

# Data-Parallel Computing Meets STRIPS

Erez Karpas, Tomer Sagi, Carmel Domshlak, Avigdor Gal, Avi Mendelson, and Moshe Tennenholtz

Technion-Microsoft Electronic Commerce Research Center

## Motivation

- Declarative query processing and user defined functions do not play well together
  - User must specify some base execution plan
  - Which optimizations are safe?

## DPSS

- Framework is based on tracking *data chunks*
- Each data chunk  $d$  is associated with the amount  $\sigma_d$  of memory it requires

A DPSS Task is described by:

- $D$  — possible data chunks, with sizes  $\sigma_d$
- $N$  — computing units, with memory  $\kappa_n$
- $A$  — computation primitives, each described by:
  - $\bar{I} \subseteq D$  — required input
  - $\bar{O} \subseteq D$  — produced output
  - $C : N \rightarrow \mathbb{R}^{0+}$  — computation cost
  - $T : N \times D \times N \rightarrow \mathbb{R}^{0+}$  — data transmission cost
- $s_0$  — the initial state of the computation
- $G$  — the goal of the computation

- A DPSS state specifies which processor holds which data chunks
- A solution is a sequence of compute / transmit / delete data actions which achieves the goal from the initial state

- The possible data chunks  $D$  and computations  $A$  may be given explicitly or described implicitly
- If they are described implicitly the sets could be infinite

## Theoretical Properties

### Expressivity:

⊙ DPSS is at least as expressive as relational algebra with aggregation

### Complexity:

⊙ Optimal data-parallel program synthesis is NP-hard, even under severe restrictions

⊙ Satisficing data-parallel program synthesis is NP-hard

⊙ Satisficing data-parallel program synthesis with no memory constraints can be solved in polynomial time, when the possible data chunks are given explicitly.

## Example: Multi Histogram

- Suppose we have a users table  $T$  with  $10^9$  users
- We want two histograms of  $T$ : by age and by relationship status

In SQL or similar:

```
SELECT COUNT(T.age) FROM T;
SELECT COUNT(T.rls) FROM T;
```

Query Execution Plan:

```
Agg (age, scan(T))
Agg (rls, scan(T))
```

- Suppose we have a user-defined function,  $DAgg$ , which aggregates by two fields simultaneously

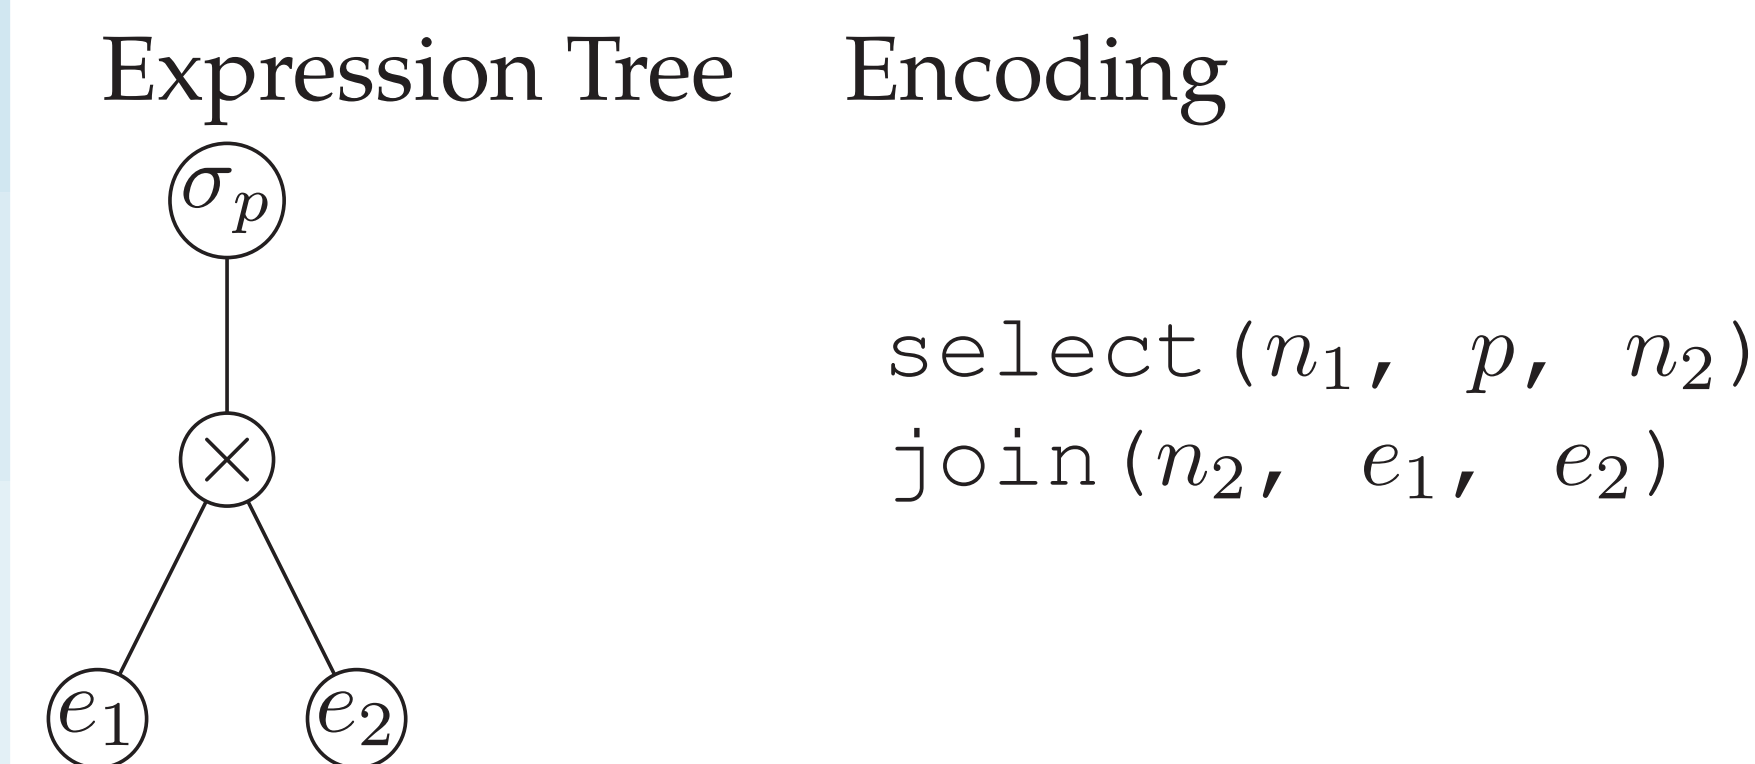
Query Execution Plan using  $DAgg$ :

```
DAgg (age, rls, scan(T))
```

- How do we come up with this execution plan?

## Compilation to STRIPS

- When computations/data chunks explicit:
  - Predicate holds( $?node, ?data$ )
  - Actions
    - compute( $?node, ?computation$ )
    - transmit( $?node, ?data, ?node2$ )
    - del( $?node, ?data$ )
  - Capacity constraints = numerical fluents
- When the computations/data chunks implicit, compilation is still possible sometimes
- When data chunks have a structure (e.g., expression trees), it is possible to represent such trees using predicates

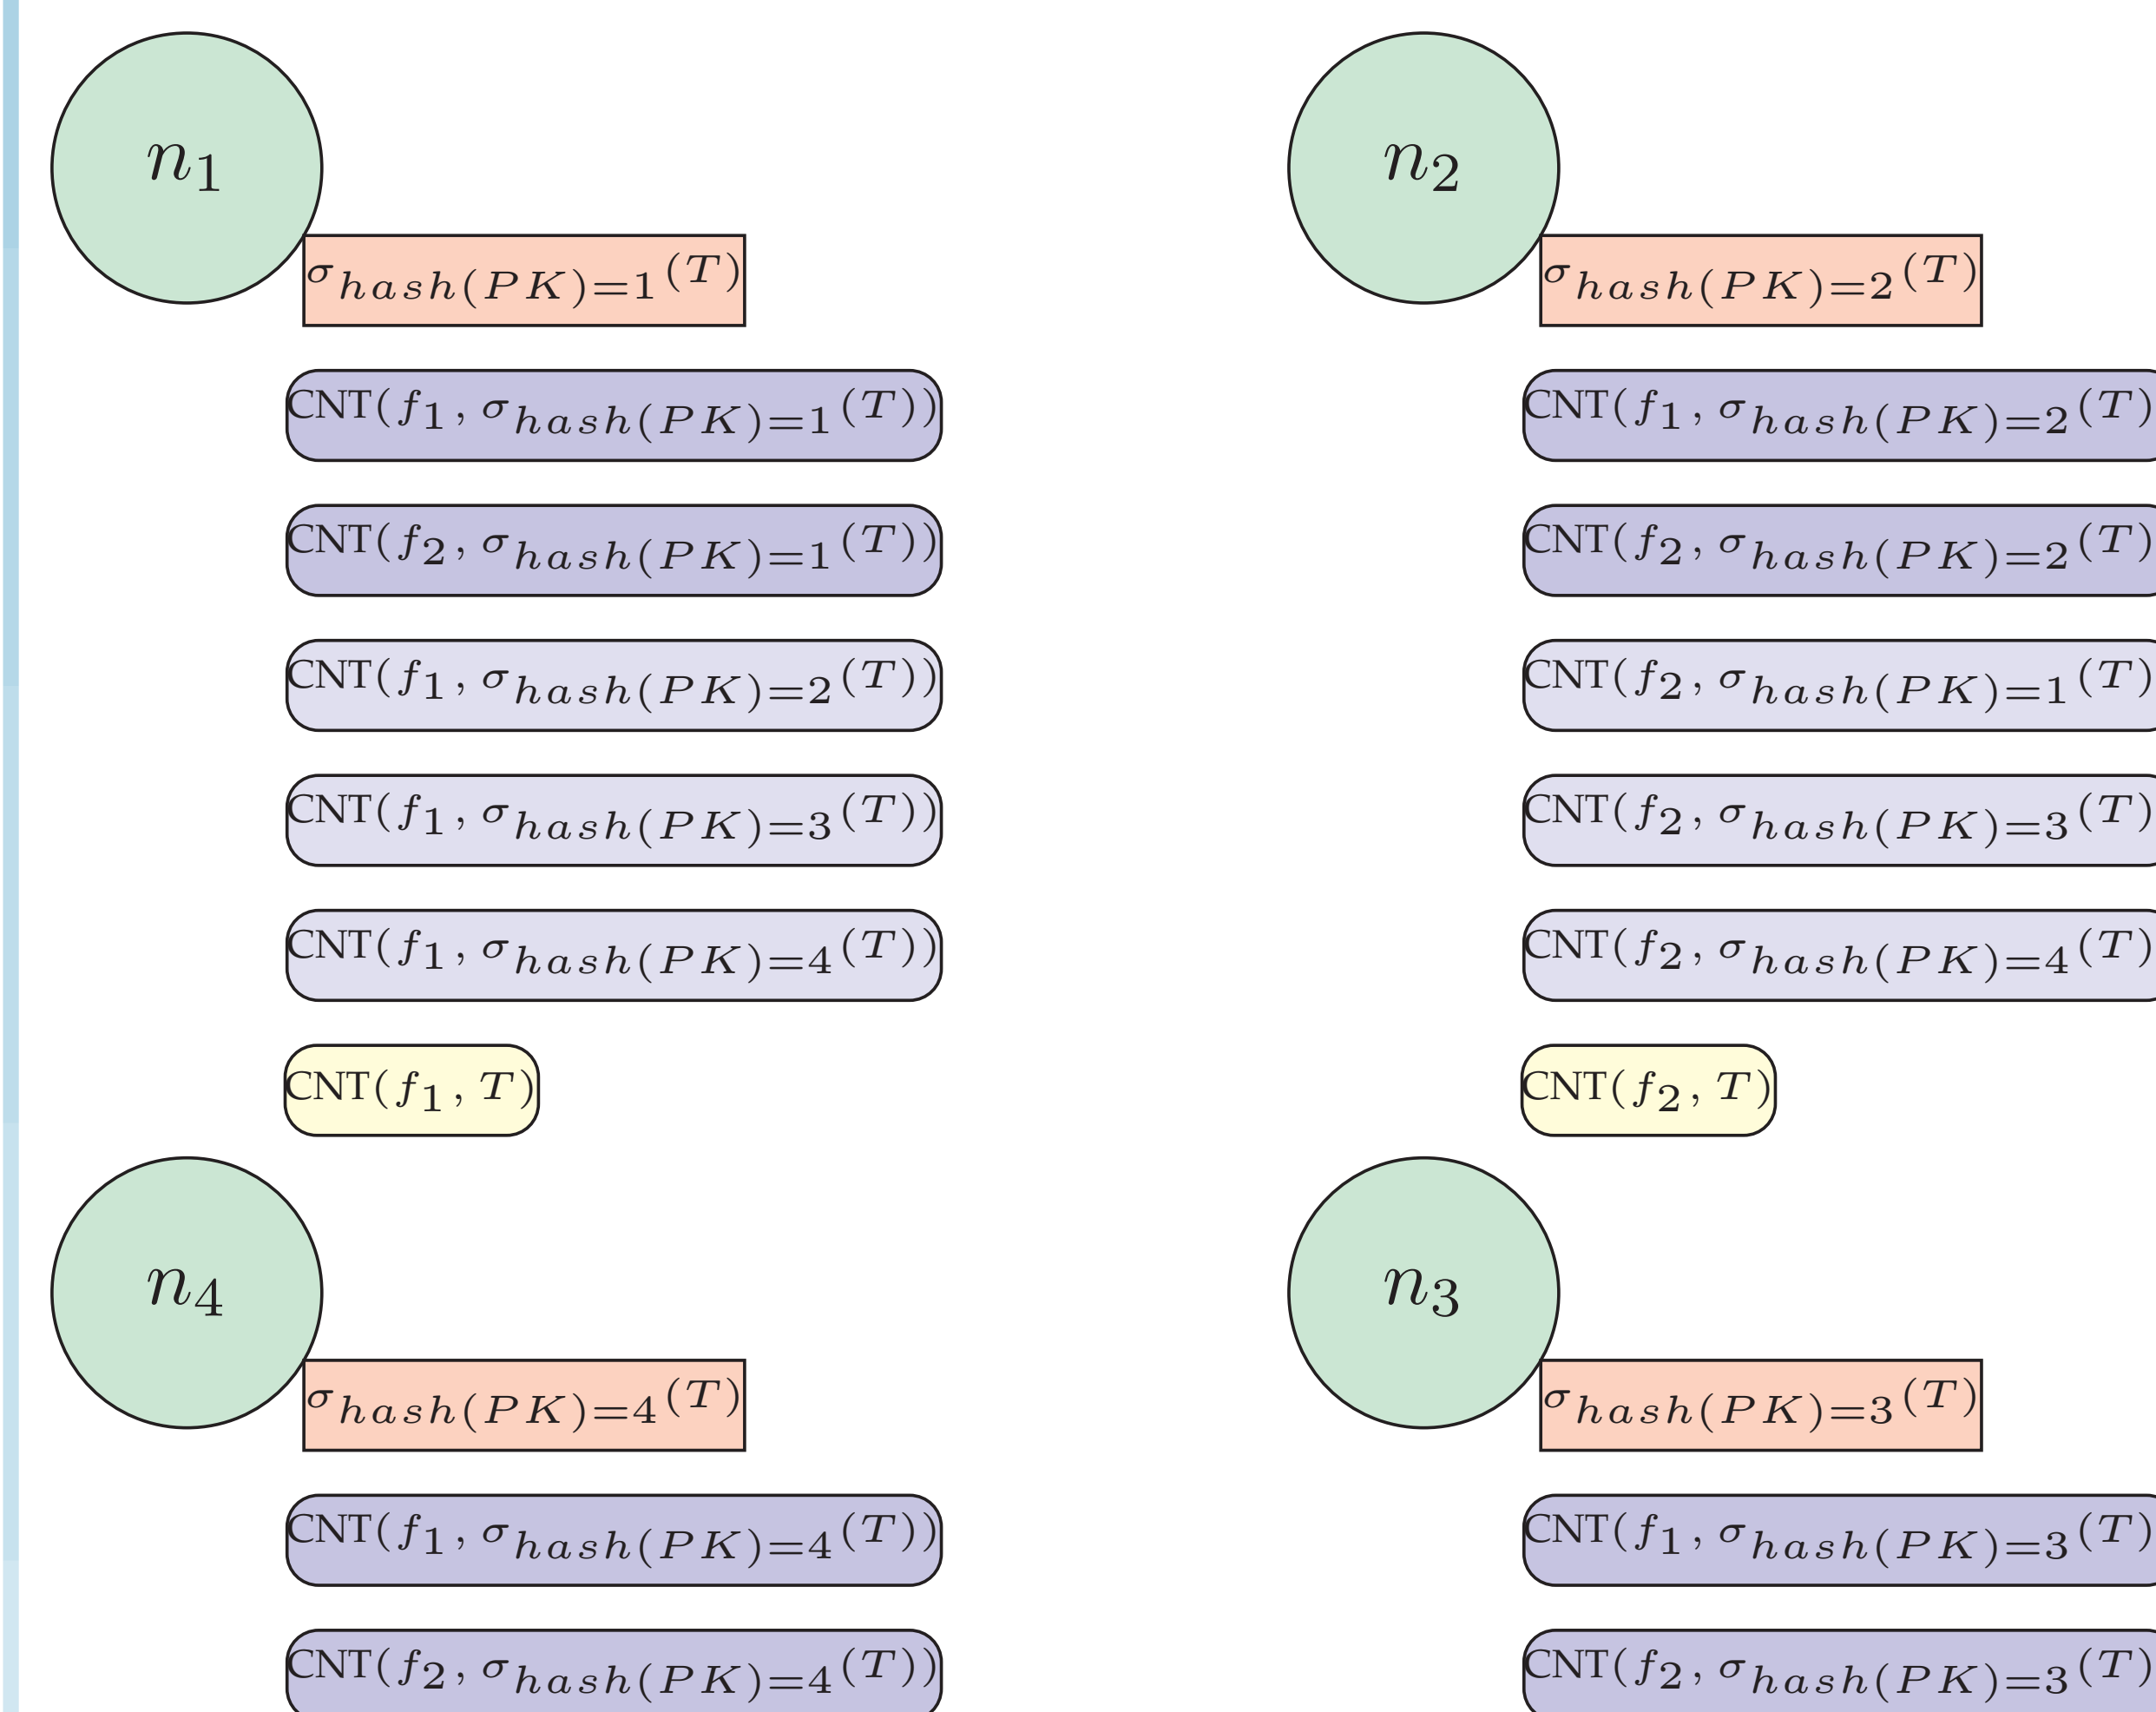


- Equivalence rules typically have limited depth, and can be encoded as operators
- More details in the paper

## Empirical Proof of Concept

- Histogram of  $F$  fields of a table divided across  $N$  processors

Example for  $N = 4, F = 2$




Query Execution Plan without  $DAgg$

```
Agg (n1, f1, sigma_hash(PK)=1(T))
Agg (n1, f2, sigma_hash(PK)=1(T))
Agg (n2, f1, sigma_hash(PK)=2(T))
Agg (n2, f2, sigma_hash(PK)=2(T))
Agg (n3, f1, sigma_hash(PK)=3(T))
Agg (n3, f2, sigma_hash(PK)=3(T))
Agg (n4, f1, sigma_hash(PK)=4(T))
Agg (n4, f2, sigma_hash(PK)=4(T))
transmit (n2, CNT(f1, sigma_hash(PK)=2(T)), n1)
transmit (n3, CNT(f1, sigma_hash(PK)=3(T)), n1)
transmit (n4, CNT(f1, sigma_hash(PK)=4(T)), n1)
merge (n1, f1)
transmit (n1, CNT(f2, sigma_hash(PK)=1(T)), n2)
transmit (n3, CNT(f2, sigma_hash(PK)=3(T)), n2)
transmit (n4, CNT(f2, sigma_hash(PK)=4(T)), n2)
merge (n2, f2)
```

Query Execution Plan using  $DAgg$

```
DAgg (n1, f1, f2, sigma_hash(PK)=1(T))
DAgg (n2, f1, f2, sigma_hash(PK)=2(T))
DAgg (n3, f1, f2, sigma_hash(PK)=3(T))
DAgg (n4, f1, f2, sigma_hash(PK)=4(T))
transmit (n2, CNT(f1, sigma_hash(PK)=2(T)), n1)
transmit (n3, CNT(f1, sigma_hash(PK)=3(T)), n1)
transmit (n4, CNT(f1, sigma_hash(PK)=4(T)), n1)
merge (n1, f1)
transmit (n1, CNT(f2, sigma_hash(PK)=1(T)), n2)
transmit (n3, CNT(f2, sigma_hash(PK)=3(T)), n2)
transmit (n4, CNT(f2, sigma_hash(PK)=4(T)), n2)
merge (n2, f2)
```

- Solved by GBFS using relaxed plan heuristic in Fast Downward
- Table below shows planning time for different values of  $F$  and  $N$



- Solutions were optimal (although this is not guaranteed)

## Acknowledgements

This work was supported by the Technion-Microsoft Electronic-Commerce Research Center.