



Mapper

An Efficient Data Transformation Operator

Paulo Carreira
University of Lisbon

Context

Source: **IMPCRED**

CREDITID	MERCH	...
FT546083	8 PCS LARGE CONTAINER MODULE X067 100 PCS COTTON CORDUROY 8 500 MTS XS03 208 63 PCS POLYESTER TUBE W987 34 PCS ARTIFICIAL FIBRE 3 PCS SPANDEX 2 400 MTS W15719 83 PCS POLYESTER TUBE W989 15 PCS ARTIFICIAL FIBRE 3 500 MTS NEW CHECK DESIGN	...
RT546084	4800 PCS IC CAN TRANSCEIVER 3.3V 8-SOIC SN65HVD232D 2400 PCS PCA9532BS-T 16-bit I2C LED DIMMER	...
...
...

Target: **MERCHDETL**

ORDRID	QTY	DESCR	...
546083	8	LARGE CONTAINER MODULE X067	...
546083	100	COTTON CORDUROY 8 500 MTS XS03 208	...
546083	63	POLYESTER TUBE W987	...
546083	34	ARTIFICIAL FIBRE	...
546083	3	SPANDEX 2 400 MTS W15719	...
546083	83	POLYESTER TUBE W989	...
546083	15	ARTIFICIAL FIBRE 3 500 MTS NEW CHECK DESIGN	...
546084	4800	IC CAN TRANSCEIVER 3.3V 8-SOIC SN65HVD232D	...
546084	2400	PCA9532BS-T 16-bit I2C LED DIMMER	...
...
...
...
...



State of the art

	General Purpose Language	ETL Tool	RDBMS				
			Union	Unpivot	Recursive Queries	PSMs	?
Declarativeness	-	+/-	+	+	+	+/-	+
Optimizability	-	-	+	+	+/-	-	+
Expressiveness	+	+/-	-	-	+	+	+

Main Contributions

1. Extending of Relational Algebra

- **Data Mapper:** An operator for expressing One-to-many data transformations

2. Logical optimization rules for the mapper operator, e.g.,

- Pushdown of projections
- Pushdown of selections

3. Physical execution algorithms for the mapper operator

- Naïve algorithm
- Shortcircuiting algorithm
- Cache-based algorithm

The mapper operator

$$\mu_F(s) = \bigcup_{u \in s} f_{A_1}(u) \times \dots \times f_{A_k}(u)$$

Source schema
IMPCRED (CREDID, MERCHANDISE)

CREDID	MERCHANDISE
PT546084	4800 PCS IC CAN TRANSCEIVER 3.3V 8-SOIC SN65HVD232D 2400 PCS PCA9532BS-T 16-bit I2C LED DIMMER
...	<i>u</i> ...

Target schema

MERCHDETL (ORDERID, QTY, DESCR)

ORDERID	QTY	DESCR
546084	4800	IC CAN TRANSCEIVER 3.3V 8-SOIC SN65HVD232D
546084	2400	PCA9532BS-T 16-bit I2C LED DIMMER
...

$$\mu_{cid_{ORDERID}, extr_{QTY, DESCR}}(s)$$

$cid_{ORDERID}(u)$

546084
$t[ORDERID]$

$extr_{QTY, DESCR}(u)$

4800	IC CAN TRANSCEIVER 3.3V 8-SOIC SN65HVD232D
2400	PCA9532BS-T 16-bit I2C LED DIMMER
$t[QTY, DESCR]$	

X

Algebraic Optimization

Pushdown of σ

$$\pi_Z(\sigma_{ORDRID < 500000 \wedge GARTYP \neq 'R'}(R))$$

1

Introduction

$$\mu_F(\sigma_{ACCTNO=CREDAC}(R))$$

Pushdown of σ

$$\sigma_{C_B}$$

Pushing select

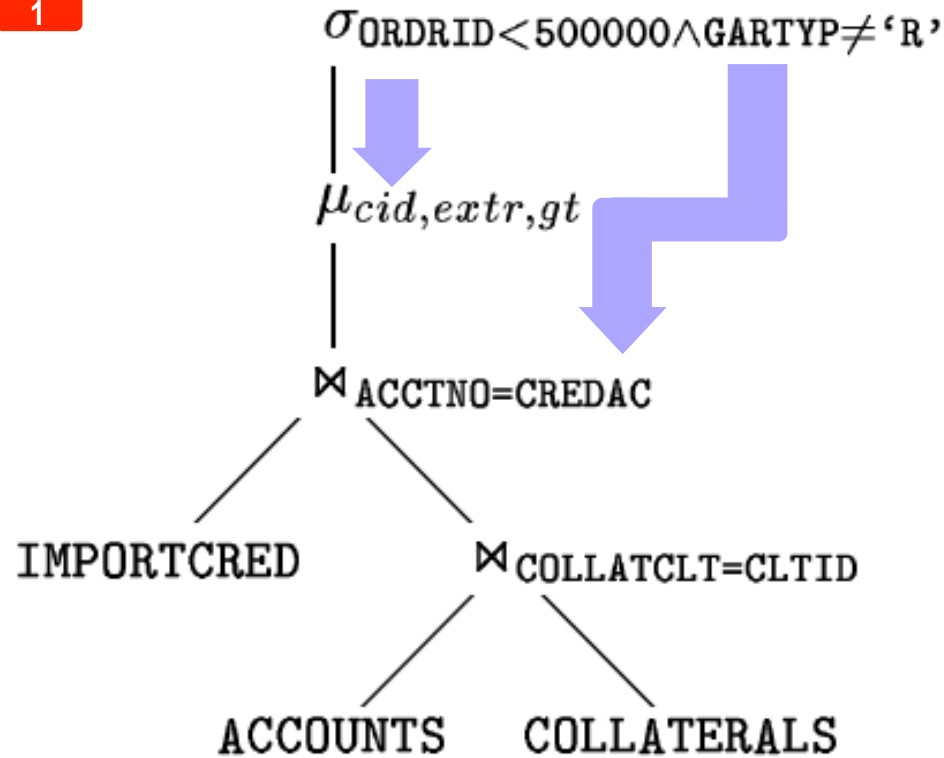
$$\sigma_{C_A}$$

Distributing μ

$$\mu_F(\sigma_{ACCTNO=CREDAC}(R))$$

Distributing μ

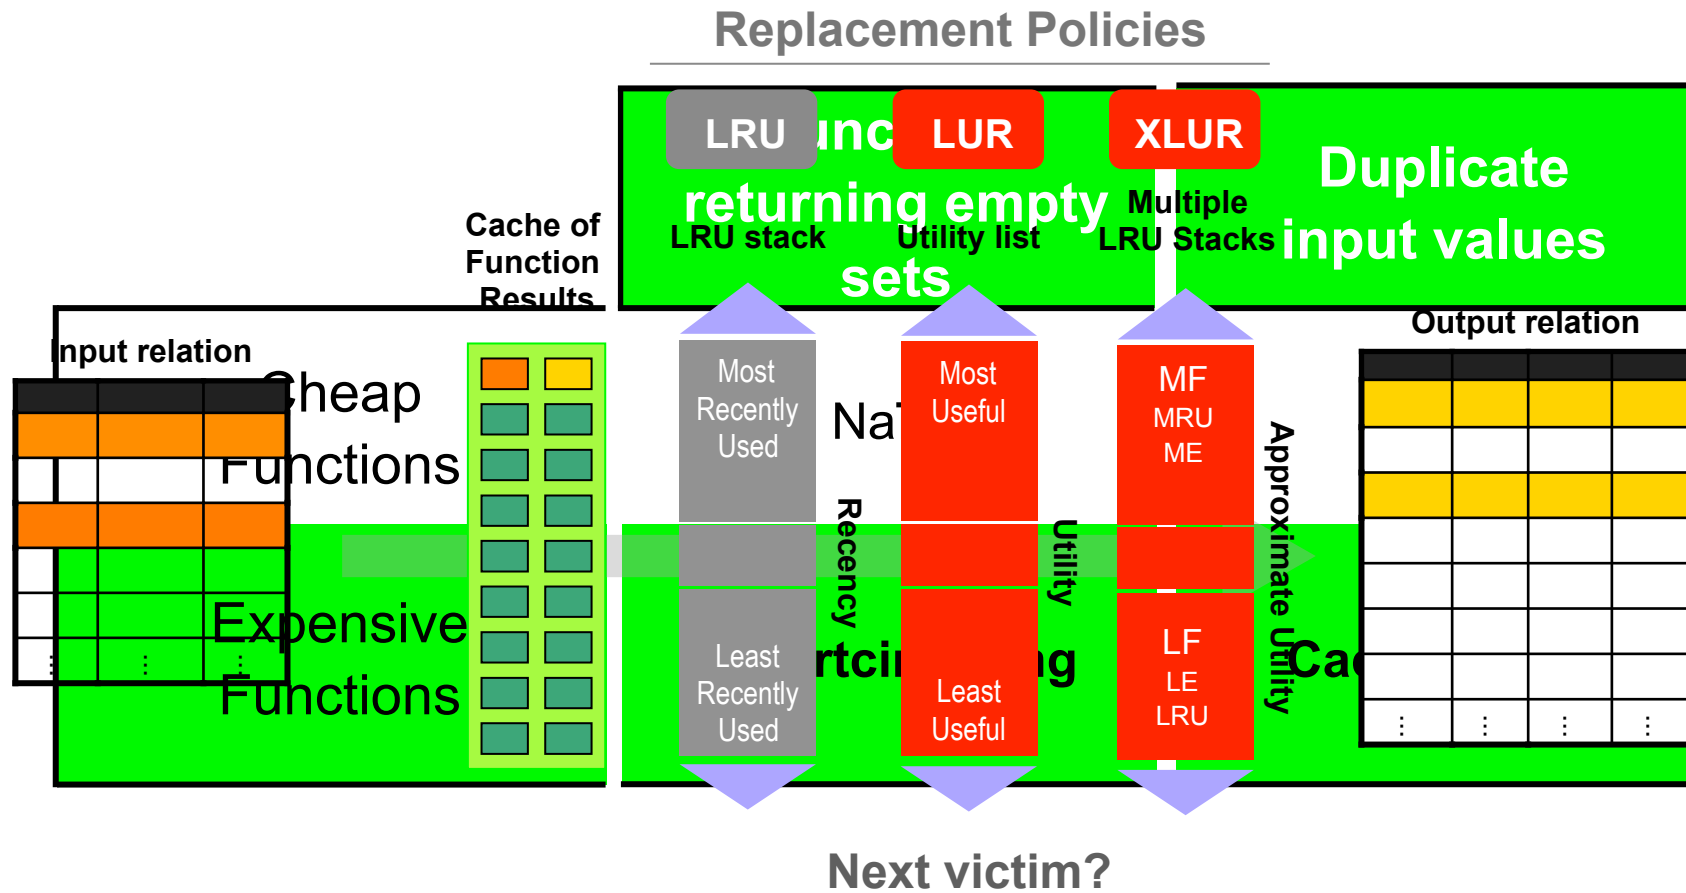
$$\mu_F(\sigma_{ACCTNO=CREDAC}(R))$$



$\sigma_{ACCTNO=CREDAC}(R)$

R'

Physical Algorithms



Experimental Validation

1. Comparison of RDBMSs with mappers

- Mappers are 1.7x – 3.6x faster than the best RDBMS solution

2. Gain of logical optimizations

- Gains of 1.3x – 5.5x for pushing selections

3. Performance of physical algorithms

- Shortcircuiting algorithm: 50x faster than the Naïve (at ~1ms/call)
- Cache-based algorithm (XLUR)
 - Interesting savings: 50% duplicates = 1.4x faster than the Naïve
 - Lightweight: virtually the same overhead as LRU
 - Successful at performing cost-based decisions: outperforms LRU for low CHRs

Conclusions

- New unary operator extension to Relational Algebra:
Data Mapper

- Addresses: One-to-many data transformations
 - Declarative
 - Optimizable
 - Expressive

- Consequences:
 - Broadens the span of application of RDBMSs
 - Significant improvement for ETL tools (Data Fusion tool)

Selected Publications

- [Carreira et al. 2007] Carreira, P., Galhardas, H., Lopes, A. & Pereira, J., “*One-to-many Data Transformations through Data Mappers*”, *Data & Knowledge Engineering Journal (DKE)*, 62(3), 483–503, Elsevier-Science, **2007**.
- [Carreira et al. 2005] Carreira, P., Galhardas, H., Pereira, J., & Lopes, A., “*Data Mapper: An operator for expressing one-to-many data transformations*”. 7th Int’l Conference on Data Warehousing and Knowledge Discovery, **DaWaK '05**, Vol. 3589 of LNCS, Springer-Verlag, **2005**.
- [Carreira et al. 2005b] Carreira, P., Galhardas, H., Lopes, A. & Pereira, J. (2005). “*Extending relational algebra to express one-to-many data transformations*”. In 20th Brazillian Symposium on Databases **SBBD '05**, **2005**.
- [Carreira & Galhardas 2004] Carreira, P., & Galhardas, “*Efficient development of data migration transformations*”, Demo paper, Proc. of the ACM Conference on the Management of Data, **SIGMOD'04**, **2004**.
- [Carreira et al. 2007b] Carreira, P., Galhardas, H., Pereira, J., Martins F. & Silva, Mário J., “*On the performance of One-to-Many Data Transformations*”, Venkatesh Ganti, Felix Naumann (Eds.): Proceedings of the 5th International Workshop on Quality in Databases, **QDB 2007 at VLDB Conference**, **2007**.



Reserve Slides

Future work

- Study mapper functions
 - Study the properties of mapper functions
 - Develop a taxonomy of mapper functions
 - Propose a library of useful mapper functions

- Enhance the physical algorithms

- Incorporate the mapper operator on an Open Source RDBMS

Extracting rows from unstructured data

Source: CITEDATA

Target: EVENTS

AUTHORS	TITLE	...
Periklis Andritsos, Ronald Fagin, Ariel Fuxman, Laura M. Haas, Mauricio A. Hernández, C. T. Howard Ho, Anastasios Kementsietsidis, Renée J. Miller, Felix Naumann, Lucian Popa, Yannis Velegrakis, Charlotte Vilarem, Ling-Ling Yan	Schema Management	...
Jorge Vieira, Jorge Bernardino, Henrique Madeira	Efficient Compression of Text Attributes of Data Warehouse Dimensions	...
...
...

NAME	TITLE	...
Periklis Andritsos	Schema Management	...
Ronald Fagin	Schema Management	...
Ariel Fuxman	Schema Management	...
Laura M. Haas	Schema Management	...
Mauricio A. Hernández	Schema Management	...
C. T. Howard Ho	Schema Management	...
Anastasios Kementsietsidis	Schema Management	...
Renée J. Miller	Schema Management	...
Felix Naumann	Schema Management	...
Lucian Popa	Schema Management	...
Yannis Velegrakis	Schema Management	...
Charlotte Vilarem	Schema Management	...
Ling-Ling Yan	Schema Management	...
...



Example: Converting lines into columns

Source: LOANEVT

LOANO	EVTYP	CAPTL	TAX	EXPNS	BONUS
1234	OPEN	0.0	0.19	0.28	0.1
1234	PAY	1000.0	0.28	0.0	0.0
1234	PAY	1250.0	0.30	0.0	0.0
1234	EARLY	550.0	0.0	0.0	0.0
1234	FULL	5000.0	1.1	5.0	3.0
1234	CLOSED	0.0	0.1	0.0	0.0

Target: EVENTS

LOANO	EVTYP	AMTYP	AMT
1234	OPEN	TAX	0.19
1234	OPEN	EXPNS	0.28
1234	OPEN	BONUS	0.1
1234	PAY	CAPTL	1000.0
1234	PAY	TAX	0.28
1234	PAY	CAPTL	1250.0
1234	PAY	TAX	0.30
1234	EARLY	CAPTL	550.0
1234	FULL	CAPTL	5000.0
1234	FULL	TAX	1.1
1234	FULL	EXPNS	5.0
1234	FULL	BONUS	3.0
1234	CLOSED	EXPNS	0.1

Implementations of one-to-many data transformations

	Bounded				Unbounded			
	RA/SQL	PSM	Recursive Queries	Unpivot	RA/SQL	PSM	Recursive Queries	Unpivot
DBX	YES	YES	YES	N/A	NO	YES	YES	NO
OEX	YES	YES	N/A	N/A	NO	YES	N/A	NO

Extensions to Relational Algebra

Type of data transformation	Unary Extension to RA	Type of function
Many-to-one	Aggregation	Aggregate Function Set \rightarrow Value
One-to-one	Extended Projection	Scalar Function Value \rightarrow Value
One-to-many	Mapper	Mapper Function Value \rightarrow Set


Mapper example

Extra slide 4

```
1: select
2:   lpad(tostr(ACCT), 4, '0') as ACCTNO,
3:   map AMOUNT, SEQNO
4:   begin
5:     var rem_amnt: numeric
6:     var seq_no: integer = 0
7:     rem_amnt = AMT
8:     loop while rem_amnt > 100 do
9:       rem_amnt = rem_amnt - 100
10:      seq_no = seq_no + 1
11:      insert rem_amnt, seq_no
12:    end loop
13:    if rem_amnt > 0 then
14:      insert rem_amnt, seq_no + 1
15:    end if
16:  end
17: from LOANS
```

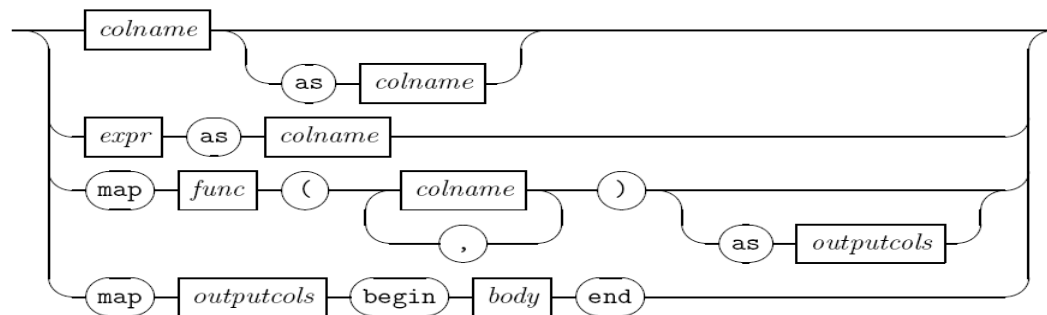
Simplified SELECT syntax for mappers

select




Syntactic details for mapper functions

mapperfunc

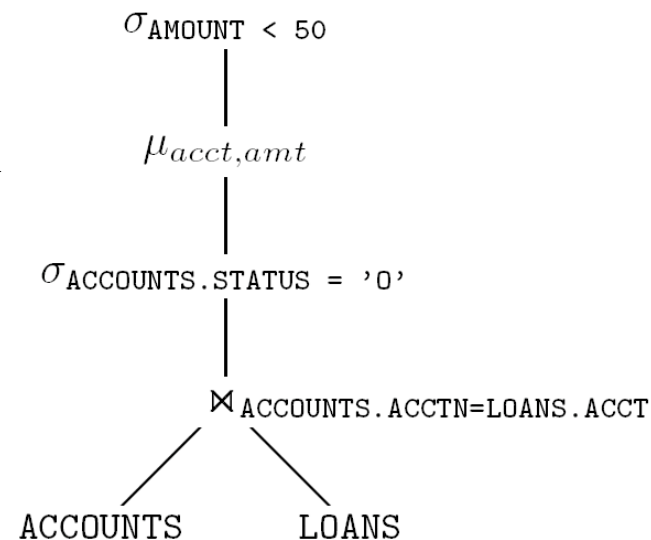



outputcols



Example: Query with the mapper operator

- 1: **select** **map** acct(ACCT) as ACCTNO,
- 2: **map** amt(AM) as AMOUNT
- 3: **from** LOANS, ACCOUNTS
- 4: **where** ACCOUNTS.ACCTN = LOANS.ACCT
- 5: **and** ACCOUNTS.STATUS = 'O'
- 6: **and** AMOUNT < 50



Union

```

1: insert into EVENTS (LOANNO, EVTYP, AMTYP, AMT)
2:   select LOANNO, EVTYP, 'CAPTL' as AMTYP, CAPTL
3:     from LOANEVT
4:    where CAPTL > 0
5:   union all
6:   select LOANNO, EVTYP, 'TAX' as AMTYP, TAX
7:     from LOANEVT
8:    where TAX > 0
9:   union all
10:  select LOANNO, EVTYP, 'EXPNS' as AMTYP, EXPNS
11:    from LOANEVT
12:   where EXPNS > 0
13:  union all
14:  select LOANNO, EVTYP, 'BONUS' as AMTYP, BONUS
15:    from LOANEVT
16:   where BONUS > 0;
  
```

Recursive Query

```

1: with repayments(digits(ACCTNO), AMOUNT, SEQNO, REMAMNT) as
2:   (select ACCT,
3:     case when base.AM < 100 then base.AM
4:     else 100 end,
5:     1,
6:     case when base.AM < 100 then 0
7:     else base.AM - 100 end
8:   from LOANS as base
9:  union all
10:  select ACCTNO,
11:    case when step.REMAMNT < 100 then
12:      step.REMAMNT
13:    else 100 end,
14:    SEQNO + 1,
15:    case when step.REMAMNT < 100 then 0
16:    else step.REMAMNT - 100 end,
17:  from repayments as step
18:  where step.REMAMNT > 0)
19: select ACCTNO, SEQNO, AMOUNT
20: from repayments as PAYMENTS
  
```

Table function

```

1: create function LOANSTOPAYMENTS
2:   return PAYMENTS_TABLE_TYPE pipelined is
3:   ACCTVALUE LOANS.ACCT%TYPE;
4:   AMVALUE LOANS.AM%TYPE;
5:   REMAMNT INT;
6:   SEQNUM INT;
7:   cursor CLOANS is
8:     select * from LOANS;
9: begin
10:  open CLOANS;
11:  loop
12:    fetch CLOANS into ACCTVALUE, AMVALUE;
13:    REMAMNT := AMVALUE;
14:    SEQNUM := 1;
15:    while REMAMNT > 100
16:      loop
17:        pipe row(PAYMENTS_ROW_TYPE(
18:          LPAD(ACCTVALUE, 4, '0'), 100.00, SEQNUM));
19:        REMAMNT := REMAMNT - 100;
20:        SEQNUM := SEQNUM + 1;
21:      end loop
22:    if REMAMNT > 0 then
23:      pipe row(PAYMENTS_ROW_TYPE(
24:        values (LPAD(ACCTVALUE, 4, '0'),
25:          REMAMNT, SEQNUM));
26:    end if
27:  end loop
28: end LOANSTOPAYMENTS

```

An iterator that scans the input relation once

With a nested while loop that inserts pipes multiple several output tuples

Algebraic rewriting rules

Pushdown of projections

$$\pi_Z(\mu_F(s)) = \pi_Z(\mu_{F'}(s))$$

Introduction of projections

$$\mu_F(s) = \mu_F(\pi_N(s))$$

Pushdown of selections

$$\sigma_{C_B}(\mu_F(s)) = \mu_F(\sigma_{C[B_1, \dots, B_m \leftarrow F|_{B_1}, \dots, F|_{B_m}]}(s))$$

Pushing selections to mapper functions

$$\sigma_{C_{A_i}}(\mu_F(s)) = \mu_{F \setminus \{f_{A_i}\} \cup \{\sigma_{C_{A_i}} \circ f_{A_i}\}}(s)$$

Distributing mappers over unions

$$\mu_F(r \cup s) = \mu_F(r) \cup \mu_F(s)$$

Distributing mappers over Cartesian products

$$\mu_F(s \times r) = \mu_{F_S}(s) \times \mu_{F_R}(r)$$

Statistics for computing the cost of One-to-many data transformations

- Given a mapper μ_F
 1. Average selectivity of the mapper (α_F)
 - Ratio of input tuples that are transformed
 2. Average fanout factor of the mapper (O_F)
 - Ratio of output tuples produced for each input tuple
 3. Average mapper function cost (C_F)
 - Cost of evaluating a mapper function

Estimates for cost-based plan selection

- Estimating the cardinality of a mapper expression

$$card(\mu_F(e)) = card(e) \cdot \alpha_F \cdot O_F$$

└──┬──┘ number of input tuples
└──┬──┘ input tuples that are not filtered
└──┬──┘ total number of output tuples

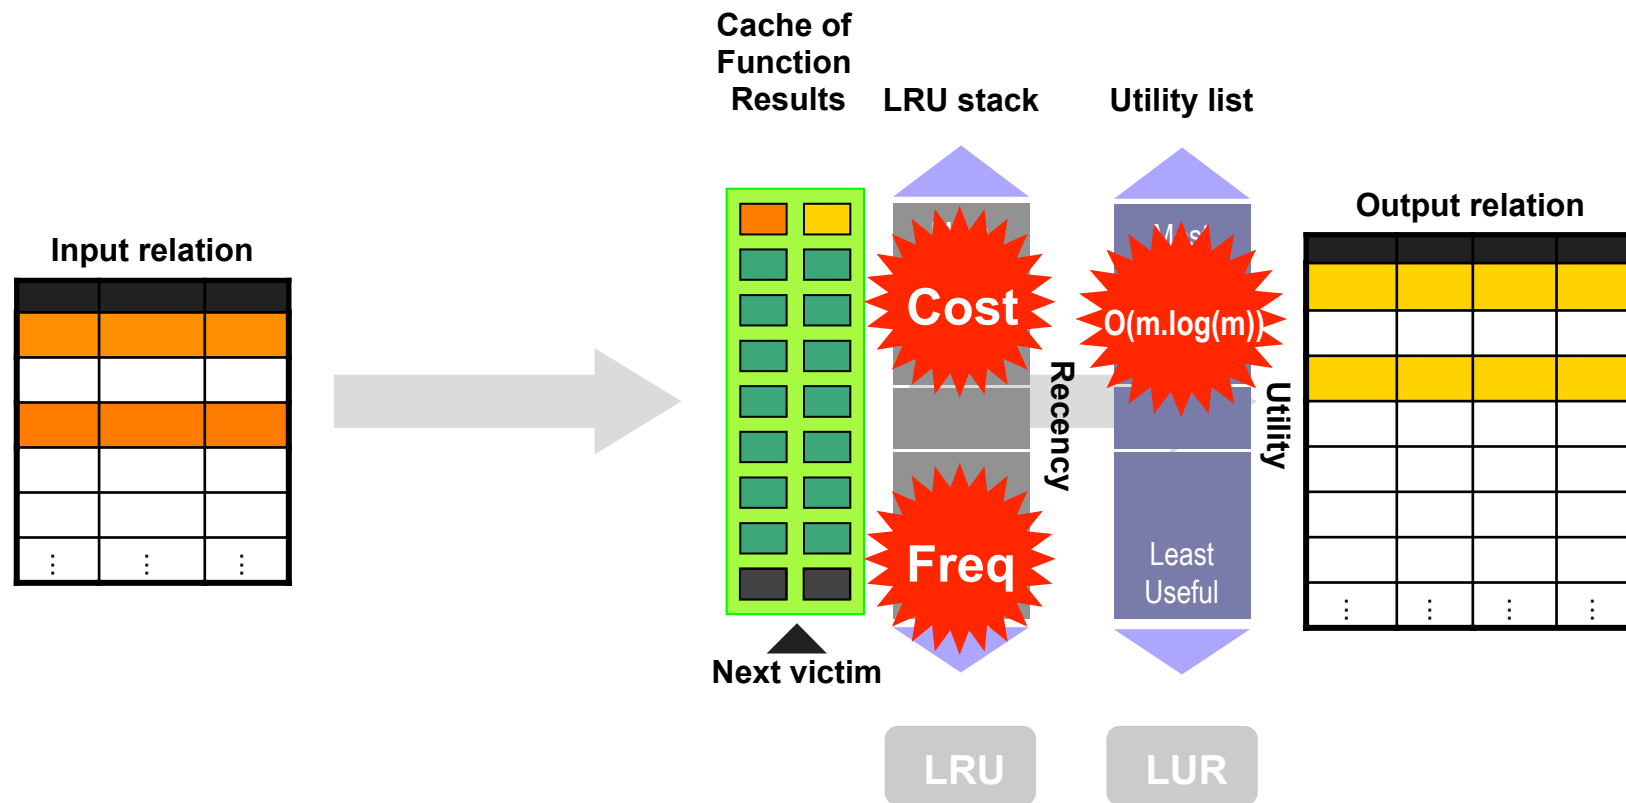
← Fanout factor
 ← Selectivity factor

- Estimating the CPU cost of a mapper expression

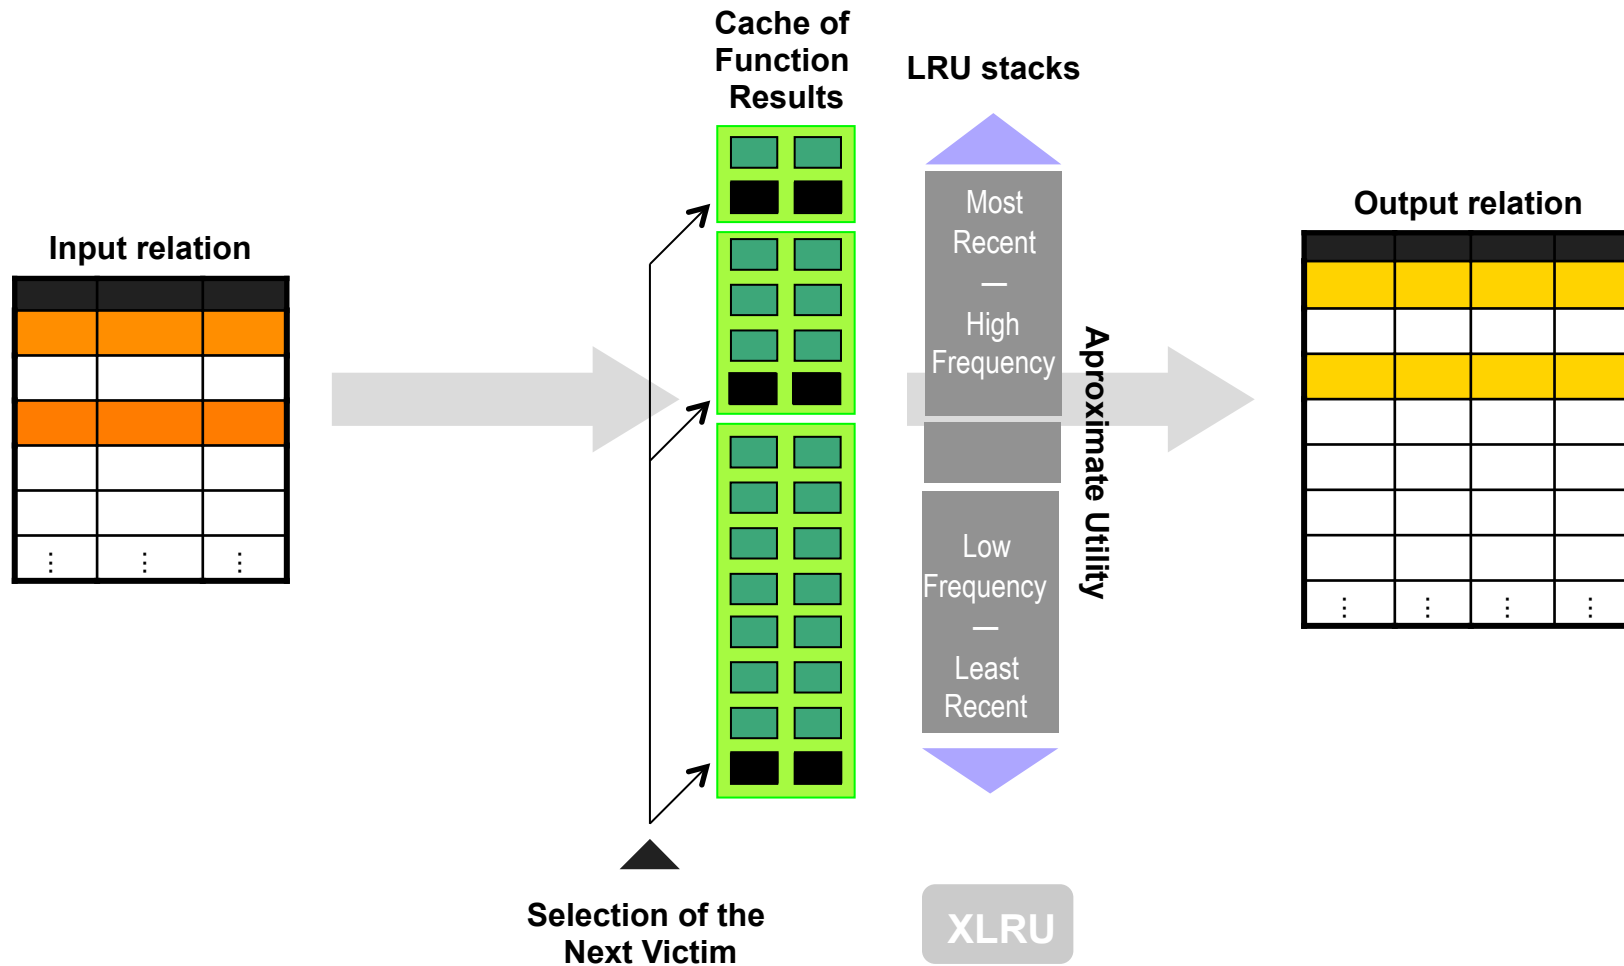
$$cost_{CPU}(\mu_F(e)) = card(e) \cdot (k \cdot O_F + C_F)$$

└──┬──┘ number of input tuples
└──┬──┘ CPU cost of the cartesian product operation
└──┬──┘ CPU cost of evaluating all the mapper functions
└──┬──┘ per-tuple CPU cost
└──┬──┘ total CPU cost


Cache-based evaluation 1/2




Cache-based evaluation 2/2



An utility metric based on statistical inference

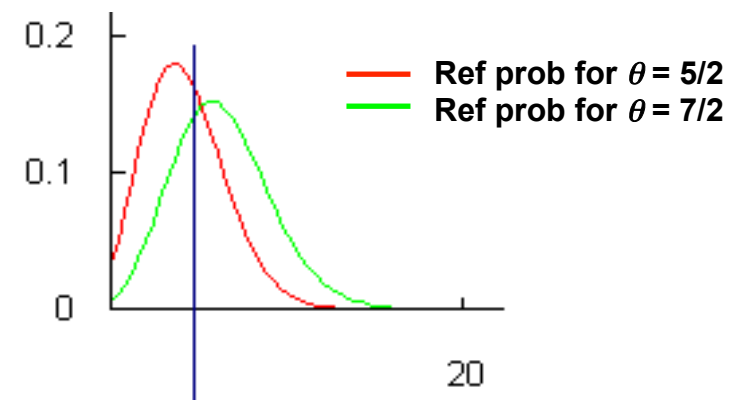


- Which entry to replace?

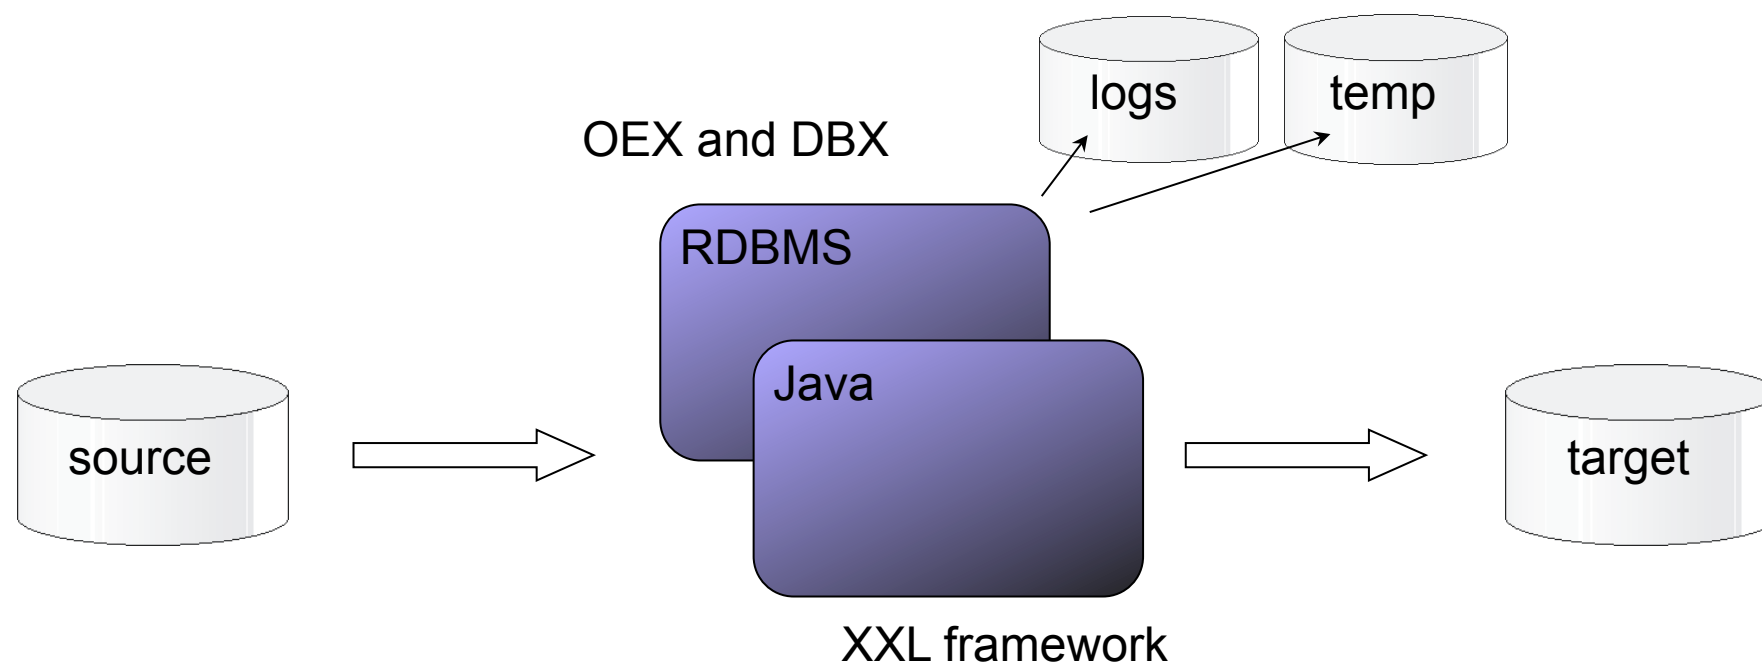


- Utility $u_{t_0}(e)$ based on:
 - Number of past references n_h
 - Cost c
 - Avg frequency θ
 - Time to last $k = (t_0 - t_l)$

- Evolution of $(1-\theta)^k$



Experimental Setup



■ Workload

- `LOANEVT` and `LOANS` synthetic input relations (on raw devices)
- Relations sizes: 100K, 500K, 1M, 2.5M, 5M (in tuples)

Settings checklist

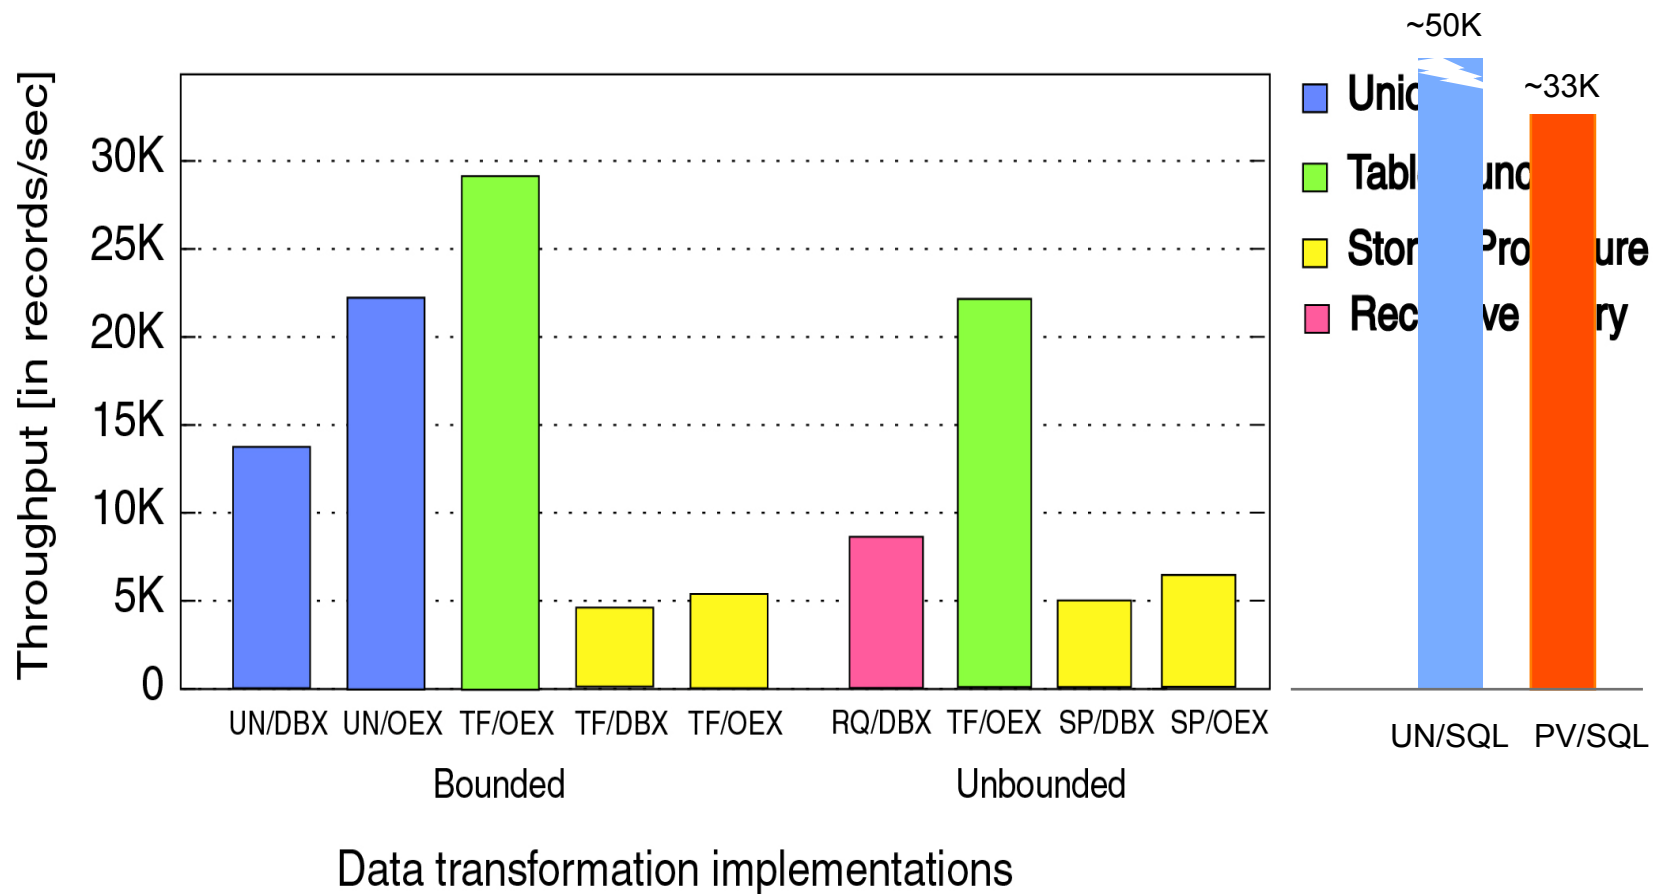
- Same configuration for SGBDs
 - Same block size
 - Same multi-block read count
 - Same size for Buffer pools
 - Same concurrency settings (DBRD and DBWR processes)

- Align the configuration of the SGBDS with Java
 - Same physical conditions for I/O
 - Same record-size
 - Same number of records per page
 - Logging disabled
 - Asynchronous I/O disabled
 - Parallel query execution disabled

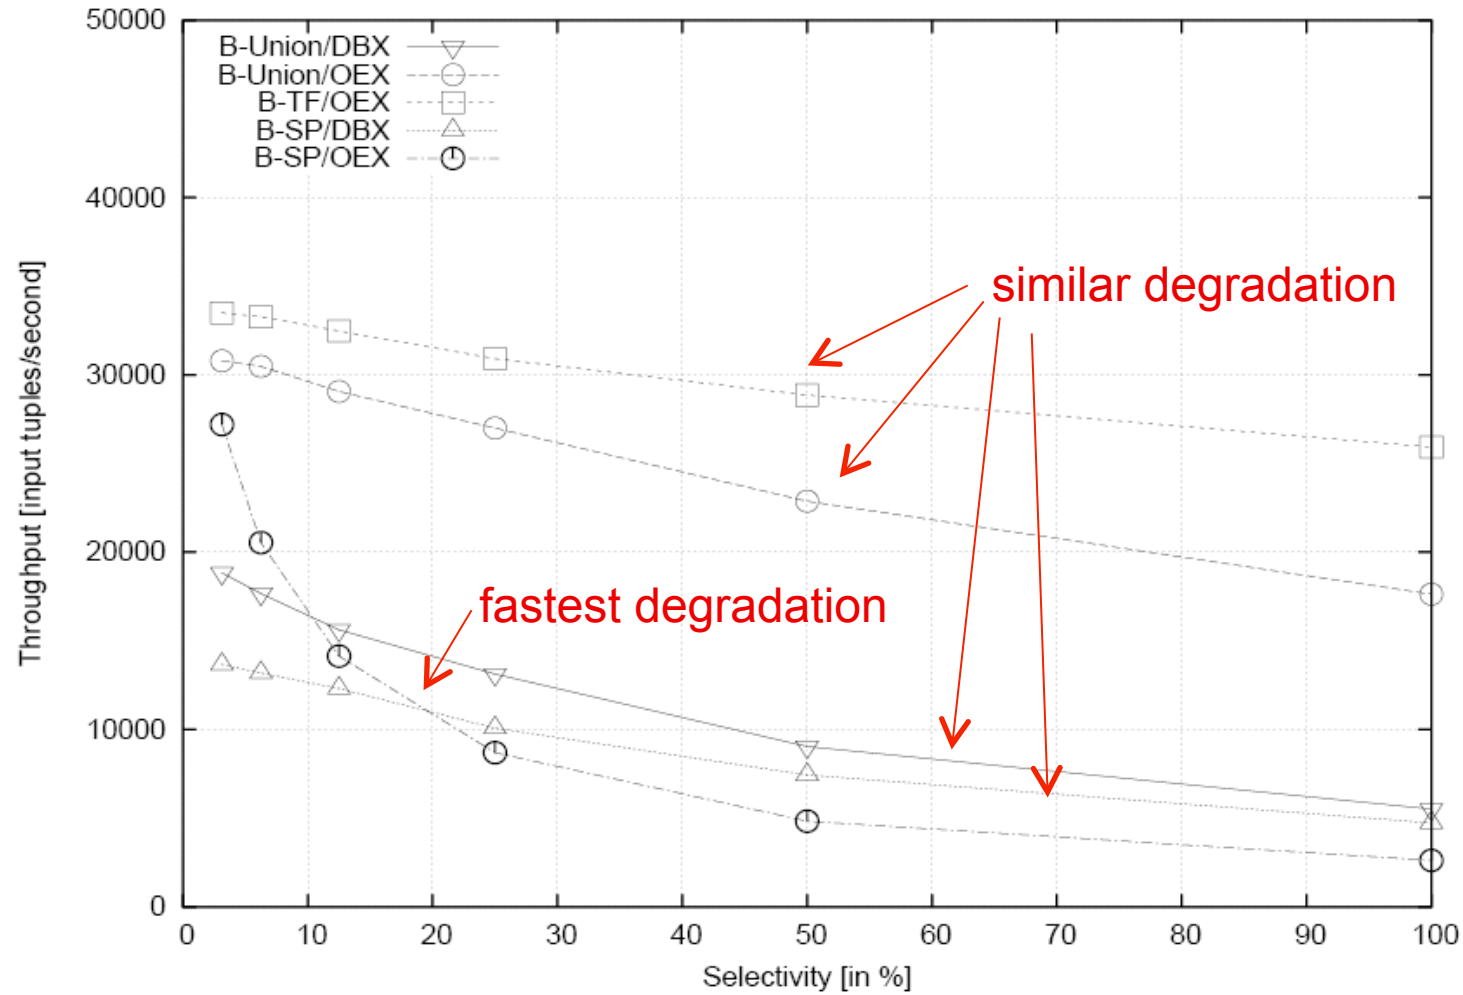
Resource consumption

	Union	Unpivot	RQ	TF
> fanout	more input more output bigger query	input once more output	input once more output more <i>'temp'</i>	input once more output
> select	same input more output	same input more output	same input more <i>'temp'</i> more output	same input more output

Throughput of one-to-many transformations



Influence of selectivity



Original v.s. optimized expression

