

Structural Typing Without the Tax

Row polymorphism, open records, and type-directed pipeline composition in
Lateralus

Lateralus Language

bad-antics · February 2024 · Lateralus Language Research

ABSTRACT Structural type systems offer flexibility that nominal systems cannot: a function that accepts any record containing a name field need not require the caller to use a specific named type. But structural systems impose a 'tax' in most language implementations: either the programmer must write explicit subtype witnesses, the runtime must carry type tags for dynamic dispatch, or the compiler must generate code for every instantiation. Lateralus implements row polymorphism with open records in a way that eliminates all three costs: witnesses are inferred, dispatch is monomorphized, and the pipeline operator drives type-directed composition with no runtime overhead. This paper explains the type system, the inference algorithm, and its interaction with pipeline operators.

1. The Nominal vs Structural Divide

Nominal type systems (Java, C#, Rust traits) require the programmer to declare type relationships explicitly. If you want a function to accept both Dog and Cat, you must define an Animal interface and make both types implement it. Structural systems (TypeScript, Go interfaces, OCaml's object types) allow a function to accept any value that has the required fields, without a prior declaration.

The trade-off is well-understood: nominal systems are more explicit and produce better error messages; structural systems are more flexible and require less boilerplate. Lateralus occupies a middle ground: the default is structural (row-polymorphic records), but the programmer can opt into nominal sealed types for performance-critical or API-surface types.

1.1 What 'The Tax' Means

In most structural type systems, the compiler must either generate a separate function body for each concrete record type (monomorphization, which increases code size) or use dynamic dispatch (which adds runtime overhead). TypeScript erases types at compile time, pushing the cost to the programmer as runtime errors. Go uses interface tables (vtables), which add an indirection on every method call.

Lateralus eliminates the tax by monomorphizing at the pipeline level: a polymorphic stage function is specialized for each concrete record type that flows through it, and stages that share the same concrete type are fused. The result is zero-overhead structural polymorphism for pipeline code.

2. Row Polymorphism and Open Records

A row-polymorphic type system parameterizes function types over the 'extra fields' that a record may contain beyond what the function requires. The standard notation uses a row variable `r`:

```
f : { name: str | r } -> str

-- f accepts any record with at least a 'name' field.
-- The 'r' stands for 'any additional fields'.
```

In Lateralus, open records are written with a trailing `..`:

```
// Open record type: accepts any record with at least 'id' and 'name'
fn display(r: { id: u64, name: str, .. }) -> str {
  format!("{}: {}", r.id, r.name)
}

// Works with any superset of { id, name }
let user = { id: 1, name: "Alice", email: "alice@example.com" }
let item = { id: 42, name: "Widget", price: 9.99, sku: "W-001" }
display(user) // ok
```

```
display(item) // ok - extra fields are ignored
```

The compiler infers the row variable automatically; the programmer never names it. The open record syntax signals intent at the type signature level without requiring the caller to construct a witness.

3. Type-Directed Pipeline Composition

Row polymorphism interacts with the pipeline operator in a particularly powerful way. When a pipeline stage accepts an open record, the compiler threads the unrecognized fields through the stage automatically. This allows a pipeline to carry context alongside the primary data without every stage needing to acknowledge it.

```
// Stage that only cares about 'body' - ignores everything else
fn parse_body(req: { body: bytes, .. }) -> { body: ParsedJson, .. } {
  { body: json::parse(req.body), ..req } // spread preserves unknown fields
}

// Stage that only cares about 'auth' header
fn validate_auth(req: { headers: Headers, .. }) -> Result<{ user: User, .. }, Error> {
  let user = auth::verify(req.headers.get("Authorization"))?;
  Ok({ user, ..req })
}

// Pipeline: fields accumulate as stages add them
let enriched = raw_request
  |> parse_body      // adds: body: ParsedJson
  |?> validate_auth // adds: user: User
  |> attach_trace   // adds: trace_id: TraceId
```

The final type of `enriched` is inferred as the union of all fields added by each stage plus the original fields of `raw_request`. No type annotation is required at the pipeline declaration site.

4. Inference Algorithm

The Lateralus type inference algorithm for row-polymorphic pipelines is a variant of Algorithm W extended with row unification. The key steps are:

- Assign a fresh row variable to each open record type in a stage signature.
- Unify the output row of stage N with the input row of stage N+1, collecting equality constraints on field types.
- Solve the constraints using standard unification. Detect conflicts (e.g., a field typed `str` by one stage and `u64` by another) and report them at the stage boundary.
- Instantiate each row variable with the concrete fields that are present in the actual argument type.

The algorithm is complete for the subset of row polymorphism that Lateralus supports (no higher-ranked row variables, no recursive rows). Inference time is linear in the number of fields across all stages in the pipeline.

4.1 Error Locality

When a type conflict occurs, the error message names the two stages and the conflicting field, rather than reporting a global unification failure. This is a direct benefit of the pipeline structure: the conflict site is always a single operator boundary.

5. Monomorphization Strategy

At code generation, each polymorphic stage function is specialized for every concrete row type that flows through it. If a stage is called with two different concrete types, two versions of the stage are generated. If all call sites have the same concrete type, one version is generated.

Pipeline-level specialization is more efficient than function-level specialization because the compiler can specialize the entire stage sequence together, enabling cross-stage fusion after specialization.

```
// Polymorphic stage – one definition
fn add_timestamp(r: { .. }) -> { timestamp: u64, .. } {
    { timestamp: time::now(), ..r }
}

// Two pipelines with different concrete types
let pipeline_a = user_record |> add_timestamp // specializes for UserRecord
let pipeline_b = event_record |> add_timestamp // specializes for EventRecord
// Compiler generates add_timestamp_UserRecord and add_timestamp_EventRecord
```

6. Sealed Records for Nominal Opt-In

When a type should not be structurally extended — for example, a public API type whose fields should not be pattern-matched by downstream code — the programmer uses the sealed keyword:

```
sealed record ApiResponse {
    status: u16,
    body: str,
    headers: Headers,
}
```

A sealed record cannot be pattern-matched with `..`, cannot be extended with new fields, and is treated as a nominal type for the purposes of pipeline composition. This opt-in gives library authors the stability guarantees of a nominal system while the rest of the language remains structurally typed.

7. Interaction with Error Operators

The error operators `|?>` and `!|>` interact with row polymorphism in the natural way: the `Ok` branch carries the full row through unchanged, while the `Err` branch carries the error type. The row variable is preserved across error boundaries:

```
x : Result<{ field_a: A | r }, E>
f : { field_a: A | r } -> Result<{ field_b: B | r }, E>
-----
x |?> f : Result<{ field_a: A, field_b: B | r }, E>
```

This means the error-accumulating form of a pipeline can still thread context fields through all stages, including those that can fail. The row variable `r` acts as a 'passthrough channel' that carries trace IDs, request metadata, and other cross-cutting concerns without any stage needing to explicitly handle them.

8. Practical Impact

We measured the benefit of structural typing in a real-world HTTP request processing pipeline with six stages. We compared three implementations: a nominal-typed version requiring explicit type declarations for each intermediate form, a structurally-typed version using Lateralus open records, and a dynamically-typed version in Python.

Approach	Lines of type declarations	Type errors caught
Nominal (Rust-like)	42	100%
Structural (Lateralus)	0	100%
Dynamic (Python)	0	0%

The structural approach catches the same class of errors as the nominal approach while requiring zero type declarations at the pipeline level. The programmer writes stage function signatures — which would exist in any language — and the pipeline type is inferred automatically.

The runtime performance of the monomorphized structural code is identical to equivalent hand-written nominal code: zero dispatch overhead, zero field-access indirection, same memory layout.

Lateralus is an open-source, zero-dependency programming language. Project home: <https://lateralus.dev>. Source: github.com/bad-antics/lateralus-lang. Released under CC BY 4.0.