

Prof. Hasso Plattner



A Course in In-Memory Data Management

The Inner Mechanics
of In-Memory Databases

October 12, 2015

This learning material is part of the reading material for Prof. Plattner's online lecture "In-Memory Data Management" taking place at www.openHPI.de. If you have any questions or remarks regarding the online lecture or the reading material, please give us a note at openhpi-imdb@hpi.uni-potsdam.de. We are glad to further improve the material.

Research Group "Enterprise Platform and Integration Concepts",
<http://epic.hpi.uni-potsdam.de>

Examples for New Enterprise Applications

This chapter gives an overview of exemplary enterprise applications that build on the changed foundations and database techniques described in this book. We will first describe how in-memory technology can be leveraged for the interactive analysis of point-of-sales data, followed by an example application for high performance in-memory genome analyses.

35.1 Interactive Analyses of Point-of-Sales Data

Retailers face not only the challenge of consolidating all the data generated by electronic point-of-sale (POS) terminals, but also to leverage the data to derive business value. Especially when the data is stored at its finest granularity recording the actual transactions with all their items, processing becomes a challenge. In a prototypical implementation described in [SFKP14], we show how in-memory technology can help to analyze POS data and how it enables new types of enterprise applications. We show that it is possible to interactively explore the transactional data set without precomputing analytical summaries while providing users with full flexibility. We present a prototypical application for interactive analyses and exploration of 8 billion records of real data from a large retail company with sub-second response times.

35.1.1 *Motivation*

Information derived from POS data can be used for the planning and verification of sales promotions, cross-selling, up-selling, loyalty programs or product placement inside stores.

As an example, we will focus on the planning of sales promotions. Consumer sales promotions are an essential part of a retailers business, as the frequency of promotions increased significantly over the last years through competition and price pressure. Nowadays, up to 20% of the total revenue of German retailers are generated through promotions [SAP11]. When planning promotions, extended analyses based on the POS data are taken into account.

However, the amount of data and the separation of today's analytical and transactional systems make it difficult to perform these analyses and limit the range of questions that can be answered with the analytical system. In the following, we outline the advantages of in-memory technology and how it can be used for interactive analyses of the transactional POS data.

35.1.2 POS Explorer

The speed and flexibility provided by the technical foundation of in-memory database technology enables the creation of completely new tools for enterprises. To illustrate the potentials of in-memory technology, we built the POS Explorer application for the interactive analysis of POS data as a proof of concept to outline the potentials of such applications.

The application is designed to support the planning and verification of promotions, while pursuing the following goals. The application should (a) work directly on the POS transactions as the finest level of granularity, (b) enable interactive access to the data with sub-second response times, (c) provide full flexibility to users and not limit queries due to missing pre-aggregated information and (d) leverage a modern javascript and HTML based user-interface for access in browsers and mobile devices.

The main functionality provided by the application is (A) a hierarchical product browser with live calculated key-figures and full filter flexibility, (B) an interactive basket analysis for selected products and (C) an interactive heat-map analysis of sales and revenue.

35.1.2.1 Transactional Data Model and Real World Data Volume

The paradigm of analyzing the transactional data as it is recorded on the level of single transactions is a fundamental design decision and in stark contrast to current approaches [TLN⁺11]. Typically, enterprises use IT-landscapes with a variety of systems and replicate data between those systems while transforming data and pre-calculating aggregated information [BFH⁺11]. The results are then used in analytical systems or for the generation of reports summarizing the information. However, it is expensive, complicated and slow to establish additional reports if the ETL process has to be adapted.

Therefore, all operations in our demo application work directly on the transactional data model, which stores single POS terminal transactions. The customer transactions are stored as a set of items, which are connected through unique transaction identifiers. Every item contains information about the specified product, the price, the sold quantity, the total revenue and if the product was sold during a promotion. Each receipt has additional information, such as the store, date and time of the transaction.

The application was developed based on a real data set from a large retail company, containing 8 billion line items from sales over two years, 500,000 products and 5,000 stores. The uncompressed size of the dataset is roughly one terabyte. As a back of the envelope calculation and assuming a hard disk bandwidth of 100 megabytes per second per disk, one would need about 10,000 disks to scan the whole dataset in a traditional disk-based row store in one second. However, the prototype application runs on a cluster of 25 nodes with 1,000 processor cores and 25 terabytes of main memory.¹ The processing power and level of parallelization of the cluster correspond to the hardware that we expect to be generally available on a single, affordable high-end computer in a few years.

35.1.2.2 On the Fly Calculation of Key Figures

As an example how the flexibility provided by in-memory technology can be leveraged, the demo application allows users to browse through products while calculating common key figures on the fly. Figure 35.1 shows a screenshot from the demo application's hierarchical product browser with the key figures relative market share, generated revenue and the basket value as well as their development over time.

The relative market share indicates the percentage of a product's sales within its product group. It is calculated by summing up all sales of products in a certain product group and comparing it with the sales of each product within that group to calculate the relative share. The generated revenue is calculated by the sum of all sales of the respective product. The basket value indicates the average value of transactions that contain a product and is calculated by the average basket size and basket value. This allows users to quickly identify products or product groups that drive revenue as the basket value is often a good indicator for products to promote.

While the user navigates through the product hierarchy or directly searches for a product, every click results in multiple queries scanning the complete data set and calculating the requested numbers. As everything is always calculated directly on the POS data and no aggregates are used, the data can be filtered on any attribute or combination of attributes, for example the size of stores, periods of time, product groups or areas. In the application,

¹ Hosted at the Future SOC Lab at the Hasso Plattner Institute.
http://www.hpi-web.de/future_soc_lab




Fig. 35.1: The prototypical application enables an interactive exploration of the transactional POS data set with 8 billion records. The screenshot shows the entry point of the tool, presenting a hierarchical product browser with key figures for every product and their development over time. Each selection step within the product hierarchy aggregates all transactions and is done interactively with response times below one second.

these filters can be adjusted directly by the user through the optional menu at the top, providing a fixed set of filters. However, it would also be possible to directly expose the complete data schema to users, allowing them to filter on any available attribute.

Despite the huge amount of data and without using pre-aggregated summaries, every query is processed with sub-second response times. The application allows to browse the presented products interactively and users can enter into a dialog with the system to quickly retrieve answers and to articulate follow up questions. This enables the exploration of data in a try and error fashion, assisting in the discovery of new insights.

35.1.2.3 Heat-Map Analysis

The calculation of key figures on the actual transactions allows for new types of analyses that take the additional information into account. As an example, we implemented a heat map visualization over the days of the week and the hours of the day, as shown in Figure 35.2. This particular visualization is useful for the verification of promotions. It is an example of data analyses




Fig. 35.2: Intuitive visualization of a heat map analysis of a product's revenue over the days and hours of the week. The analysis leverages the granularity of data at the transaction level and provides interactivity through sub-second query processing even on 8 billion records.




Fig. 35.3: Intuitive visualization of the result of a basket analysis, whereas the circle visualizes products that are most commonly sold together with the product of interest and the size of the area indicates the average basket values.

that are only feasible with fast access to the raw POS data. We can visualize hourly aggregates for the common sales of any two products, while allowing users to interactively filter the considered time-range, store-size and region without the need for pre-aggregated cubes that would lead to immense storage and maintenance costs.

35.1.2.4 Basket Analysis

A basket analysis determines relationships of products that are most commonly sold together. Instead of using old traditional list views, Figure 35.3 shows a new intuitive visualization of the result of a basket analysis. The circle visualizes the relationships of the most commonly sold products with a product of interest, whereas the size of the area can be used to visualize another dimension as for example the average basket size or value. Like the calculation of key figures, the basket analysis runs interactively and can be executed on any arbitrarily filtered data set. The quality of a rule $A \rightarrow B$ can be drawn from its *support* and *confidence* values. The support value refers to the number of transactions that contain product A , whereas the confidence value reflects the number of transactions that contain product A and B .

Since the receipts are distributed across the computing cluster based on the transaction identifier, the database system can evaluate the intermediate results on each partition and aggregate only the counts, thereby keeping the network traffic between the computing-nodes low and resulting in low response times.

35.1.2.5 Advantages over Traditional Approaches

In traditional analytical systems the set of possible questions and their detail level are defined at the implementation and execution of the Extract-Transform-Load (ETL) process. The trade-off is between a level of detail and storage and processing requirements for creating the data. As a result, the cubes contain the data up to a certain detail level in all dimensions. With a product hierarchy of 500,000 products, any cube that cross-relates all of these products would contain 250 billion data fields for each dimension. In this setting, storing a relation between two products on a monthly level would require 3 trillion data fields for one year. Due to these effects, today's basket analyses have to focus on subsets of products, e.g. to analyze the products of a few product groups at a time.

The main point of our solution is not that the presented analyses were not possible before, but that users are now able to perform these analyses interactively with sub-second response times even on large amounts of data and that the user has full flexibility and can filter and aggregate data on the fly as needed.

35.2 References

- [BFH⁺11] Brigitte Bläser, Karl Fleckstein, Stefan Hoffmann, Oliver Schall, and Olaf Rutz. SAP Enterprise Data Warehouse for Point of Sales

- Data Optimized for IBM DB2 for Linux, UNIX, and Windows on IBM Power Systems. *SAP SCN, DOC-14457*, 2011.
- [SAP11] SAP AG und GfK Panel Services Deutschland. Professionalisierung des Promotionsmanagements im deutschen LEH. *Eine Gemeinschaftsstudie von GfK und SAP*, 2011.
- [SFKP14] D. Schwalb, M. Faust, J. Krueger, and H. Plattner. Leveraging In-Memory Technology for Interactive Analyses of Point-of-Sales Data. *Workshop on Big Data Customer Analytics in Conjunction with ICDE*, 2014.
- [TLNJ11] T J Teorey, S S Lightstone, T Nadeau, and H V Jagadish. *Database Modeling and Design: Logical Design, Fifth Edition*. The Morgan Kaufmann Series in Data Management Systems. Elsevier Science, 2011.

