


Ferry — Database-Supported Program Execution


ID	NAME	DEPT	SALARY
1	Alex	DE	300
2	Bert	DE	100
3	Cora	DE	200
4	Drew	US	200
5	Erik	US	400
6	Fred	US	300
7	Gina	US	200
8	Herb	NL	600
9	Ivan	NL	400
10	Jill	UK	500


Your Programming Language's New Friend

Ferry explores query compilation technology that can turn off-the-shelf relational database systems into efficient and scalable co-processors for your programming language's runtime.

- (1) The **host PL prescribes data types and operations**,
- (2) **Ferry compiles fragments of data-intensive computation into bundles of SQL:1999 queries**—for efficient, possibly parallel, execution on an associated database back-end.

Ferry  `map, filter, nth, take, drop, append, zip, unzip, all, any, concat, sort, group, partition, sum, max, min, ...`
`e1.map (x => e2)`

LINQ  `Select, SelectMany, Where, ElementAt, Skip, Take, TakeWhile, Join, Concat, OrderBy, GroupBy, Zip, Any, Max, ...`
`e1.Select (x => e2)`

Ruby  `collect, flatten, select, drop, take, zip, all?, any?, grep, max, min, sum, sort_by, group_by, partition, ...`
`e1.collect {|x|e2}`

With Ferry, you continue to use the idioms and operations of your favorite programming language to describe computation over **ordered** and **arbitrarily nested** records, lists, arrays, or maps.

Ferry's **list comprehension-based** compilation scheme is as compositional as your programming language's semantics and will not break when facing complex programs.

Order and Nesting on 1NF RDBMSs

Ferry has been designed to support the compilation and execution of programs over ordered, nested data structures. To realize this on top of unordered First Normal Form (1NF) SQL back-ends,

- (1) Ferry **embeds order information in the data** itself (if relevant), and
- (2) represents **nested data structures via flat tables containing surrogate key values**—much like in the 1980s.

Count the [·]s to Count the Queries

The (static) **type of your program's result** determines **how many database queries Ferry will emit** to compute that result. Nothing else impacts the query bundle size. Period.

If the result type T features n list constructors $[·]$, Ferry prepares a bundle of n **independent queries** for execution.

(With Microsoft's LINQ to SQL, the number of submitted queries may depend on the database instance size. Yes, that may be *many* queries.)

Example: Result Type T ($n = 4$)

$[(int, [str], [(bool, [(int, int)])])]$



Programming with Tables

"Who are the two top-paid staff members per department in Table Employees?"

Ferry permits a compositional programming style that glues operations over tables and ordered lists, e.g. `zip(·, ·)`. Nesting occurs naturally. Somewhat close to Haskell. And XQuery. And LINQ. And SQL.

Ferry Sample Program P: Two Top-Paid Employees per Department

```
let e = table Employees (id int, name string,
                        dept string, salary int)
    with key (id)
in
  for x in e
  group by x.dept
  return (the (x.dept),
         take (2, for y in zip (x.name, x.salary)
                              order by y.salary descending
                              return y))
```

Result Type of Program P

$[(str, [str, int])]$




Intermediate Language: A Table Algebra Dialect

ρ	serialization point (sub-plan root)
$\pi_{a_1, b_1, \dots, a_n, b_n}$	project onto columns b_i , rename b_i into a_i
σ_p	select rows satisfying predicate p
$\rho_{a,v}$	attach column a containing constant value v
\times	Cartesian product, join with predicate p
\cup	disjoint union, difference
δ	eliminate duplicate rows
$\rho_{a:(b_1, \dots, b_n)/c}$	group rows by c , then attach row rank in a
$\rho_{a:(b_1, \dots, b_n)}$	attach result of application $*$ (b_1, \dots, b_n) in a
$\text{agg}_{a:(b)/c}$	group rows by c , then attach aggregated b in a
ρ_R	access database-resident table R
abc	literal table with columns a, b, c

Ferry compiles programs into a particularly simple dialect of a table algebra whose design **mimics the actual capabilities of modern relational query processors**. (No wishful thinking here.)

The compiler implements **loop lifting**, a translation technique that derives **set-oriented plans from iterative programs**. Primary goal: operate the database back-end in its efficient bulk processing mode.

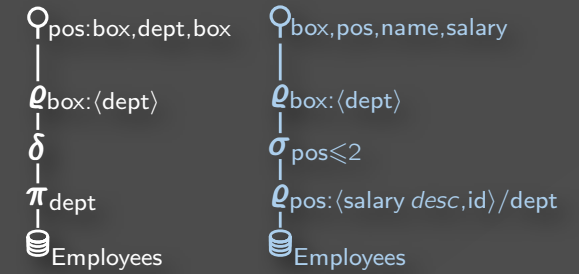
Shape of Algebraic Plan Bundle for Type T



Ferry's algebraic compiler yields a bundle of **independent queries**. The database back-end **may execute these queries in any order** it sees fit—or even in parallel.

(With Microsoft's LINQ to SQL, your back-end may experience an avalanche of sequentially submitted queries, each entailing an execution context switch.)

Algebraic Plan Bundle for Program P



Two independent algebraic plans compute the list contents for their associated list constructors. (Note: Plan \blacksquare computes a single list, Plan \blacksquare computes the contents of *many* nested lists—**with column box**, we may stitch the pieces together later on.)

In **column pos**, the plans embed **order in the data**: positional access (`take(n , ·)`) or alignment (`zip(·, ·)`) indeed makes sense.


Target Language: SQL:1999, Off the Shelf

<code>[ROW_NUMBER() DENSE_RANK()] OVER (PARTITION BY ... ORDER BY ...)</code>	implementation of order semantics and nested lists
<code>[RIGHT OUTER] JOIN</code>	implementation of iteration
<code>UNION ALL</code>	sequence construction
<code>WITH ... AS ...</code>	basic block forming

Ferry comes equipped with a **code generator deriving standard SQL:1999 common table expressions** from algebraic plans.

Such code generators have been devised for IBM DB2 V9, MS SQL Server 2008, CWI's MonetDB column store, and Ix Systems' kdb+. If your query engine can order and then number table rows, it is most probably a suitable Ferry back-end.

Lazy and/or Partial Plan Evaluation



The result has a **sensible interpretation per table**, each representing **all lists** of its associated type constructor. Table \blacksquare already delivers

$[(42, @_1, @_2), (43, @_3, @_4), (44, @_5, @_6), \dots]$

If your program does not inspect lists beyond nesting level 1 yet, there is no need to compute Tables \blacksquare , \blacksquare , or \blacksquare right now.

Fully Evaluated Result of Program P

pos	dept	box	box	pos	name	salary
			@ ₁	@ ₁	Alex	300
			@ ₂	@ ₂	Cora	200
1	DE	@ ₁	@ ₂	@ ₁	Erik	400
2	US	@ ₂	@ ₁	@ ₂	Fred	300
3	NL	@ ₃	@ ₁	@ ₃	Herb	600
4	UK	@ ₄	@ ₂	@ ₄	Ivan	400
			@ ₄	@ ₄	Jill	500

Both tables **contribute pieces to the result** of Program P and may be sensibly computed, consumed, and interpreted on their own.

The **surrogate keys @_k in columns box** enable the runtime system to assemble the complete result. And: without columns **box|pos**, Ferry's encoding of **flat unordered lists and regular DBMS tables coincide**.

Ferry vs. LINQ to SQL vs. Ruby— How Much Faster Do You Want To Be Today?

Queries Emitted for Program P

# Depts.	Ferry # Queries	LINQ # Queries
1 000	2	1 001
10 000	2	10 001
100 000	2	100 001

Execution Times for Program P

# Depts.	Ferry (sec)	LINQ (sec)
1 000	0.140	0.265
10 000	0.296	1.230
100 000	1.948	10.486

Ruby 1.9 Variant of Program P

# Depts.	Ruby (Heap Execution) (sec) load + execute
1 000	0.188 + 0.031
10 000	1.240 + 0.350
100 000	11.600 + 8.312

Ferry supports a data model and language semantics that is so close to Microsoft's **Language Integrated Query (LINQ)** that you can use Ferry as a replacement for the Microsoft LINQ to SQL provider. If you do, the **family of applicable Standard Query Operators (SQOs) grows**, all SQOs adhere to **proper order semantics**, and **no query avalanches** will be generated.

If your programs process significant amounts of database-resident data, Database-Supported Program Execution can save the host PL's runtime system from a lot of work. And from a job it has not been designed for.

Ferry—Database-Supported Program Execution.

Heap Image of the Result of Program P

```
[
  ("DE", [ ("Alex", 300), ("Cora", 200) ]),
  ("US", [ ("Erik", 400), ("Fred", 300) ]),
  ("NL", [ ("Herb", 600), ("Ivan", 400) ]),
  ("UK", [ ("Jill", 500) ])
]
```

A merge of the sorted Tables \blacksquare and \blacksquare , performing a left outer join, suffices to materialize the result on the heap of the host PL.

On-demand materialization may be implemented in terms of **holes**: just scan Table \blacksquare to produce $[("DE", \square), \dots]$, where hole \square represents a suspension that will scan Table \blacksquare once it is referenced.