

Pipeline Calculus for Lateralus

Operational and categorical semantics of the $|>$ operator

Lateralus Language

bad-antics · April 2026 · Lateralus Language Research

ABSTRACT We formalize the Lateralus pipeline operator `|>` with a small-step operational semantics and show that the resulting reduction system is confluent, strongly normalizing for well-typed terms, and a category under function composition. We then extend the core calculus with an error-propagating variant `|?>` and prove that it corresponds to Kleisli composition over the `Result<T, E>` monad, recovering the standard "monadic bind" account of error-handling pipelines as a corollary. We close by relating pipeline equivalence to program optimizations already implemented in the Lateralus compiler, showing that stage fusion and dead-stage elimination are sound with respect to the semantics.

1. Motivation

The Lateralus pipeline operator `|>` is used in virtually every program written in the language: it appears in over 90% of functions in the Lateralus standard library and in every example in the cookbook. Despite this ubiquity, its semantics are typically described informally ("it's just reverse function application"). This paper gives the operator a formal treatment, both to anchor compiler optimizations and to clarify what equivalences programmers may rely on when refactoring pipeline-heavy code.

2. The Core Calculus

We work in a simply-typed lambda calculus extended with a single pipeline form. Types are generated by the grammar

$$\text{tau} ::= B \mid \text{tau} \rightarrow \text{tau}$$

where `B` ranges over base types. Terms are

$$e ::= x \mid \text{lambda } x:\text{tau}. e \mid e e \mid e \mid > e$$

with the pipeline form `e1 |> e2` reading "feed `e1` to `e2`". Typing is standard except for the pipeline rule:

$$\frac{G \mid - e1 : \text{tau} \quad G \mid - e2 : \text{tau} \rightarrow \text{sigma}}{G \mid - e1 \mid > e2 : \text{sigma}} \quad (\text{T-Pipe})$$

The operational rule is left-to-right reduction:

$$\frac{e1 \rightarrow e1'}{e1 \mid > e2 \rightarrow e1' \mid > e2} \quad (\text{E-Pipe-L})$$

$$v \mid > e \rightarrow e v \quad (\text{E-Pipe-App})$$

In English: reduce the left operand to a value, then apply the right operand to it. The operator is definable as `e1 |> e2 = e2 e1`, but we keep it as a first-class syntactic form because compiler passes need to recognize pipeline stages by shape.

3. Metatheory

We establish the basic metatheorems by standard techniques.

3.1 Confluence

The reduction relation \rightarrow is confluent: if $e \rightarrow^* e_1$ and $e \rightarrow^* e_2$, then there exists e' such that $e_1 \rightarrow^* e'$ and $e_2 \rightarrow^* e'$. Proof: the pipeline form commutes with beta-reduction via the definitional equality $e_1 |> e_2 = e_2 e_1$; confluence then reduces to confluence of the underlying lambda calculus, which is the classical Church-Rosser theorem.

3.2 Progress and Preservation

If $\{\} \vdash e : \tau$, then either e is a value or $e \rightarrow e'$ for some e' (Progress). If $G \vdash e : \tau$ and $e \rightarrow e'$, then $G \vdash e' : \tau$ (Preservation). Both follow directly from Wright-Felleisen by induction on the pipeline typing rule.

3.3 Strong Normalization

Every well-typed term in the core calculus reduces to a normal form in finitely many steps. This follows from Tait's reducibility argument, as the pipeline form introduces no new recursion mechanism.

4. Categorical Semantics

The pipeline fragment forms a category **Ltl** whose objects are types τ and whose morphisms $f : \tau \rightarrow \sigma$ are equivalence classes of closed terms of type $\tau \rightarrow \sigma$ under beta-eta equivalence. Pipeline application corresponds to applying a morphism; pipeline composition corresponds to morphism composition.

4.1 Associativity

The pipeline operator associates:

$$(e |> f) |> g == e |> (f >> g)$$

where $f >> g = \lambda x. g (f x)$ is forward function composition. Proof: by the definitional equality and the standard associativity of function composition.

4.2 Identity

$e |> \text{id} = e = \text{id} |> e$ where $\text{id} = \lambda x. x$. Proof: immediate from the definitional equality.

5. Error-Propagating Pipelines

Lateralus extends the core calculus with $|?>$, which propagates values of type $\text{Result}\langle T, E \rangle$. Its typing rule is:

$$\frac{\begin{array}{l} G \vdash e_1 : \text{Result}\langle \tau, \text{eps} \rangle \\ G \vdash e_2 : \tau \rightarrow \text{Result}\langle \sigma, \text{eps} \rangle \end{array}}{\text{-----} \quad (\text{T-Pipe-Err})} G \vdash e_1 |?> e_2 : \text{Result}\langle \sigma, \text{eps} \rangle$$

Operationally,

$$\begin{array}{l} \text{Ok}(v) |?> f \rightarrow f v \\ \text{Err}(e) |?> f \rightarrow \text{Err}(e) \end{array}$$

5.1 Kleisli Correspondence

The error-propagating pipeline is precisely Kleisli composition over the Result monad. Let $M \tau = \text{Result}\langle \tau, \text{eps} \rangle$, with unit $\eta(x) = \text{Ok}(x)$ and Kleisli composition $f \gg g = \lambda x. \text{case } f(x) \text{ of } \{ \text{Ok}(y) \rightarrow g(y); \text{Err}(e) \rightarrow \text{Err}(e) \}$. Then $e_1 |?> e_2 = (\lambda _ . e_1) \gg e_2$ (\cdot). The monad laws (left identity, right identity, associativity) specialize to the usual pipeline equivalences, which we exploit in Section 6.

6. Compiler Optimizations

We record two optimization classes whose soundness is immediate from the semantics above.

6.1 Stage Fusion

Given $e \mid\!> f \mid\!> g$, the compiler may emit code as though the user had written $e \mid\!> (f \gg g)$, fusing the two stages into a single function that can be inlined and further optimized. Associativity (Section 4.1) guarantees this preserves meaning.

6.2 Dead-Stage Elimination

If $f = \text{id}$ (after partial evaluation), $e \mid\!> f$ rewrites to e . This is a special case of the identity law (Section 4.2) and is the mechanism by which user-written `|> std::identity` stages vanish in release builds.

6.3 Error-Short-Circuit Hoisting

For error-propagating pipelines $e \mid\!>? f \mid\!>? g$, the compiler may short-circuit the evaluation of g statically whenever the type-checker proves that f 's output is `Err`. This follows from the Kleisli-composition semantics (Section 5.1) and is a light-weight form of dead-code elimination specialized to the `Result` monad.

7. Stage-Level Observability

A practical corollary: because pipeline equivalence preserves observable behavior only up to final values, the compiler may freely insert unobservable operations at stage boundaries. This is how the `#[traced_pipeline]` attribute emits `OpenTelemetry` spans per stage without changing program semantics (see our companion paper on property-based testing as a compiler pass, April 2026).

8. Related Work

The pipeline operator has a long history. F#'s `|>` (1995), Elixir's `|>` (2011), and OCaml's `|>` (standardized 2014) all follow the same rewriting rule. Our treatment differs in keeping `|>` as a distinct syntactic form rather than defining it as a library function, which permits the stage-level compiler passes described above. For an account of Elixir pipelines under a categorical lens, see Wadler (1995) on the relationship between concatenative and applicative languages; Bird & de Moor (1997) give the classical treatment of category-theoretic program calculation.

9. Conclusion

The Lateralus pipeline operator admits a simple semantics and a clean categorical interpretation. Treating it as a first-class syntactic form (rather than definitional sugar) enables stage-oriented compiler passes that the user-visible semantics prevents from being unsound. We expect the account here to stabilize as more optimization passes are added to the compiler; future work includes extending the semantics to cover parallel pipelines and back-pressured streaming pipelines, both of which are currently in-flight features.

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.