Hasso Plattner Institut

IT Systems Engineering | Universität Potsdam

A Case for Online Mixed Workload Processing

Jens Krüger, Christian Tinnefeld, Martin Grund, Alexander Zeier, Hasso Plattner

Introduction

- Data model acted as main guideline for DBMS development
- For a long time: few DBMSs for many different applications
- Recent DBMS discussions: heavily application oriented
  - Two examples:
    - Stonebraker et al– One size fits all
    - Vogels et al- VLDB Keynote 2007 / Amazon Dynamo
- DBMS for Enterprise Applications
  - Beginning of 1990s separation into OLTP and OLAP
  - But, we claim:

OLTP-style workloads also require the ability to frequently compute OLAP-style aggregate queries

Enterprise Applications

1. Company estimates future demand for its products –
   Demand Planning / Supply Chain Management
- Customers contact the company and place orders –
  Sales Order Processing / Enterprise Resource Planning
- The availability of the requested products have to be checked –
  Available-to-Promise / Supply Chain Management
- Customers fall behind on their payments –
  Dunning / Financial Accounting
- The company wants to analyze its sales performance –
  Sales Analytics / Enterprise Resource Planning

Demand Planning

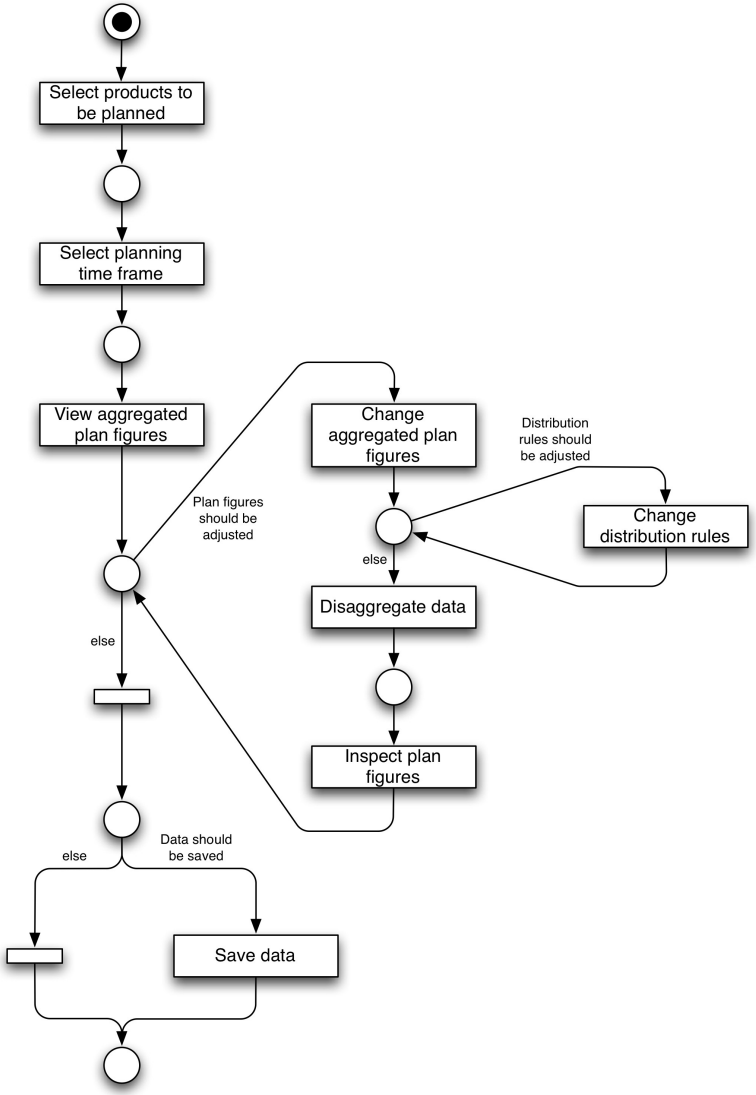

Available-to-Promise

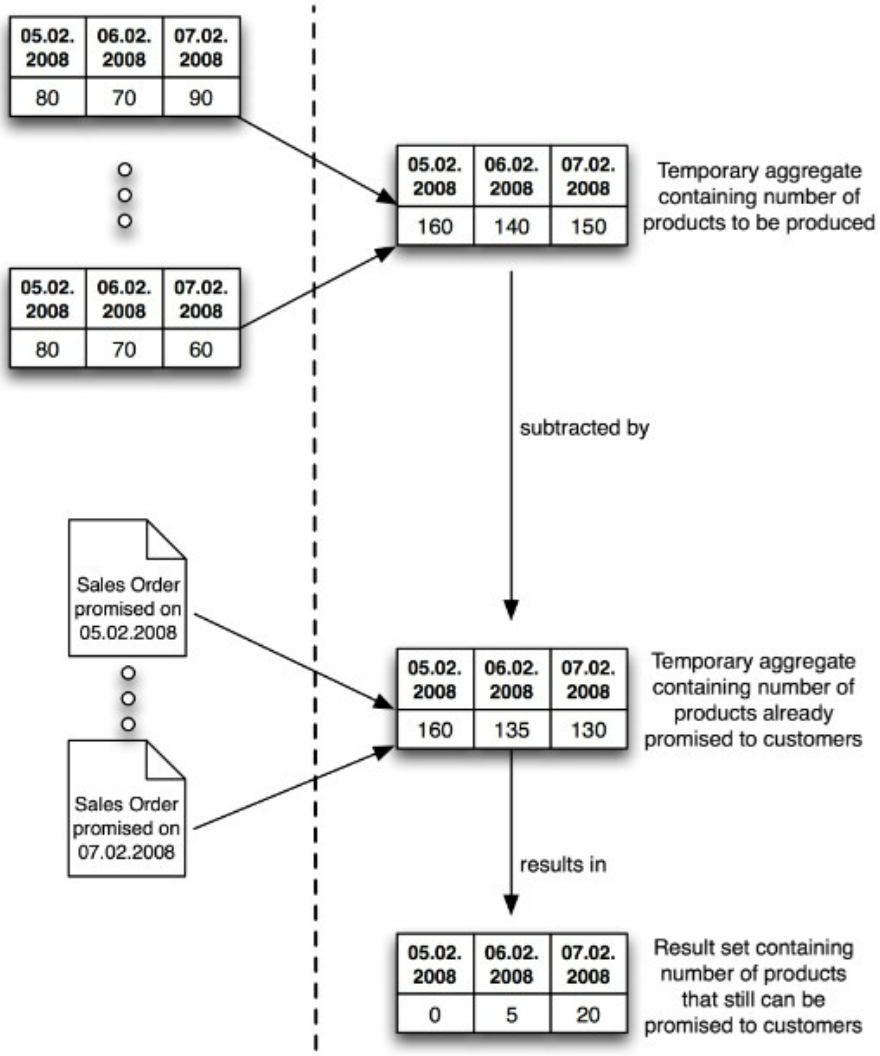

Dunning

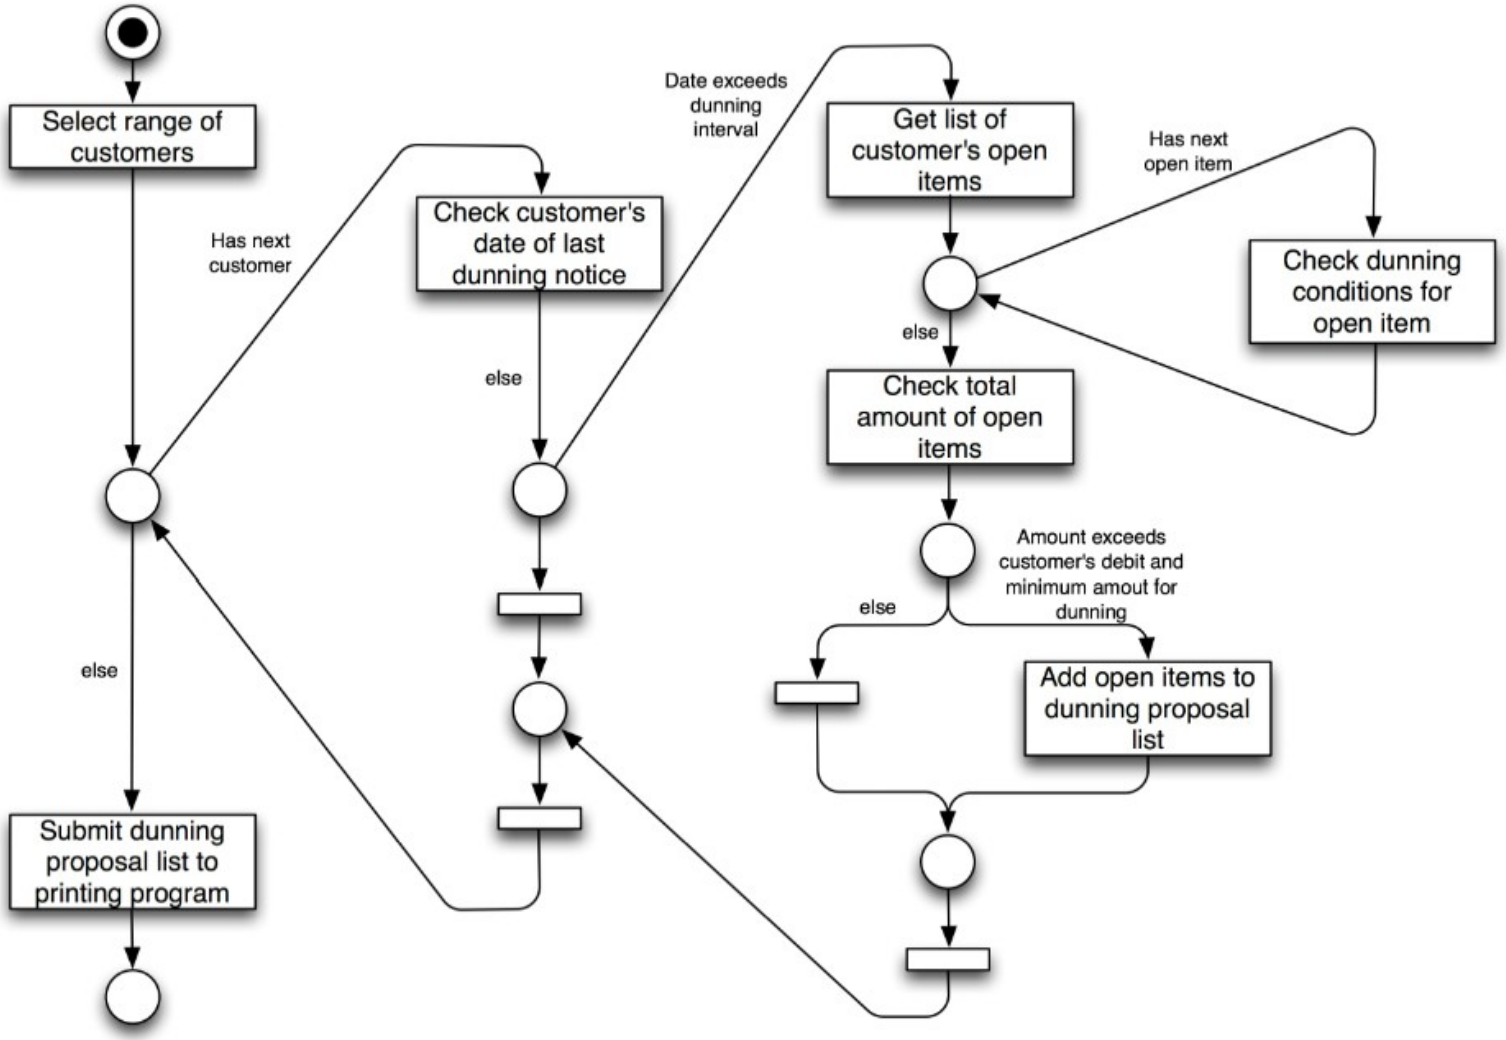
Select range of customers
Has next customer
Check customer's date of last dunning notice
Date exceeds dunning interval
Get list of customer's open items
Has next open item
Check dunning conditions for open item
else
Check total amount of open items
else
Amount exceeds customer's debit and minimum amout for dunning
else
Add open items to dunning proposal list
else
Submit dunning proposal list to printing program

Comparison of Application Characteristics

	Demand Planning	Sales Order Processing	Available to Promise	Dunning	Sales Analysis
Granularity of Data	Transactional	Transactional	Transactional	Transactional	Pre-Aggregated
Operations on Data	Read & Write	Read & Write	Read & Write	Read & Write	Read-Only
Preprocessing of Data	No	No	No	No	Yes
Timeframe of Data	Historical & Recent	Recent Only	Historical & Recent	Historical & Recent	Historical & Recent
Update Cycles of Data	Always Up-to-Date	Always Up-to-Date	Always Up-to-Date	Always Up-to-Date	Cyclic Updates
Amount of Data per Query	Large	Small	Large	Large	Large
Complexity of Queries	High	Standard	High	High	High
Predictability of Queries	Medium	High	Medium	Medium	Low
Response Time of Queries	Seconds	Seconds	Seconds	Seconds to Hours	Seconds to Hours

OLTP Characteristics are colored light grey

OLAP Characteristics are colored dark grey

Online Mixed Workload Processing

- Large amount of data is needed to perform transactional query
- Nothing new, mixed workloads are a well established topic
- But: here they are originated by a single application.

Granularity of Data	Transactional
Operations on Data	Read & Write
Preprocessing of Data	No
Timeframe of Data	Historical & Recent
Update Cycles of Data	Always Up-to-Date
Amount of Data per Query	Large
Complexity of Queries	High
Predictability of Queries	Medium
Response Time of Queries	Seconds

Conclusion

- Not only Web 2.0 companies need application-specific data stores
- Mixed workloads are originated by a single application
- Huge potential for increasing performance and functionality of enterprise applications by supporting mixed workload
- Future Work
  - Need for a Mixed Workload Benchmark
    - TPC-E and TPC-H claim one part of the process separately
    - But, OLTP-style workloads also require the ability to frequently compute OLAP-style aggregate queries
  - DBMS Draft for Mixed Workloads
    - Read-optimized, in-memory columnar store with transactional support

Backup Slides

Application Characteristics at Amazon

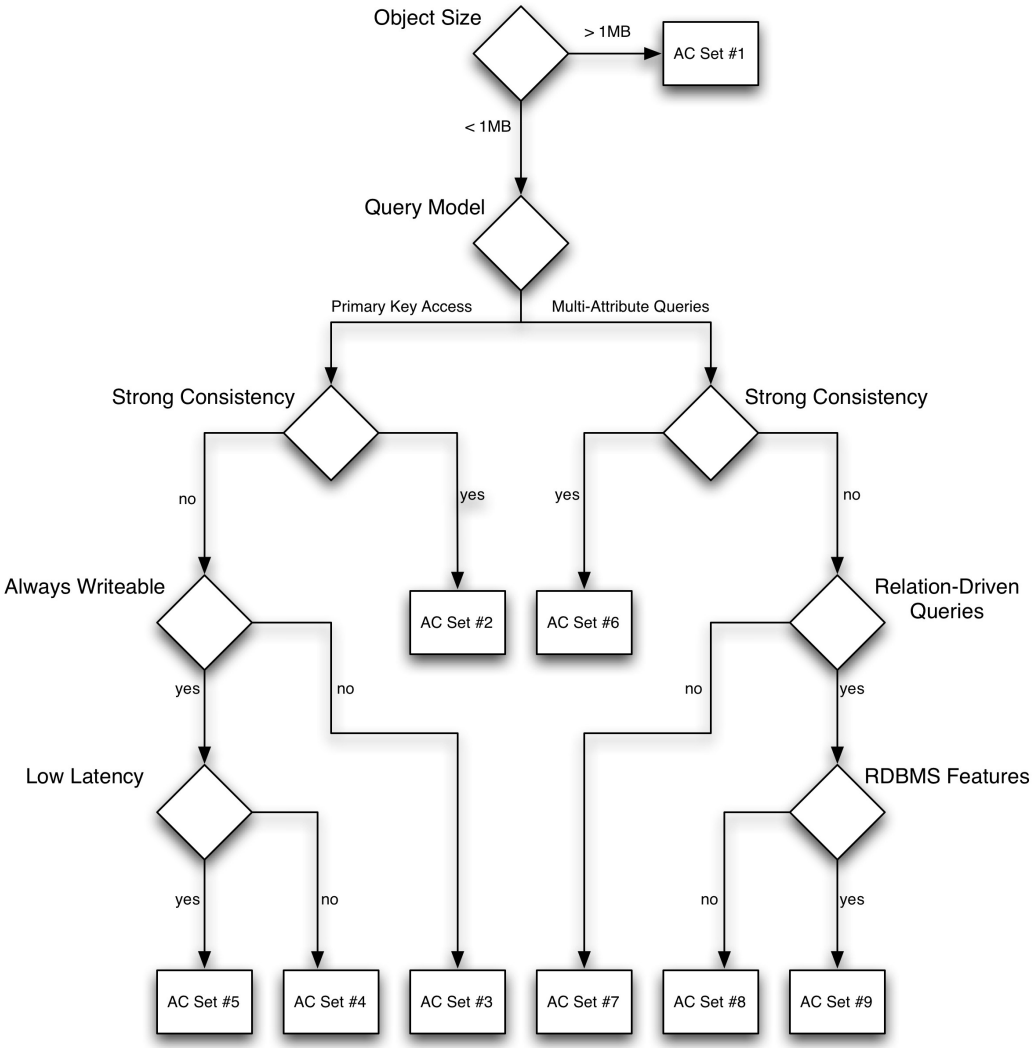

Sales Order Processing Data Logs