

# Interactive Performance Monitoring of a Composite OLTP and OLAP Workload

Anja Bog<sup>1</sup>, Kai Sachs<sup>2</sup>, Hasso Plattner<sup>1</sup>

<sup>1</sup> Hasso Plattner Institute, University of Potsdam, *first.last@hpi.uni-potsdam.de*

<sup>2</sup> SAP AG, Germany, *kai.sachs@sap.com*

2012 ACM SIGMOD/PODS @ Scottsdale, Arizona, USA

## A Benchmark Using a Novel Workload Mix

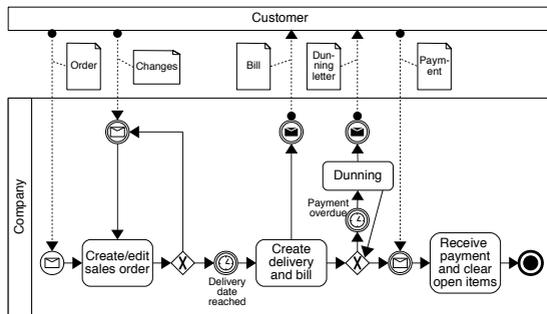
- In need of a benchmark for composite OLTP and OLAP systems to evaluate and compare currently evolving hybrid OLTP/OLAP systems with existing solutions
- Design goals for optimizing OLTP and OLAP systems are in conflict: efficient recording of business events, high throughput for many small read and write transactions vs. analyses of large amounts of data within a single query, complex query designs are counteracted by preparation of data

### Workload Mix

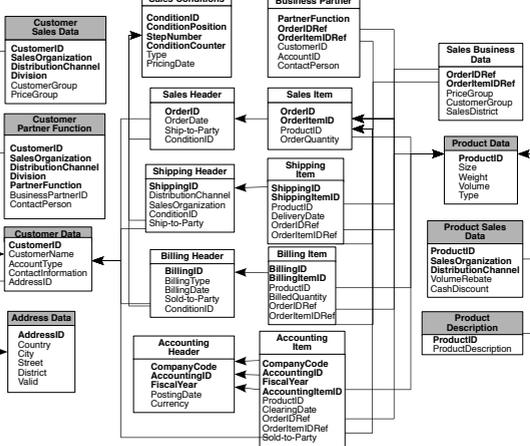
- Additional parameter in a composite benchmark
- Controls the share of OLTP and OLAP-style queries
- Client types: OLTP, OLAP and mixed
- Configurable number of clients for each type

## The Benchmark Schema and Queries

### Scenario



### Database Schema



### Queries

Type	Query	Share	Profile
read/write OLTP	Sales Order	30%	7+ MD* selections, 5+ TD* inserts
	Shipping	27%	2+ TD selections, 2+ TD inserts
	Billing	25%	4+ TD selections, 5+ TD inserts
	Payment	18%	1 TD selection, 1+ TD update
read OLTP	Sales Order by Key	random/configurable	TD Selection on key, header item join
	Sales Order by Period		Range selection, no join
	Open Bills (Items)		Selection by non-key attributes, header item join, order by FK attribute and PK attribute
	Customer Details		MD Selection on key, 2-table join
OLAP	Product Details	random/configurable	MD Selection on key, 2-table join
	Daily Flash		TD aggregate range selection, header item join, group by 3 sales hierarchy attributes, order by aggregate
	Avg. Order Processing Time		TD aggregate select on independent range sub-select (4-table header item join), group by sales hierarchy, order by aggregate
	Order Delivery Fulfillment		TD aggregate range select, 4-table header item join, group by sales hierarchy, order by aggregate, dependent aggregate range sub-select with header item join
	Days Sales Outstanding		TD aggregate range select, header item join, group by currency, order by aggregate, dependent aggregate range sub-select with header item join

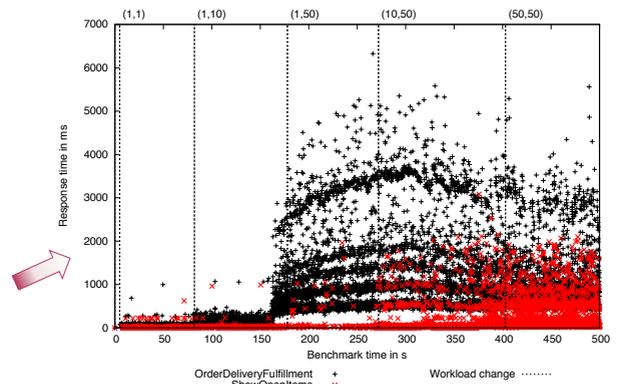
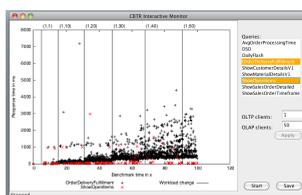
\*MD = master data, TD = transaction data

Schema statistics as taken from a productive system that is the base for the benchmark

- 18 tables: 7 master data, 11 changing with transactions
- 2632 columns, the smallest table with 5, the largest with 327 columns

## The Interactive Performance Monitoring Tool

- Interactively monitor the impact of live workload changes on response time
- Change client configuration: number of OLTP and OLAP clients running concurrently
- Switch between display of different queries or select multiple queries to monitor



## References

A Framework for Simulating Combined OLTP and OLAP Workloads. A. Bog, M. Domschke, J. Mueller, A. Zailer. 16th International Conference on Industrial Engineering and Engineering Management (IE&EM), China, 2009.

A mixed transaction processing and operational reporting benchmark. A. Bog, H. Plattner, A. Zailer. Information Systems Frontiers Journal, Springer, pp. 1-15, 2010.

Benchmarking database design for mixed OLTP and OLAP workloads. A. Bog, K. Sachs, H. Plattner. 2nd International Conference on Performance Evaluation & Benchmarking (TPCTC), USA, 2011.

Normalization in a mixed OLTP and OLAP workload scenario. A. Bog, K. Sachs, H. Plattner. 3rd Technology Conference on Performance Evaluation & Benchmarking (TPCTC), USA, 2011.