^z : \text{num}[u^z]} \text{T-Pow} \quad (15)$$

The same-unit side condition also applies to comparisons and equality on numeric quantities (e.g., $e_1 < e_2$ requires $u_1 = u_2$).

Numeric suffix constructors such as $2i$ for complex injection are handled by elaboration into ordinary core terms (e.g., function/constructor application); after elaboration they type-check using the existing rules for applications and overloaded operators.

K. Bidirectional Type Inference

Lento employs a bidirectional type inference algorithm **num** that combines both type checking and type synthesis. This approach allows the compiler to infer types of expressions without requiring explicit type annotations.

To support local type inference without global unification (which is undecidable with dependent refinements), Lento utilizes a bidirectional discipline. We split the typing judgment into two modes: Synthesis (\Rightarrow) and Checking (\Leftarrow).

a) *Synthesis* (\Rightarrow): In synthesis mode, the compiler generates a type τ for an expression e based on the environment Γ .

$$\frac{x : \tau \in \Gamma}{\Gamma \vdash x \Rightarrow \{\nu : \tau \mid \nu = x\}} \text{Var} \quad (16)$$

b) *Checking* (\Leftarrow): In checking mode, the compiler verifies expression e against a known type τ . Crucially, this allows for “smart” literals that adapt to their expected width.

$$\frac{\Gamma \vdash e \Rightarrow \tau' \quad \Gamma \vdash \tau' <: \tau}{\Gamma \vdash e \leq \tau} \text{Subsumption} \quad (17)$$

L. Operational Semantics

Lento uses a call-by-value, big-step operational semantics. The core reduction rule for let-binding illustrates the flow-sensitive nature of the environment update.

$$\frac{\Gamma \vdash e_1 \uparrow v_1 \quad \Gamma[x \mapsto v_1] \vdash e_2 \Downarrow v_2}{\Gamma \vdash \text{let } x = e_1 \text{ in } e_2 \Downarrow v_2} \text{E-Let} \quad (18)$$

VI. OPERATIONAL SEMANTICS

This section defines the small-step operational semantics of Lento using a standard evaluation relation $e \rightarrow e'$ over expressions.

VII. COMPILATION

To effectively translate Lento’s high-level abstractions and rich type system into efficient executable code, a sophisticated multistage compilation process is required.

A. Overview

Including staging, elaboration, lowering, optimization, code generation, and profiling with feedback for re-optimization.

Interpretation:

B. Objects (nodes)

- S : Source program
- S^{ct} : Staged/expanded source (after hygienic macros / MSP / const-eval)
- A : AST / syntax tree
- K : “Semantic IR” typed core calculus (post-elaboration, bidirectional typing, holes resolved/registered)
- I : Mid-level IR (e.g., SSA-ish or semantic dialect)
- I^* : Optimized IR
- E : Final artifact (object file / executable / WASM module)

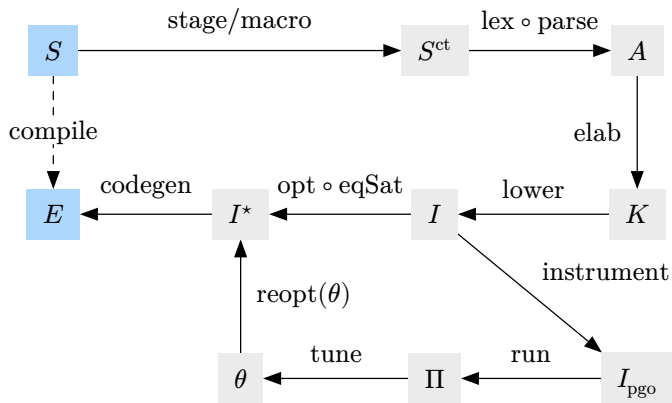


Fig. 4. Internal compilation pipeline with stages from source code to final artifact highlighted in blue

- $K_{\text{spec}}, I_{\text{spec}}$: Specialized (partial-evaluated) variants
- I_{pgo} : Instrumented IR
- Π : Profiling data
- θ : Optimization parameters / decisions derived from Π

C. Morphisms (edges)

- stage/macro: Hygienic macros, MSP phase splitting, const-eval fragment
- lex o parse: Lexer+parser
- elab: Elaboration/type inference (bidirectional) into core, including numeric-suffix resolution in a dedicated suffix namespace (producing ordinary core terms and marking unit metadata for erasure before lower)
- lower: Lowering into IR
- opt: Standard optimizations + fusion + specialization triggers + etc.
- eqSat/rewrite: Equality saturation / rewrite-rule engine for library-driven optimization
- partialEval: Specialization w.r.t. known inputs
- instrument, run, tune: PGO loop with a “pause” awaiting external execution input
- codegen: Code generation to target architecture (e.g., x86-64, WASM, etc.)

Linking with external libraries and runtime is a separate concern that may occur at the IR or final artifact stage depending on the design of the runtime and FFI.

REFERENCES

- [1] A. J. Kennedy, “Programming languages and dimensions,” technical report UCAM-CL-TR-391, Apr. 1996. doi: 10.48456/tr-391.
- [2] A. Kennedy, “Dimension types,” *Programming Languages and Systems – ESOP ’94*, vol. 788. in Lecture Notes in Computer Science, vol. 788. Springer, pp. 348–362, 1994. doi: 10.1007/3-540-57880-3_23.
- [3] A. Kennedy, “Types for Units-of-Measure: Theory and Practice.” [Online]. Available: http://typesatwork.imm.dtu.dk/material/TaW_Paper_TypesAtWork_Kennedy.pdf
- [4] R. Atkey, N. Ghani, F. Nordvall Forsberg, T. Revell, and S. Staton, “Models for Polymorphism over Physical Dimension,” in *13th International Conference on Typed Lambda Calculi and Applications (TLCA 2015)*, in Leibniz International Proceedings in Informatics (LIPIcs), vol. 38. Schloss Dagstuhl – Leibniz-Zentrum f’ur Informatik, 2015, pp. 45–59. doi: 10.4230/LIPIcs.TLCA.2015.45.
- [5] S. Westrick, M. Fluet, M. Rainey, and U. A. Acar, “Automatic Parallelism Management,” *Proc. ACM Program. Lang.*, vol. 8, no. POPL, Jan. 2024, doi: 10.1145/3632880.
- [6] S. P. Midkiff, *Automatic Parallelization*. in Synthesis Lectures on Computer Architecture. Cham: Springer International Publishing, 2022, doi: 10.1007/978-3-031-01736-0.
- [7] N. Polikarpova, I. Kuraj, and A. Solar-Lezama, “Program synthesis from polymorphic refinement types,” in *Proceedings of the 37th ACM SIGPLAN Conference on Programming Language Design and Implementation*, in PLDI ’16. Santa Barbara, CA, USA: Association for Computing Machinery, 2016, pp. 522–538. doi: 10.1145/2908080.2908093.
- [8] P. M. Rondon, M. Kawaguci, and R. Jhala, “Liquid types,” in *Proceedings of the 29th ACM SIGPLAN Conference on Programming Language Design and Implementation*, in PLDI ’08. Tucson, AZ, USA: Association for Computing Machinery, 2008, pp. 159–169. doi: 10.1145/1375581.1375602.
- [9] H. Xi and F. Pfenning, *Advanced Types in Programming Languages*. in Titolo collana. Cambridge University Press, 2016. [Online]. Available: <https://dl.acm.org/doi/book/10.5555/1076265>
- [10] B. Pierce, *Types and Programming Languages*. in MIT Press Books. MIT Press, 2002. [Online]. Available: <https://dl.acm.org/doi/book/10.5555/509043>

- [11] M. Woo-Kawaguchi, "High-Level Liquid Types," Doctoral dissertation, 2016. [Online]. Available: <https://escholarship.org/uc/item/7d8525sz>
- [12] M. Kawaguchi, P. Rondon, A. Bakst, and R. Jhala, "Deterministic parallelism via liquid effects," in *Proceedings of the 33rd ACM SIGPLAN Conference on Programming Language Design and Implementation*, in PLDI '12. Beijing, China: Association for Computing Machinery, 2012, pp. 45–54. doi: 10.1145/2254064.2254071.
- [13] F. Pfenning, "Lecture Notes on Subtyping," 2023, [Online]. Available: <https://www.cs.cmu.edu/~fp/courses/15836-f23//lectures/08-subtyping.pdf>