

Compiling PL/SQL Away

CIDR 2020

Christian Duta • Denis Hirn • Torsten Grust
University of Tübingen

How is Everyone Enjoying Their PL/SQL?

From time to time I seek comfort in the success stories that tell how users benefit from the technology we've created ...

PL/SQL: Computation Close to the Data (↓ Robot Walk)

```
CREATE FUNCTION walk(origin coord, win int, loose int, steps int) RETURNS int AS
$$
DECLARE
  reward int = 0;
  location coord = origin;
  movement text = '';
  roll float;
BEGIN
  -- move robot repeatedly
  FOR step IN 1..steps LOOP
    -- where does the Markov policy send the robot from here?
    movement = (SELECT p.action
                  FROM   policy AS p
                  WHERE  location = p.loc);
    -- compute new location of robot,
    -- robot may randomly stray from policy's direction
    roll = random();
    location =
      (SELECT move.loc
       FROM   (SELECT a.there AS loc,
                      COALESCE(SUM(a.prob) OVER lt, 0.0) AS lo,
                      SUM(a.prob) OVER leq AS hi
               FROM   actions AS a
               WHERE  location = a.here AND movement = a.action
               WINDOW leq AS (ORDER BY a.there),
                      lt AS (leq ROWS UNBOUNDED PRECEDING
                           EXCLUDE CURRENT ROW)
              ) AS move(loc, lo, hi)
       WHERE  roll BETWEEN move.lo AND move.hi);
    -- robot collects reward (or penalty) at new location
    reward = reward + (SELECT c.reward
                       FROM   cells AS c
                       WHERE  location = c.loc);
    -- bail out if we win or loose early
    IF reward >= win OR reward <= loose THEN
      RETURN step * sign(reward);
    END IF;
  END LOOP;
  -- draw: robot performed all steps without winning or losing
  RETURN 0;
END;
$$ LANGUAGE PLPGSQL;
```

PL/SQL: $\frac{1}{2}$ Imperative PL + $\frac{1}{2}$ SQL (\downarrow N-Body Simulation)

```
CREATE FUNCTION force(b body, theta float) RETURNS point AS
$$
DECLARE
    force    point := point(0,0);
    G CONSTANT float := 6.67e-11;
    Q        barneshut[];
    node     barneshut;
    children  barneshut[];
    dist     float;
    dir      point;
    grav     point;
BEGIN
    -- enter Barnes-Hut tree at the root
    node = (SELECT t
            FROM   barneshut AS t
            WHERE  t.node = 0);
    Q = array[node];
    -- iterate while there are Barnes-Hut nodes to consider
    WHILE cardinality(Q) > 0 LOOP
        node = Q[1];
        Q    = Q[2:];
        dist = node.center->b.pos;
        dir  = node.center - b.pos;
        grav = point(0,0);
        -- bodies separated by walls do not affect each other
        IF NOT EXISTS (SELECT 1
                       FROM   walls AS w
                       WHERE  (b.pos <= b.pos ## w.wall) <>
                           (node.center <= node.center ## w.wall)) THEN
            grav = (G * b.mass * node.mass / dist^2) * dir;
        END IF;
        -- Barnes-Hut optimization: approximate effect of distant bodies
        IF (node.node IS NULL) OR (width(node.bbox) / dist < theta) THEN
            force = force + grav;
        ELSE
            -- inspect area at higher resolution: descend into subtrees
            children = (SELECT array_agg(t)
                       FROM   barneshut AS t
                       WHERE  t.parent = node.node);
            Q = Q || children;
        END IF;
    END LOOP;
    -- return aggregated force on body
    RETURN force;
END;
$$ LANGUAGE PLPGSQL STABLE STRICT;
```

Elements of PL/SQL: Stateful Variables

```
CREATE FUNCTION  $f(\dots)$  RETURNS  $\tau$  AS
$$
DECLARE
     $V_1$   $\tau_1$ ;
     $V_2$   $\tau_2$ ;
BEGIN
    =====
     $V_1$  = =====;
     $V_2$  = ===== $V_1$ ====;
    =====
     $V_1$  = == $V_1$ == $V_2$ ====;
    =====
END;
$$
```



- Statement sequencing (straight-line control flow)
- Variable references and variable updates

Elements of PL/SQL: Complex and Iterative Control Flow

```
CREATE FUNCTION  $f(\dots)$  RETURNS  $\tau$  AS  
$$
```


```
=====  
WHILE ===== LOOP
```

```
=====  
IF ===== THEN
```

```
=====  
ELSE  
  EXIT;  
END IF;
```

```
=====  
END LOOP;
```

```
=====  
$$
```



- Arbitrarily complex control flow (via IF...THEN...ELSIF, CASE...WHEN, LOOP, WHILE, FOR, EXIT, CONTINUE, ...)

.....



CREATE FUNCTION $f(\dots)$ RETURNS τ AS
\$\$


<pre>===== V₁ = [REDACTED]; =====</pre>	<pre>] simple SQL expression</pre>
<pre>===== V₂ = ([REDACTED]); =====</pre>	<pre>] embedded SQL query Q</pre>

\$\$


- PL/SQL expressions:
 1. “Simple” SQL expressions (evaluated w/o planning)
 2. Embedded SQL queries (require planning + execution)

From SQL to PL/SQL And Back Again

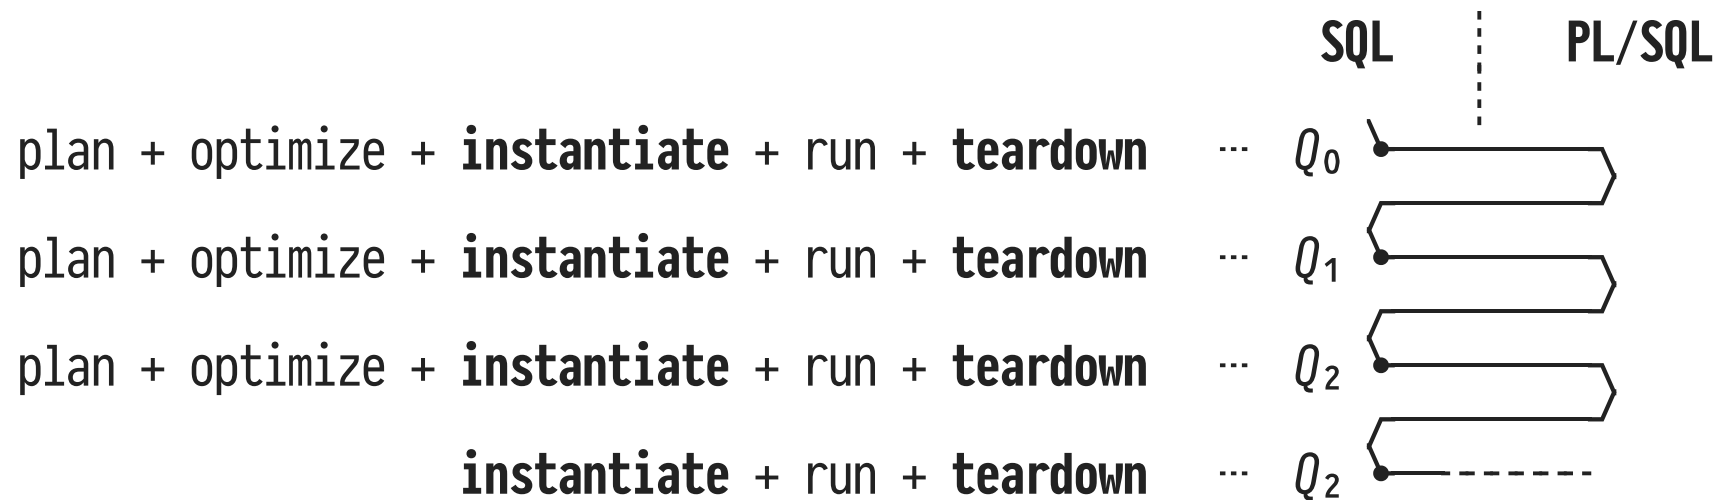
CREATE FUNCTION $f(\dots)$ RETURNS τ AS
\$\$
=====
 $v_1 = ($  $);$] embedded
SQL Q_1
=====
=====
=====
=====
=====
=====
=====
 $v_2 = ($  $);$] embedded
SQL Q_2
=====
=====
\$\$(
SELECT $f(\dots);$] top-level
SQL Q_0




SQL \leftrightarrow PL/SQL (And Again, Again, Again, Again, ...)



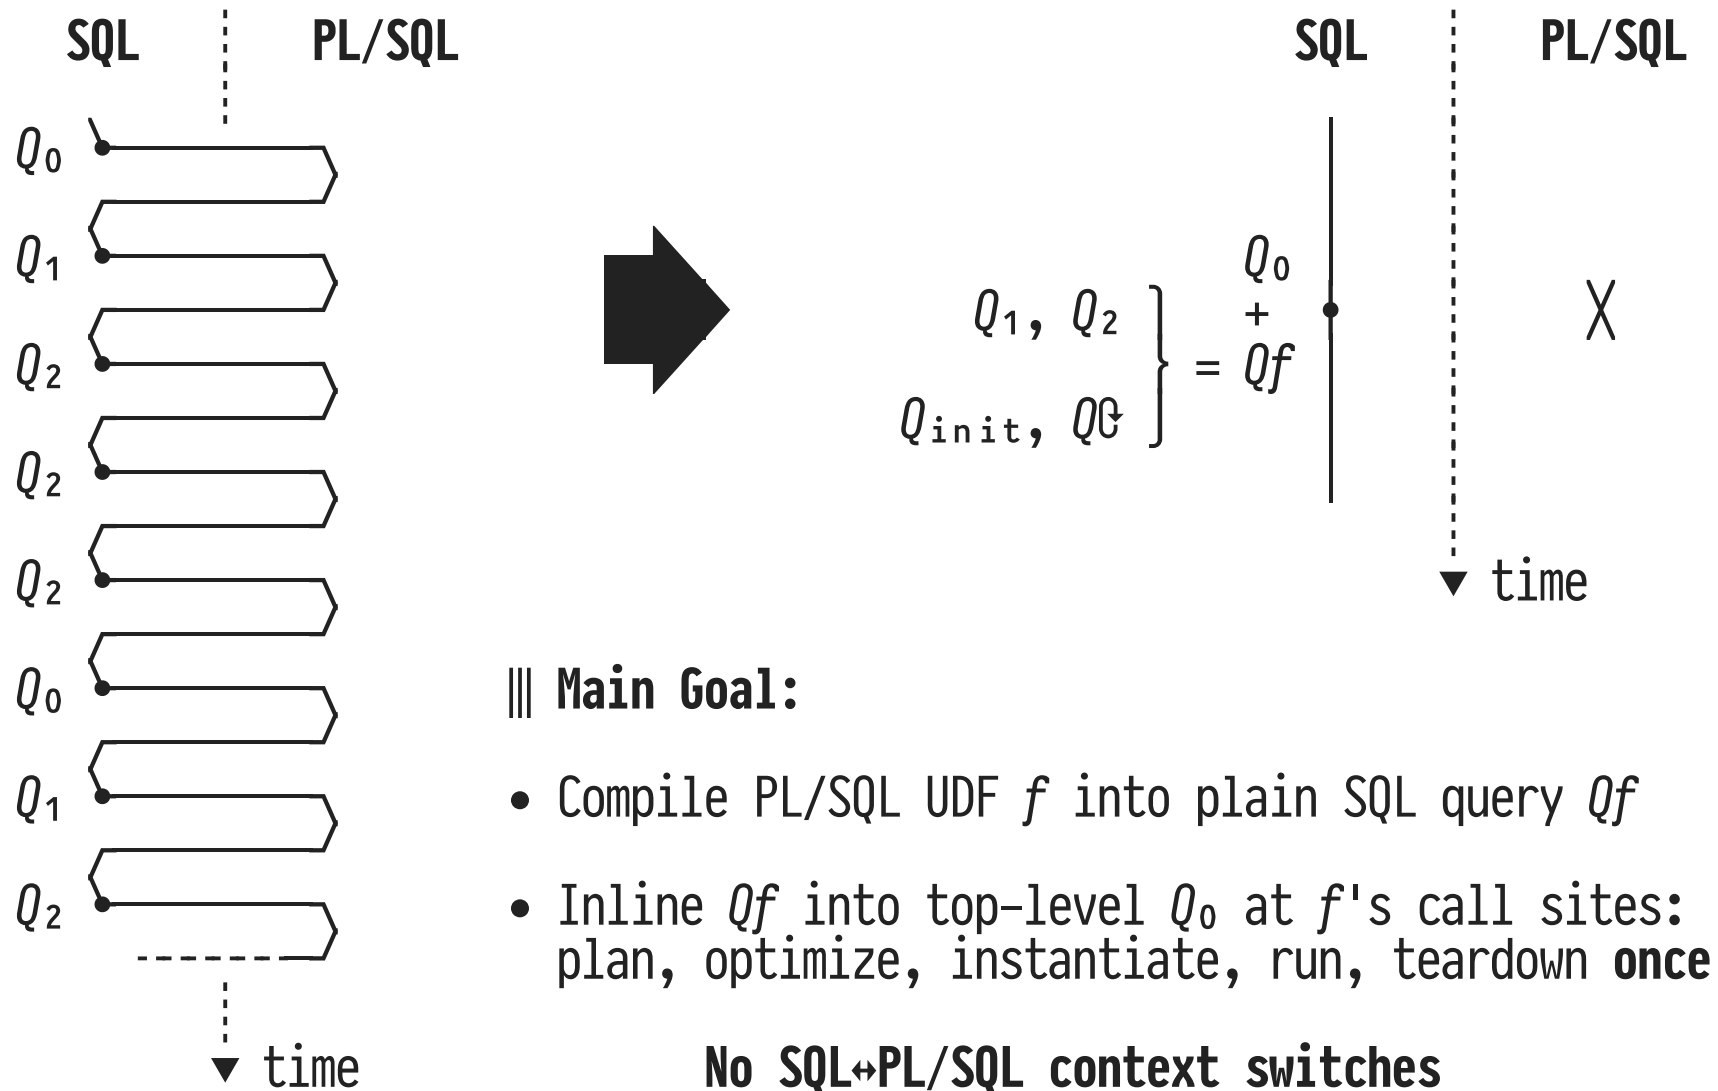
SQL↔PL/SQL Context Switches Are Costly





PostgreSQL 11.3 Runtime Profile 🕒



Compiling PL/SQL Away



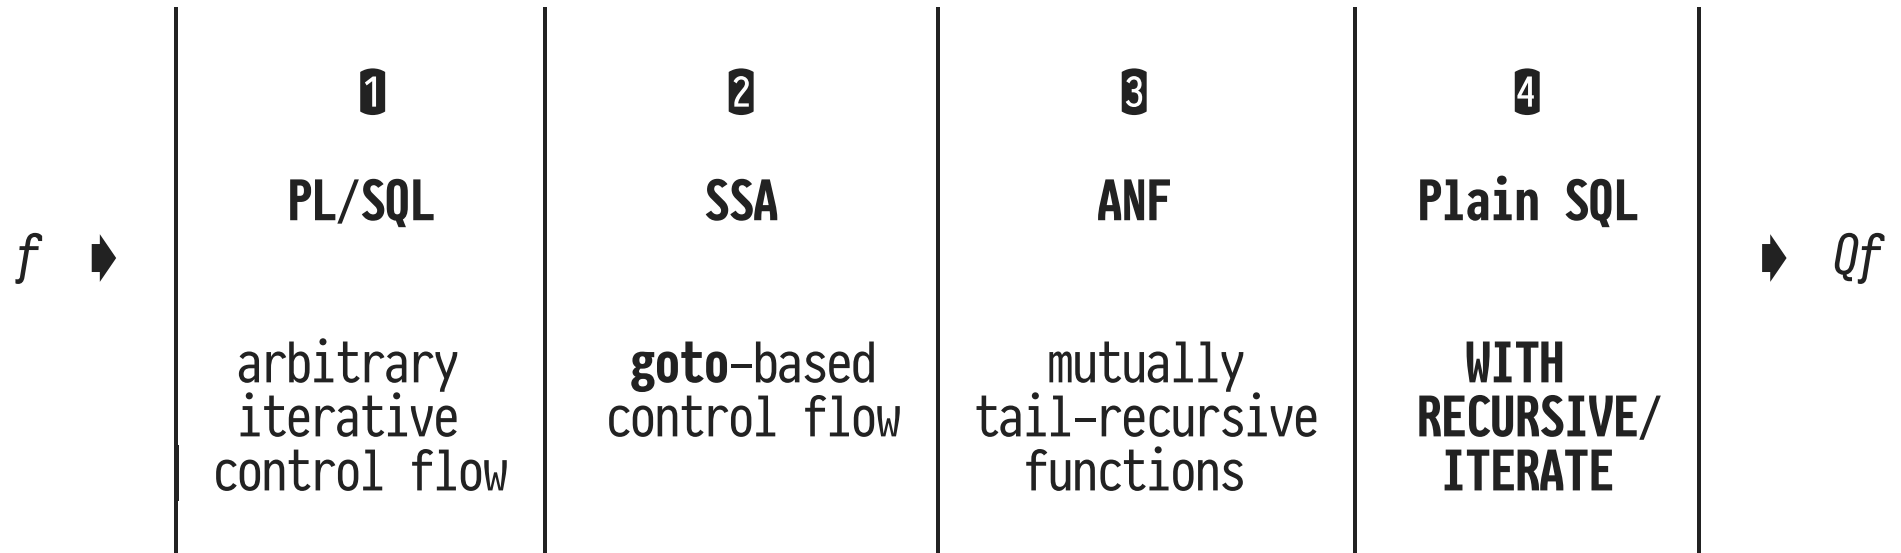
If f Is Iterative, Qf Is Recursive



||| **Challenge:**

Map arbitrary control flow into **WITH RECURSIVE**

PL/SQL to Plain SQL: Left to Right



(⚠ Compiler folks tend to go right to left ◀ instead...)

① From PL/SQL to Single Static Assignment Form

PL/SQL (x^n)	➡	SSA
<pre>CREATE FUNCTION f(x int, n int) RETURNS int AS \$\$ DECLARE i int = 0; p int = 1; BEGIN WHILE i < n LOOP p = p * x; i = i + 1; END LOOP; RETURN p; END; \$\$</pre>		<pre>f(x,n): i₀ ← 0; p₀ ← 1; while: i₁ ← ϕ(i₀, i₂); p₁ ← ϕ(p₀, p₂); if i₁ < n then goto loop; else goto exit; loop: p₂ ← p₁ * x; i₂ ← i₁ + 1; goto while; exit: return p₁;</pre>

- Control flow expressed via **goto**, variables assigned once

② From SSA to Administrative Normal Form

SSA	➡	ANF \Leftrightarrow
<pre>f(x,n): i₀ ← 0; p₀ ← 1; while: i₁ ← ϕ(i₀,i₂); p₁ ← ϕ(p₀,p₂); if i₁ < n then goto body; else goto exit; body: p₂ ← p₁ * x; i₂ ← i₁ + 1; goto while; exit: return p₁;</pre>		<pre>f(x,n) = let i₀ = 0 in let p₀ = 1 in while(i₀,p₀,x,n) while(i₁,p₁,x,n) = let t₀ = i₁ ≥ n in if t₀ then p₁ [← exit] else body(i₁,p₁,x,n) body(i₁,p₁,x,n) = let p₂ = p₁ * x in let i₂ = i₁ + 1 in while(i₂,p₂,x,n)</pre>

- Mutually recursive functions, **tail calls** only

③ From Mutual to Direct Recursion

ANF \leftrightarrow



ANF \mathcal{G}

```
f(x,n) =  
  let i0 = 0 in  
  let p0 = 1 in  
    while(i0,p0,x,n)
```

```
while(i1,p1,x,n) =  
  let t0 = i1 ≥ n in  
  if t0 then p1  
    else loop(i1,p1,x,n)
```

```
loop(i1,p1,x,n) =  
  let p2 = p1 * x in  
  let i2 = i1 + 1 in  
    while(i2,p2,x,n)
```

```
f(x,n) =  
  let i0 = 0 in  
  let p0 = 1 in  
    run(0,i0,p0,x,n)
```

```
run(fn,i1,p1,x,n) =  
  if fn = 0 then  
    let t0 = i1 ≥ n in  
    if t0 then p1  
      else run(2,i1,p1,x,n)
```

```
else  
  let p2 = p1 * x in  
  let i2 = i1 + 1 in  
    run(0,i2,p2,x,n)
```

- Single recursive function *run()*, **tail calls** only

③ From Mutual to Direct Recursion

ANF \leftrightarrow

```
f(x,n) =  
  let i0 = 0 in  
  let p0 = 1 in  
    while(i0,p0,x,n)  
  
while(i1,p1,x,n) =  
  let t0 = i1 ≥ n in  
  if t0 then p1  
    else body(i1,p1,x,n)  
  
body(i1,p1,x,n) =  
  let p2 = p1 * x in  
  let i2 = i1 + 1 in  
    while(i2,p2,x,n)
```



ANF \hookrightarrow

```
f(x,n) =  
  let i0 = 0 in  
  let p0 = 1 in  
    run(i0,p0,x,n)  
  
run(i1,p1,x,n) =  
  let t0 = i1 ≥ n in  
  if t0 then p1  
    else  
      let p2 = p1 * x in  
      let i2 = i1 + 1 in  
        run(i2,p2,x,n)
```

- Single recursive function *run()*, **tail calls** only

④ From Tail Recursion to WITH RECURSIVE

ANF \mathcal{G}


```
f(x,n) =
  let i0 = 0 in
    let p0 = 1 in
      run(i0,p0,x,n)
```

```
run(i1,p1,x,n) =
  let t0 = i1 ≥ n in
    if t0 then p1
    else
      let p2 = p1 * x in
        let i2 = i1 + 1 in
          run(i2,p2,x,n)
```

$\left[\begin{array}{l} Q_{init} \end{array} \right]$

$\left[\begin{array}{l} ? \end{array} \right]$

$\left[\begin{array}{l} Q\mathcal{G} \end{array} \right]$



- Resulting tail recursion now matches the restricted control flow implemented by **WITH RECURSIVE** ✓

④ From Tail Recursion to WITH RECURSIVE

ANF \mathcal{Q}



Plain SQL

```
f(x,n) =
  let i0 = 0 in
  let p0 = 1 in
  run(i0,p0,x,n) ] Qinit .....
```

```
run(i1,p1,x,n) =
  let t0 = i1 ≥ n in
  if t0 then p1
  else
    let p2 = p1 * x in
    let i2 = i1 + 1 in
    run(i2,p2,x,n) ] Q $\mathcal{Q}$  .....
```

WITH RECURSIVE

run("call?",i₁,p₁,x,n,result) AS (

```
  [ SELECT true, 0, 1, x, n, □ ]
```

UNION ALL

SELECT iter.*

FROM run, LATERAL (

```
  [ SELECT false, □, □, □, □, p1
    WHERE i1 ≥ n
    UNION ALL
    SELECT true, i1+1, p1*x, x, n, □
    WHERE i1 < n ]
```

) AS iter("call?",i₁,p₁,x,n,result)

WHERE run."call?")

TABLE run;


Invoking a PL/SQL Function \equiv Running the Recursive CTE

Table *run* (for $f(2,5) = 2^5$)

call?	i_1	p_1	x	n	result
true	0	1	2	5	\square
true	1	2	2	5	\square
true	2	4	2	5	\square
true	3	8	2	5	\square
true	4	16	2	5	\square
true	5	32	2	5	\square
false	\square	\square	\square	\square	32◀


- Look at this table either as...
 - ... the SSA function's **trace of variable states**, or
 - ... the ANF function's **call stack**

SQL Recursion vs. PL/SQL Iteration (Intra-Function)




Scaling the Number of PL/SQL↔SQL Context Switches

#invocations



Markov DP-based Robot Walk



N-Body Simulation

When WITH RECURSIVE Does Too Much

Table *run*

#iter	call?	i_1	p_1	x	n	result
	true	0	1	2	5	□
	true	1	2	2	5	□
	true	2	4	2	5	□
:	true	3	8	2	5	□
	true	4	16	2	5	□
	true	5	32	2	5	□
	false	□	□	□	□	32

Diagram illustrating the recursive process. A vertical arrow labeled #iter points downwards. A curved arrow indicates the flow from the bottom row (false) back to the top row (true). The table shows the state of variables i_1 , p_1 , x, n, and result across iterations. The final result is 32.

```
WITH RECURSIVE run AS (
    :
)
SELECT result
FROM run
WHERE NOT "call?"
```

- Since tail recursion *does not look upwards the stack*, a single-row “stack” suffices

If We Had “WITH TAIL RECURSIVE”...

#iter

Table *run*

call?	i ₁	p ₁	x	n	result
true	0	1	2	5	□
true	1	2	2	5	□
⋮					
true	5	32	2	5	□
false	□	□	□	□	32


↓

↓
WITH ITERATE *run* **AS** (
⋮
)
SELECT result
FROM *run*
WHERE **NOT** "call?"

||| **WITH ITERATE:**

- Singleton table *run*
- Allocates no buffer space

WITH ITERATE: Space-Efficient Iteration



- Local change to PostgreSQL, two implementation avenues


Compiling PL/SQL Away

- Once compiled, **no traces of PL/SQL remain**
 - Thus, zero PL/SQL↔SQL context switches occur
- **Source-to-source compilation** *on top* of the RDBMS
 - Applies to Oracle and SQL Server just as well
- Run compilation chain...
 - ...up to ANF: derive plain SQL UDFs from PL/SQL
 - ...all the way: execute PL/SQL on RDBMSs that implement SQL:1999 but do not support UDFs at all (**SQLite3**)

Compiling PL/SQL Away

CIDR 2020

Christian Duta • Denis Hirn • Torsten Grust
University of Tübingen

 @Teggy | db.inf.uni-tuebingen.de/team/grust