

DuckPL: A Procedural Language in DuckDB

Bringing PL/pgSQL to DuckDB

30th of January 2026  DuckDB Developer Meeting #1

Denis Hirn

 denis.hirn@uni-tuebingen.de

University of Tübingen

 kryonix

Denis Hirn? kryonix?

Long-time DuckDB contributor since early 2020:

- Recursive/Materialized CTEs
 - Some re-architecting of query decorrelation
 - Various bug fixes and optimizations
-

Research focus:

- Database systems
 - ▶ query optimization,
 - ▶ execution engine design, and
 - ▶ User-Defined Function optimization, aka. **how to get rid of them**
- Compilers and programming languages
- Bridging the gap between both fields

State of DuckDB UDFs

Supported ✓

```
1 | CREATE MACRO add(a, b) AS a + b;
```

Not Supported ✗

```
1 | CREATE MACRO sequence(n) AS
2 |   IF n < 0 THEN
3 |     do some stuff
4 |   ELSE
5 |     do other stuff
6 |   END IF;
7 |   do some more
8 | CREATE TABLE something AS (column type);
9 | RETURN something;
```

MACROs:

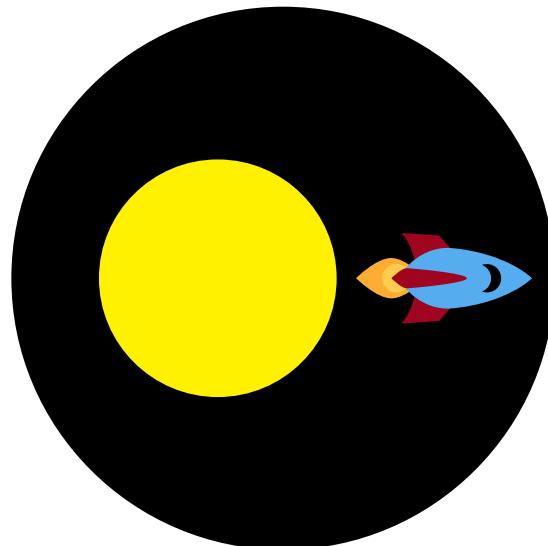
- simple textual replacements
- no full-fledged procedural language for writing user-defined functions (UDFs) yet

Python, R, and other languages can be used to write UDFs, but:

- Require an external runtime
- Breaks the “Single-File, Zero-Dependency Database” promise

Introducing DuckPL!

Procedural PL/SQL, Native to DuckDB.



Procedural PL/SQL for DuckDB!

Procedural Logic (DuckPL)

```
1 CREATE FUNCTION collatz(y BIGINT) RETURNS BIGINT AS $$  
2 DECLARE  
3     steps BIGINT := 0;  
4     x BIGINT := y;  
5 BEGIN  
6     WHILE x > 1 LOOP  
7         IF x % 2 = 0 THEN  
8             x := x / 2;  
9         ELSE  
10            x := 3 * x + 1;  
11        END IF;  
12        steps := steps + 1;  
13    END LOOP;  
14    RETURN steps;  
15 END;  
16 $$;  
17 SELECT collatz(5);
```

 *Easy!*

Recursive CTE (Pure SQL)

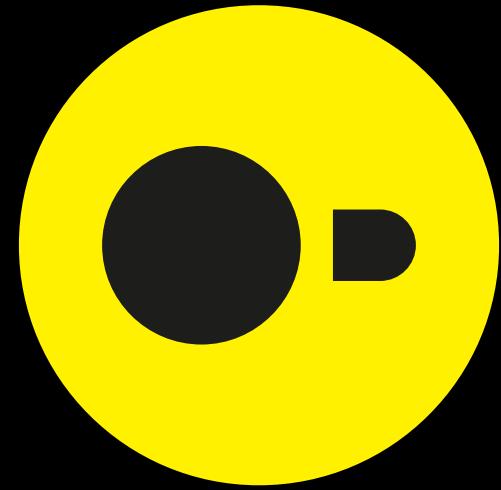
```
1 SELECT  
2     (WITH RECURSIVE collatz_cte(x, steps) AS (  
3         SELECT 5 AS x, 0 AS steps  
4         UNION ALL  
5         SELECT  
6             CASE WHEN x % 2 = 0  
7                 THEN x / 2  
8                 ELSE 3 * x + 1  
9             END AS x,  
10            steps + 1 AS steps  
11        FROM collatz_cte  
12        WHERE x > 1  
13    )  
14    SELECT steps  
15    FROM collatz_cte  
16    WHERE x = 1  
17 ) AS collatz;
```

 *Not so easy...*

*For average users

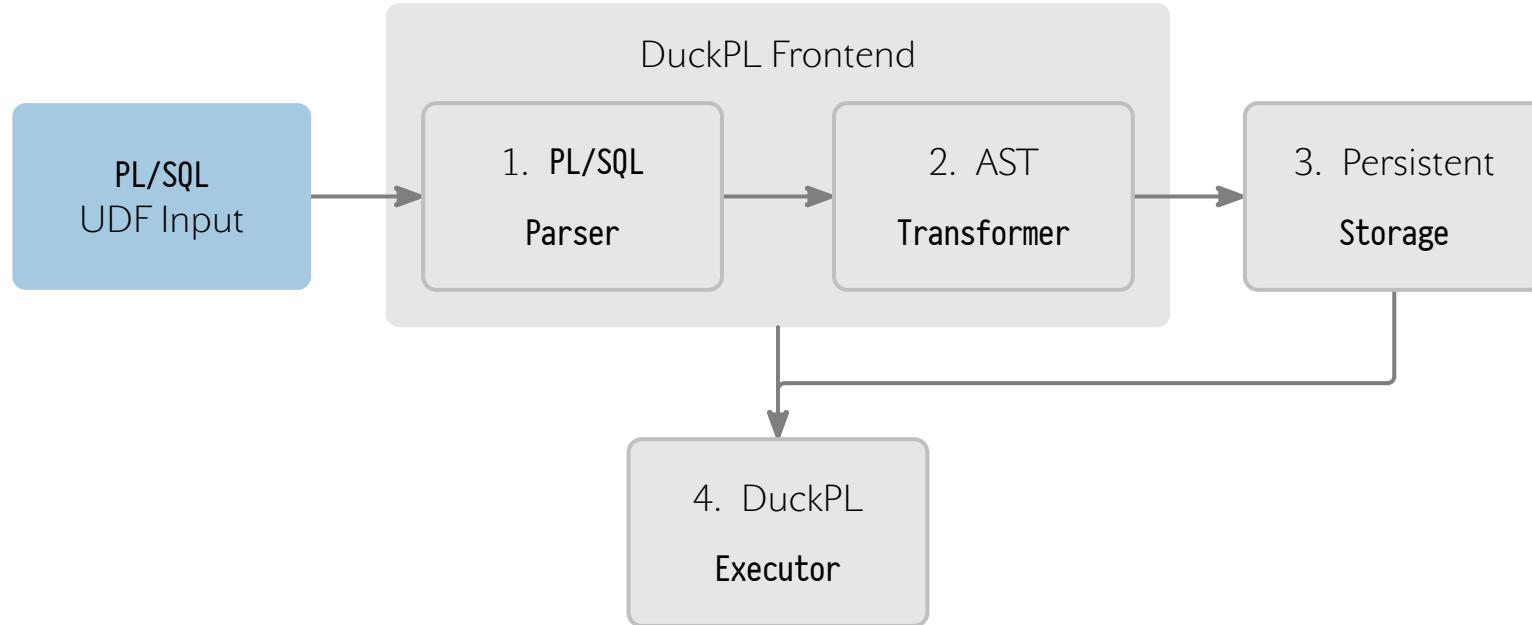
No external runtimes required Imperative programming right inside the system

PL/pgSQL compatibility Migration of existing codebases to DuckDB made *easy*



Demo Time!

Implementing DuckPL: Key Components & Architecture



Parser Extension Custom parser for `CREATE FUNCTION` statements

Operator Extension Custom `Bind` and `Plan` for `CREATE` operations

Parsing: The Missing Pieces

```
1 CREATE FUNCTION collatz(y BIGINT)
2 RETURNS BIGINT
3 AS $$
```

4 **DECLARE**

```
5   steps BIGINT := 0;
6   x BIGINT := y;
```

7 **BEGIN**

```
8   WHILE x > 1 LOOP
9     IF x % 2 = 0 THEN
10      x := x / 2;
11    ELSE
12      x := 3 * x + 1;
13    END IF;
14    steps := steps + 1;
15  END LOOP;
16  RETURN steps;
17 END;
18 $$;
```

1. **Incomplete CREATE FUNCTION parsing support:** DuckDB's SQL parser lacks grammar rules to parse `AS $$... $$` functions.
2. **No Language Parser:** The PL/pgSQL body is a generic string literal and must be parsed separately.



Consequence We need to parse **both** the `CREATE FUNCTION` statement **and** the PL/pgSQL function body.

PL/pgSQL Parsing: The `libpg_query` Approach

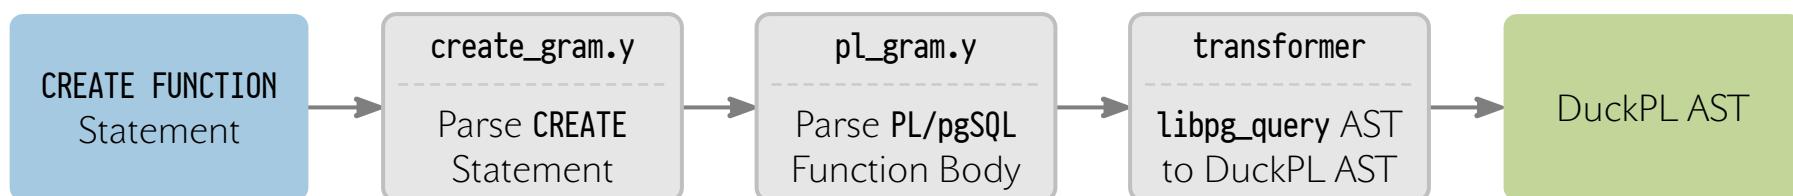
Phase 1: Parsing

- DuckPL needs **two** parsers:
 1. for the `CREATE FUNCTION` statement
 2. for the PL/pgSQL function body.
- We can reuse the existing `libpg_query` parser for both!
 - ▶ leverage missing pieces for `CREATE FUNCTION` parsing
 - ▶ reuse the existing PL/pgSQL parser as-is

Phase 2: AST Transformation

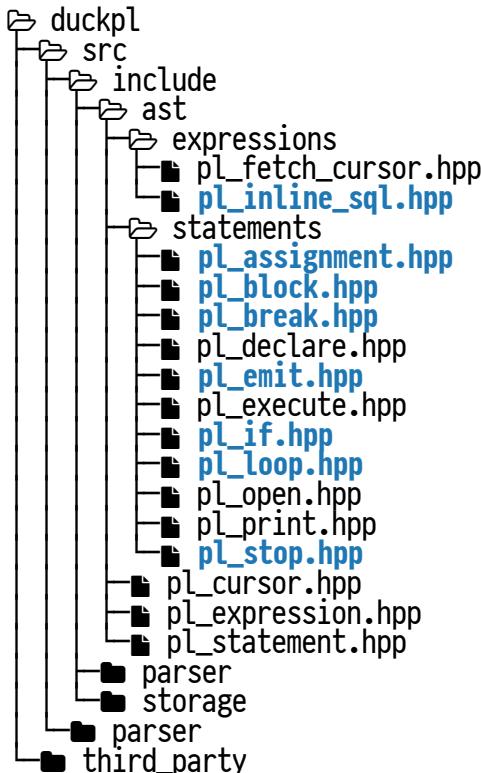
- Implement `transformer` from `libpg_query` AST to DuckPL AST

This mirrors **exactly** how DuckDB's original `SQL` parser was built in 2018!



Soon™: Rip this apart and use a PEG-based parser!

Interlude: DuckPL AST – The Universal Internal Representation



DuckPL IR: **minimalist** and **syntax-agnostic** to support multiple source languages.

- Complex constructs like **FOR/WHILE** loops desugar to **LOOP + IF + BREAK**
 - Eliminates **FOR**, **WHILE**, **CURSOR**, **ARRAY** loops, *etc.*
- No **CASE** statements: Everything simplifies to an **IF** statement.

Source Language (PL/pgSQL)

```
1 WHILE counter < 10 LOOP
2   counter := counter + 1;
3 END LOOP;
4 RETURN counter;
```

Internal DuckPL IR AST

```
1 loop {
2   if (counter >= 10) { break; }
3   let counter = counter + 1;
4 }
5 emit counter;
6 stop;
```



Simplifies Interpreter Reduces complexity of control flow handling

Future Language Support A new procedural language (PL/Python, PL/Duck) requires just **Transformer** \mapsto DuckPL AST

Simplifies Compilation Easier to compile DuckPL AST to SQL

Storage: Persisting and Registering DuckPL UDFs

Persistent Storage

- UDFs are stored in `duckpl_functions` table.
- AST is serialized and stored as a **BLOB**.
- Avoids unnecessary re-parsing.
- Similar to `DuckLake` macros.

Registration on Startup

- Load stored UDFs from `duckpl_functions`.
- Deserialize AST and register in the catalog.
- UDFs are immediately available without parsing.

```
1 | CREATE TABLE duckpl_functions (
2 |   function_id BIGINT PRIMARY KEY,
3 |   function_uuid UUID,
4 |   function_num_args INT,
5 |   function_arg_names TEXT[],
6 |   function_arg_types TEXT[],
7 |   function_return_types TEXT[],
8 |   function_returns_set BOOLEAN,
9 |   function_name TEXT NOT NULL,
10 |  function_src TEXT,
11 |  function_body BLOB);
```



Execution: The Stack-Driven Interpreter

Simple *tree-walk* interpreter, but avoids recursive calls by using an **explicit stack of frames**:

State Management Execution can be **paused** and **resumed** at any point

No C++ recursion No stack depth limits, no risk of stack overflows

```
1 ①→ LOOP {  
2   IF NOT x > 1 THEN  
3     BREAK;  
4   END IF;  
5   IF x % 2 = 0 THEN  
6     ...  
7   }  
8   ...
```

```
1 ①→ LOOP {  
2 ②→ IF NOT x > 1 THEN  
3     BREAK;  
4   END IF;  
5   RETURN NEXT x;  
6   ...  
7 }  
8 ...
```

```
1 ①→ LOOP {  
2   IF NOT x > 1 THEN  
3     BREAK;  
4   END IF;  
5 ③→ RETURN NEXT x;  
6   ...  
7 }  
8 ...
```

①→ PLLoop *	Execute
... Stack Frames ...	

②→ PLIf *	Execute
① PLLoop *	Resume

③→ PLEmit *	Execute
① PLLoop *	Resume

This design allows DuckPL to **stream** results efficiently without buffering everything in memory.

Execution: Streaming Output Like an Operator

```
1 CREATE FUNCTION infinite()
2 RETURNS SETOF BIGINT AS $$ 
3 DECLARE
4     i BIGINT := 0;
5 BEGIN
6     LOOP
7         i := (i + 1) % 1000;
8         RETURN NEXT i;
9     END LOOP;
10 END
11 $$;
12
13 -- Create 10 DataChunks:
14 SELECT *
15 FROM infinite()
16 LIMIT 10 * 2048;
```

PostgreSQL: Buffers **all results** before returning **anything**

- Will never return anything from `infinite()` function
- Leads to memory ballooning 
- Cannot be interrupted (e.g., via `LIMIT`)



DuckPL's interpreter is **Fully streaming**:

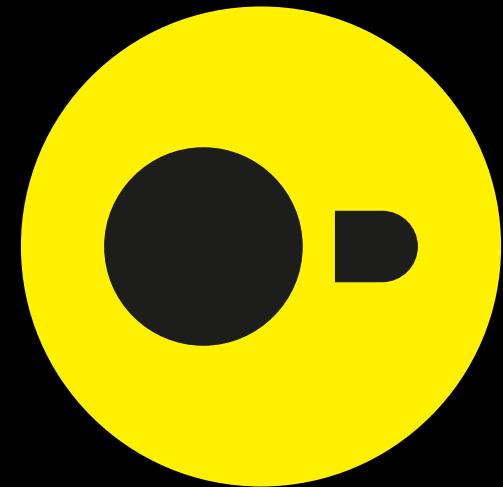
Memory Efficiency No unnecessary buffering of results

- `RETURN NEXT` statements write result into output chunk

Interruptibility When chunk is full, return `OperatorResultType::HAVE_MORE_OUTPUT`

- Interpreter **pauses** and **resumes** when requested:
Enabled by explicit stack of frames design
- Allows **early termination** (e.g., via `LIMIT`)

Architectural Fit Follows same idea as physical operators in DuckDB 



Execution: Expression Fast-Path

```
1 CREATE FUNCTION collatz(y BIGINT)
2 RETURNS BIGINT AS $$ 
3 DECLARE
4     steps BIGINT := 0;
5     x BIGINT := y;
6 BEGIN
7     WHILE x > 1 LOOP
8         IF x % 2 = 0 THEN
9             x := x / 2;
10        ELSE
11            x := 3 * x + 1;
12        END IF;
13        steps := steps + 1;
14    END LOOP;
15    RETURN steps;
16 END;
17 $$;
18 SELECT collatz(5);
```

The Slow Way:

Method Wrapping every expression in a **SELECT <...>** statement

Bottleneck Triggers the **Full SQL Pipeline** (Binding, Optimization, Execution) for *every single expression* 😱

This becomes *super slow* 🚫 without optimization!

The Fast Path:

💡 Use **ExpressionExecutor** for simple expressions

Not supported.. But how we do it anyway 😊:

1. Prepare a dummy **SELECT x > 1** statement
2. Extract the expression **x > 1** from prepared statement
3. Cache an **ExpressionExecutor** instantiated with **x > 1**
4. Execute against a **DataChunk** containing local variables

Result: 🎉 It's *fast* 🚦 now (we've seen speedups of 30x).

* This is vastly simplified; We have to do a lot more work to prepare the expression properly to make it cacheable.

DuckPL Feature Support

Supported

- Scalar/Table-valued UDFs
- Variables and Assignments
- All Data Types
- Composite Types types like `lineitem`
- Control Flow
 - `IF`
 - `LOOP`, `WHILE`, and `FOR` loops
 - `BREAK` and `CONTINUE`
 - `RETURN` and `RETURN NEXT`
- **Cursors (**`FETCH INTO`**)**
- Debugging (`RAISE INFO`)

Planned

- Aggregate/Window UDFs
- Exception handling
- Transactions `COMMIT`, `ROLLBACK`
- **UDF Optimizer**

Compilation to Pure SQL

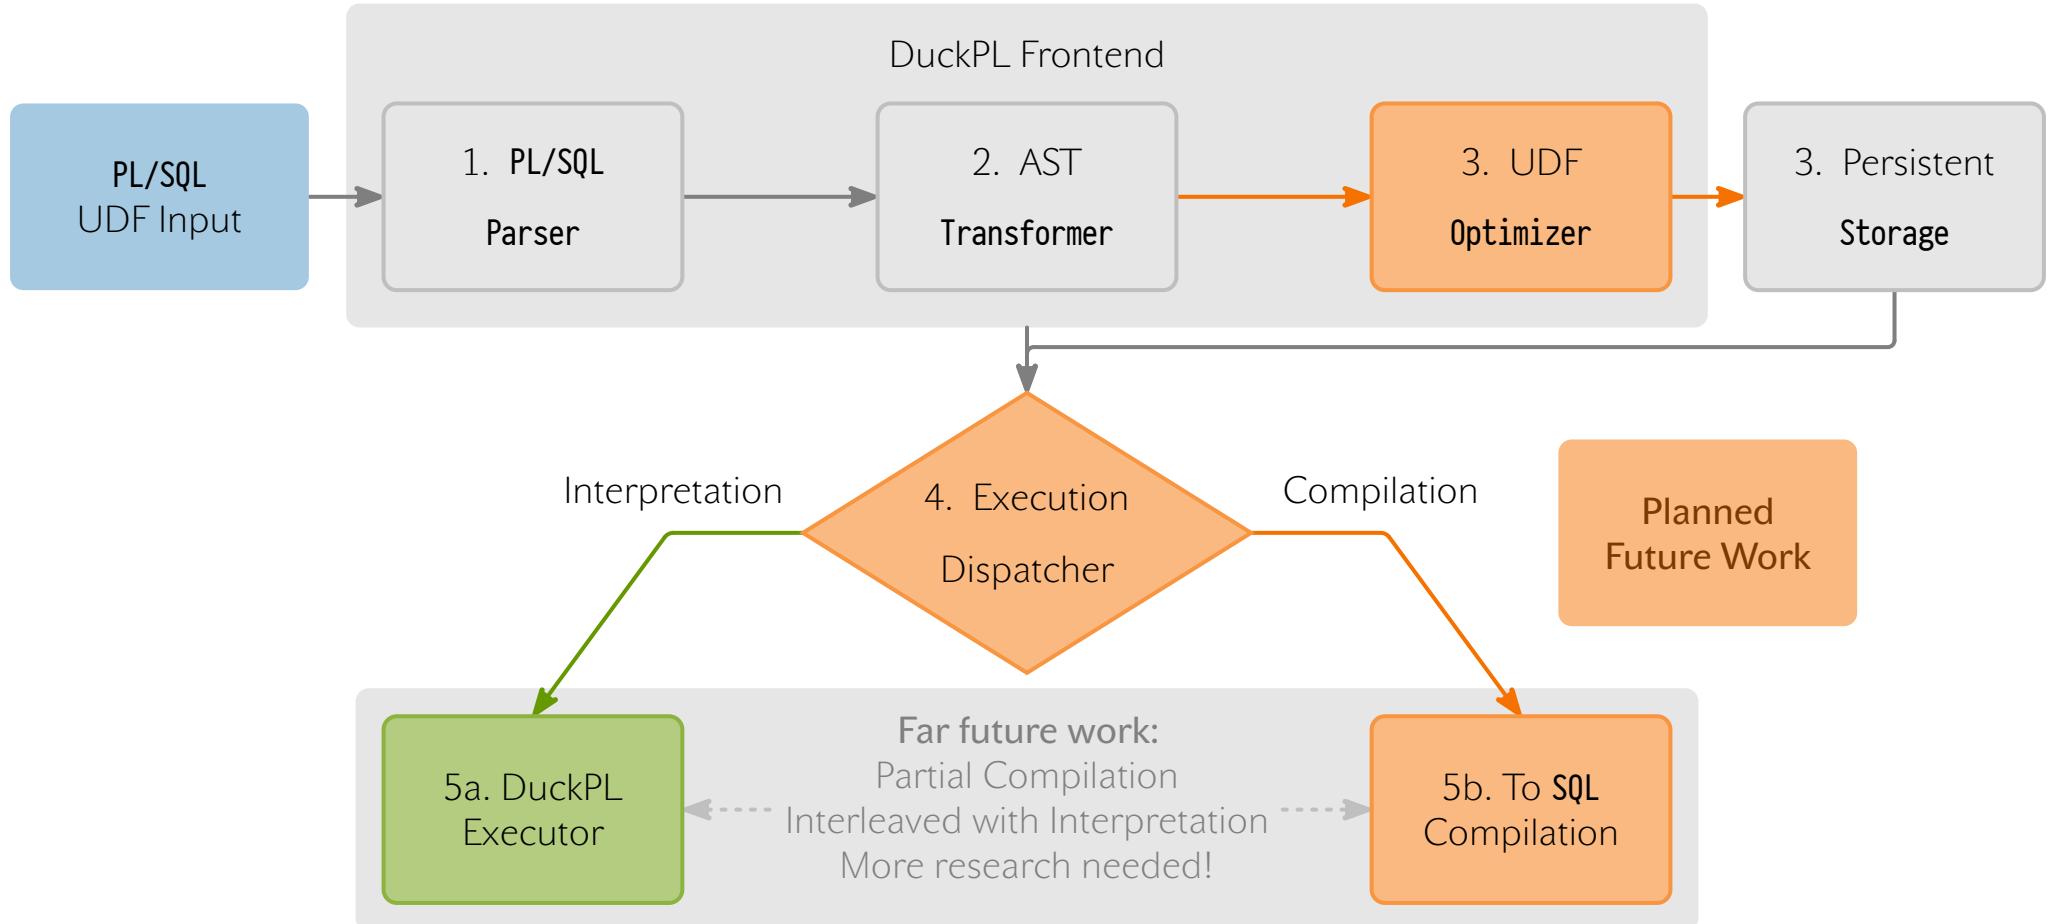
- Massively improve performance
- Leverage DuckDB's execution engine
- Use `WITH RECURSIVE` for complex control flow

Not for now or never *

- Dynamic SQL (Use `query(...)`)
- Advanced cursor features like `SCROLL` and `MOVE`
- Triggers

* PRs welcome!  As soon as DuckPL becomes open source.

DuckPL Architecture: The Path to Hybrid Execution



Future Work – The Vision for DuckPL

Interactive Mode Allow DuckPL statements directly in the **CLI**, for a **REPL-like** experience:

```
> duckdb
D LET y = 0 :: BIGINT;
D FOR i IN 1..10:
    LET x = (SELECT RANDOM());
    IF x > 0.5:
        LET y = y + 1;
D PRINT y;
5
D |
```

Modern Syntax Follow **friendly SQL** idea for PL syntax

- Add a secondary, lightweight syntax
- Add **PL/Python** frontend

Next-Gen Parser Move to PEG based **PL/pegSQL** for better DuckDB integration

Compiled Execution Integrate our **UDF compilation** research into DuckPL

Production Readiness Improve **error messages** and **debugging support**

Vectorized Interpretation Implement **vectorized interpretation**

Advanced Features Support **table-valued variables**

There is **so much** more to do!

Conclusion

Compatibility First

Bring PL/pgSQL functionality to the DuckDB ecosystem.

The Win:

- Compatibility layer for existing PL/pgSQL codebases.
- Minimal learning curve for Postgres users.
- Works with existing tools and scripts immediately.

Smart Execution

Built using **tried-and-tested** techniques from DuckDB's history.

The Win:

- Stack-driven interpretation: enables streaming (no memory ballooning).
- No external runtimes.
- Ship your database ↔ ship your code. No dependencies.

The Vision

Designed with advanced optimization techniques in mind:

Apply **Automatic UDF Compilation and Inlining*** research.

The Win:

Native Speed Massive improvements through compilation to SQL.

Hybrid Execution Interleaved interpretation for best performance and full feature set.

* Which I extensively worked on during my PhD—so I'm biased 😊



DuckPL will be open-sourced soon!

