

---

## **Materialized Views**

---

# Materialized Views

---

**Rada Chirkova**

*North Carolina State University*

*USA*

*chirkova@csc.ncsu.edu*

**Jun Yang**

*Duke University*

*USA*

*junyang@cs.duke.edu*

**now**

the essence of **know**ledge

Boston – Delft

## Foundations and Trends<sup>®</sup> in Databases

*Published, sold and distributed by:*

now Publishers Inc.  
PO Box 1024  
Hanover, MA 02339  
USA  
Tel. +1-781-985-4510  
[www.nowpublishers.com](http://www.nowpublishers.com)  
[sales@nowpublishers.com](mailto:sales@nowpublishers.com)

*Outside North America:*

now Publishers Inc.  
PO Box 179  
2600 AD Delft  
The Netherlands  
Tel. +31-6-51115274

The preferred citation for this publication is R. Chirkova and J. Yang, Materialized Views, Foundation and Trends<sup>®</sup> in Databases, vol 4, no 4, pp 295–405, 2011

ISBN: 978-1-60198-622-1  
© 2012 R. Chirkova and J. Yang

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, mechanical, photocopying, recording or otherwise, without prior written permission of the publishers.

Photocopying. In the USA: This journal is registered at the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923. Authorization to photocopy items for internal or personal use, or the internal or personal use of specific clients, is granted by now Publishers Inc for users registered with the Copyright Clearance Center (CCC). The 'services' for users can be found on the internet at: [www.copyright.com](http://www.copyright.com)

For those organizations that have been granted a photocopy license, a separate system of payment has been arranged. Authorization does not extend to other kinds of copying, such as that for general distribution, for advertising or promotional purposes, for creating new collective works, or for resale. In the rest of the world: Permission to photocopy must be obtained from the copyright owner. Please apply to now Publishers Inc., PO Box 1024, Hanover, MA 02339, USA; Tel. +1-781-871-0245; [www.nowpublishers.com](http://www.nowpublishers.com); [sales@nowpublishers.com](mailto:sales@nowpublishers.com)

now Publishers Inc. has an exclusive license to publish this material worldwide. Permission to use this content must be obtained from the copyright license holder. Please apply to now Publishers, PO Box 179, 2600 AD Delft, The Netherlands, [www.nowpublishers.com](http://www.nowpublishers.com); e-mail: [sales@nowpublishers.com](mailto:sales@nowpublishers.com)

**Foundations and Trends<sup>®</sup> in  
Databases**

Volume 4 Issue 4, 2011

**Editorial Board**

**Editor-in-Chief:**

**Joseph M. Hellerstein**

*Computer Science Division*

*University of California, Berkeley*

*Berkeley, CA*

*USA*

*hellerstein@cs.berkeley.edu*

**Editors**

Anastasia Ailamaki (EPFL)

Michael Carey (UC Irvine)

Surajit Chaudhuri (Microsoft Research)

Ronald Fagin (IBM Research)

Minos Garofalakis (Yahoo! Research)

Johannes Gehrke (Cornell University)

Alon Halevy (Google)

Jeffrey Naughton (University of Wisconsin)

Christopher Olston (Yahoo! Research)

Jignesh Patel (University of Michigan)

Raghu Ramakrishnan (Yahoo! Research)

Gerhard Weikum (Max-Planck Institute)

## Editorial Scope

**Foundations and Trends<sup>®</sup> in Databases** covers a breadth of topics relating to the management of large volumes of data. The journal targets the full scope of issues in data management, from theoretical foundations, to languages and modeling, to algorithms, system architecture, and applications. The list of topics below illustrates some of the intended coverage, though it is by no means exhaustive:

- Data Models and Query Languages
- Query Processing and Optimization
- Storage, Access Methods, and Indexing
- Transaction Management, Concurrency Control and Recovery
- Deductive Databases
- Parallel and Distributed Database Systems
- Database Design and Tuning
- Metadata Management
- Object Management
- Trigger Processing and Active Databases
- Data Mining and OLAP
- Approximate and Interactive Query Processing
- Data Warehousing
- Adaptive Query Processing
- Data Stream Management
- Search and Query Integration
- XML and Semi-Structured Data
- Web Services and Middleware
- Data Integration and Exchange
- Private and Secure Data Management
- Peer-to-Peer, Sensornet and Mobile Data Management
- Scientific and Spatial Data Management
- Data Brokering and Publish/Subscribe
- Data Cleaning and Information Extraction
- Probabilistic Data Management

### Information for Librarians

Foundations and Trends<sup>®</sup> in Databases, 2011, Volume 4, 4 issues. ISSN paper version 1931-7883. ISSN online version 1931-7891. Also available as a combined paper and online subscription.

## Materialized Views

Rada Chirkova<sup>1</sup> and Jun Yang<sup>2</sup>

<sup>1</sup> *North Carolina State University, Department of Computer Science,  
Raleigh, North Carolina, 27695-8206, USA, [chirkova@csc.ncsu.edu](mailto:chirkova@csc.ncsu.edu)*

<sup>2</sup> *Duke University, Department of Computer Science, Durham, North  
Carolina, 27708-0129, USA, [junyang@cs.duke.edu](mailto:junyang@cs.duke.edu)*

### Abstract

Materialized views are queries whose results are stored and maintained in order to facilitate access to data in their underlying base tables. In the SQL setting, they are now considered a mature technology implemented by most commercial database systems and tightly integrated into query optimization. They have also helped lay the foundation for important topics such as information integration and data warehousing. This monograph provides an introduction and reference to materialized views. We cover three fundamental problems: (1) maintaining materialized views efficiently when the base tables change, (2) using materialized views effectively to improve performance and availability, and (3) selecting which views to materialize. We also point out their connections to a few other areas in database research, illustrate the benefit of cross-pollination of ideas with these areas, and identify several directions for research on materialized views.

## Contents

---

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>Maintaining Materialized Views</b>	<b>9</b>
2.1	Algorithmizing and Implementing Maintenance	13
2.2	Information Available to Maintenance	22
2.3	Materialization Strategies	28
2.4	Timing of Maintenance	33
2.5	Other Issues of View Maintenance	36
<b>3</b>	<b>Using Materialized Views</b>	<b>43</b>
3.1	Background and Theory	43
3.2	Incorporation into Query Optimization	52
3.3	Using Views for Data Integration	54
<b>4</b>	<b>Selecting Views to Materialize</b>	<b>59</b>
4.1	View Selection to Speed Up Queries	61
4.2	Implementation in Commercial Database Systems	71
<b>5</b>	<b>Connections to Other Problems</b>	<b>75</b>
<b>6</b>	<b>Conclusion and Future Directions</b>	<b>83</b>
	<b>References</b>	<b>87</b>

# 1

---

## Introduction

---

Materialized views are a natural embodiment of the ideas of *pre-computation* and *caching* in databases. Instead of computing a query from scratch from base data, a database system can use results that have already been computed, stored, and maintained. The ability of materialized views to speed up queries benefit most database applications, ranging from traditional querying and reporting to web database caching [255], online analytical processing [89], and data mining [208, 367]. By reducing dependency on the availability of base data, materialized views have also laid much of the foundation for information integration [45, 139, 270] and data warehousing [89, 242] systems. Because of their wide applicability, materialized views are a well-studied topic with both a rich research literature and mature commercial implementations. Our goal of this monograph is to provide an accessible introduction and reference to this topic, explain its core ideas, highlight its recent developments, and point out its sometimes subtle connections to other research topics in databases.

**Background** A database *view* is defined by a query. When evaluated, this *view definition query* returns the contents of the view. Database

## 2 Introduction

users can pose queries or define other views over views just as they can over regular database tables. Conceptually, when the database system executes a query involving views, references to views are replaced by their definition queries (recursively, if a view is defined using other views), yielding a final query involving only regular “base tables.”

---

**Example 1.1 (Adapted from [188]).** Consider a retailer with multiple stores across the globe. The stores are grouped into multiple geographic regions for administrative and accounting purposes. The retailer consolidates its inventory and sales information across all stores into a single relational database for auditing and analysis purposes. Consider some of the tables in such a database and their cardinality:

- `pos(itemID, storeID, date, qty, price)`, with one billion ( $10^9$ ) rows, records point-of-sale transactions. There is one row for every item sold in a transaction, with the ID of item, the ID of the store selling it, the date of sale, the quantity sold, and the unit price.
- `stores(storeID, city, region)`, with 100 rows, records information about each store: namely, its ID, city, and region.
- `items(itemID, name, category, cost)`, with 50,000 rows, records information about each item: namely, its ID, name, product category, and cost per unit.

The following view, defined over `pos`, computes the total sales revenue generated for each item by each store:

```
CREATE VIEW TotalByItemStore(itemID,storeID,total) AS
  SELECT itemID, storeID, SUM(qty*price)
  FROM pos GROUP BY itemID, storeID;
```

Suppose a business analyst wants to know the total revenue generated by each store for each item category. This query can be written against the above view as:

```
SELECT storeID, category, SUM(total)           -- (Q1v)
FROM TotalByItemStore, items
```

```
WHERE TotalByItemStore.itemID = items.itemID
GROUP BY storeID, category;
```

When evaluating this query, the database system conceptually expands it to the following equivalent query involving only base tables:

```
SELECT storeID, category, SUM(qty*price)           -- (Q1)
FROM pos, items
WHERE pos.itemID = items.itemID
GROUP BY storeID, category;
```

---

Traditionally, views are “virtual” — the database system stores their definition queries but not their contents. Virtual views are often used to control access and provide alternative interfaces to base tables. They also support *logical data independence*: when the base table schema changes, views can be redefined to use the new schema, so application queries written against these views will continue to function.

Over the years, however, the concept and practice of *materialized views* have steadily gained importance. We materialize a view by storing its contents (though many cases call for alternative materialization strategies; see Section 2.3). Once materialized, a view can facilitate queries that use it (or can be rewritten to use it), when the base tables are expensive or even unavailable for access.

---

**Example 1.2.** Continuing with Example 1.1, suppose we materialize the view `TotalByItemStore`. Now, query (Q1v) can be evaluated by joining `items` with the materialized contents of `TotalByItemStore`. This evaluation strategy is more efficient than joining `items` and `pos`, because `TotalByItemStore` has up to  $50,000 \times 100 = 5 \times 10^6$  rows, compared with  $10^9$  rows in `pos`.

Although (Q1) is not originally written over `TotalByItemStore`, it is possible to recognize that (Q1) can be rewritten as (Q1v) to take advantage of the materialized `TotalByItemStore`.

---

## 4 Introduction

**Key Problems** Example 1.2 above illustrates one important question in the study of materialized views: how to *answer queries using views*, especially when the queries are not originally written in terms of the materialized views. The next natural question is, given a database workload (queries and modifications) as well as resource and performance requirements, how to *select what views to materialize* in the first place. Instead of relying on database administrators and application developers to answer these questions in an ad hoc fashion, we prefer a more systematic and automatic approach.

Materialized views are not free. Not only do they take additional space, but they also require maintenance: as base tables change, the materialized view contents become outdated. Thus, a third important question is how to *maintain materialized views* to keep them up to date with respect to the base tables. The most straightforward way to maintain a materialized view is to recompute its definition query over the new database state whenever the base tables change. However, in practice, since the numbers of rows changed are often small compared with the sizes of the entire base tables, *incremental view maintenance* — the practice of computing and applying only incremental changes to materialized views induced by base table changes — may work better than recomputation.

---

**Example 1.3.** Continuing with Examples 1.1 and 1.2, suppose that five rows have been inserted into base table `pos` as the result of a sale transaction  $\delta$  involving five different items at a particular store. Recomputing the materialized view `TotalByItemStore` from scratch would be both unnecessary and expensive, because most of its  $5 \times 10^6$  rows are not affected by  $\delta$ . With incremental maintenance, loosely speaking, we only need to identify the five affected rows in `TotalByItemStore` and increment their `total` by `qty*price` of their corresponding new `pos` rows inserted by  $\delta$ .

---

To recap, the examples above reveal three key problems concerning materialized views: how to maintain them (*view maintenance*), how to use them (*answering queries using views*), and how to select them (*view*

*selection*). Solutions and techniques developed for these questions over the years have made materialized views an indispensable technology that greatly enhances the performance and features of database systems and many data-intensive applications, such as those mentioned in the opening of this section. Most commercial database systems now offer built-in support of materialized views; for other systems there exist popular recipes for “simulating” support of materialized views.

The ideas underlying materialized views are simple: e.g., precomputation, caching, and incremental evaluation. However, the great database tradition of *declarative querying* is what distinguishes materialized views from generic applications of these ideas in other contexts, and makes materialized views especially interesting, powerful, and challenging at the same time. Thanks to standardized, declarative database languages with clean semantics, study of materialized views has generated a rich body of theory and practice, aimed at providing efficient, effective, automated, and general solutions to the three key problems above.

**Scope and Organization** There is a vast body of literature on materialized views dating back to 1980s, not to mention work related to or influenced by it. There have also been other authoritative references to the topic, most notably the 1999 book edited by Gupta and Mumick [188], the 2001 survey by Halevy [203] on answering queries using views, as well as relevant entries in the recently compiled *Encyclopedia of Database Systems* [291]. We intend this monograph to serve as an accessible introduction and reference to the topic of materialized views for database researchers. In addition to covering the core ideas behind materialized views, we will highlight recent developments (especially those since 2000), and discuss connections to other more recent research topics in databases. This monograph is a more of a pedagogical text than a manual: given a problem, instead of presenting one definitive solution (which in many cases may not be clear or even exist), we walk the readers through the line of reasoning and research developments leading to better understanding of the problem. Therefore, this monograph should be used as a companion to, rather than a substitute for, the literature on materialized views.

## 6 Introduction

The breadth of work on materialized views is as daunting as its depth. Different data models and query languages — object-oriented, semistructured, spatiotemporal, streaming, probabilistic, just to name a few — give rise to a multitude of problem settings that sometimes call for specialized techniques. To make this monograph approachable and focused, we limit its scope mostly to nonrecursive SQL views; we also assume that readers are familiar with relational and bag algebras (see standard database textbooks such as [159, 335], or, for quick reference, [370] and [174], respectively). Our hope is that the core ideas we cover will help readers in further exploring other problem settings.

As mentioned earlier, materialized views now have mature implementations in most commercial database systems. In fact, the database industry has contributed significantly — in many cases as leaders — to the research literature. Written primarily with a research audience in mind, this monograph focuses on the research literature (including contributions from the industry) rather than the product specifics. While we give a high-level overview of commercial implementations, we offer no in-depth comparison of product features.

We note that materialized views are but one form of *derived data* — the result of applying some transformation, structural or computational, to base data. The use of derived data to facilitate access to base data is a recurring theme in computer science. Besides materialized views, other examples include caches, replicas, indexes, and synopses [16, 123]. Despite differences in the form, complexity, and precision of derived data, the three fundamental questions remain: how to use derived data, what to maintain as derived data, and how to maintain them. Oftentimes, ideas and techniques developed for one form of derived data can be adapted and applied to another setting with interesting benefits. The repertoire of techniques for materialized views has been enriched by ideas from other forms of derived data. At the same time, many research areas, old and new alike, have drawn insights from materialized views, implicitly or explicitly. This monograph will highlight a few examples of such cross-pollination.

In the remainder of this monograph, Section 2 covers the *view maintenance* problem. Section 3 covers the *view use* problem. Section 4

covers the *view selection* problem. Section 5 explores connections between materialized views and a few other topics. Section 6 concludes with our perspectives on the current state and future directions of the study of materialized views.

## References

---

- [1] S. Abiteboul, R. Hull, and V. Vianu, *Foundations of Databases*. Addison-Wesley, 1995.
- [2] A. Aboulnaga and K. Salem, “Report: 4th international workshop on self-managing database systems (SMDB 2009),” *IEEE Data Engineering Bulletin*, vol. 32, no. 4, pp. 2–5, 2009.
- [3] B. Adelberg, H. Garcia-Molina, and B. Kao, “Applying update streams in a soft real-time database system,” in *Proceedings of the 1995 ACM SIGMOD International Conference on Management of Data*, pp. 245–256, San Jose, California, USA, May 1995.
- [4] M. E. Adiba and B. G. Lindsay, “Database snapshots,” in *Proceedings of the 1980 International Conference on Very Large Data Bases*, pp. 86–91, Montreal, Quebec, Canada, October 1980.
- [5] F. Afrati and R. Chirkova, “Selecting and using views to compute aggregate queries,” in *Proceedings of the 2005 International Conference on Database Theory*, pp. 383–397, Edinburgh, UK, January 2005.
- [6] F. Afrati and R. Chirkova, “Selecting and using views to compute aggregate queries,” *Journal of Computer and System Sciences*, vol. 77, no. 6, pp. 1079–1107, 2011.
- [7] F. Afrati, C. Li, and J. Ullman, “Generating efficient plans for queries using views,” in *Proceedings of the 2001 ACM SIGMOD International Conference on Management of Data*, pp. 319–330, Santa Barbara, California, USA, June 2001.
- [8] F. N. Afrati, R. Chirkova, M. Gergatsoulis, B. Kimelfeld, V. Pavlaki, and Y. Sagiv, “On rewriting XPath queries using views,” in *Proceedings of the 2009*

- International Conference on Extending Database Technology*, pp. 168–179, Saint Petersburg, Russia, March 2009.
- [9] F. N. Afrati, R. Chirkova, M. Gergatsoulis, and V. Pavlaki, “View selection for real conjunctive queries,” *Acta Informatica*, vol. 44, no. 5, pp. 289–321, 2007.
- [10] F. N. Afrati, M. Damigos, and M. Gergatsoulis, “Union rewritings for XPath fragments,” in *Proceedings of the 2011 International Database Engineering and Applications Symposium*, pp. 43–51, Lisbon, Portugal, September 2011.
- [11] F. N. Afrati, C. Li, and P. Mitra, “Answering queries using views with arithmetic comparisons,” in *Proceedings of the 2002 ACM Symposium on Principles of Database Systems*, pp. 209–220, Madison, Wisconsin, USA, June 2002.
- [12] F. N. Afrati, C. Li, and J. D. Ullman, “Using views to generate efficient evaluation plans for queries,” *Journal of Computer and System Sciences*, vol. 73, no. 5, pp. 703–724, 2007.
- [13] F. N. Afrati and V. Pavlaki, “Rewriting queries using views with negation,” *AI Communications*, vol. 19, no. 3, pp. 229–237, 2006.
- [14] P. K. Agarwal, J. Xie, J. Yang, and H. Yu, “Scalable continuous query processing by tracking hotspots,” in *Proceedings of the 2006 International Conference on Very Large Data Bases*, pp. 31–42, Seoul, Korea, September 2006.
- [15] C. C. Aggarwal, ed., *Data Streams: Models and Algorithms*. Springer, 1st ed., November 2006.
- [16] C. C. Aggarwal and P. S. Yu, “A survey of synopsis construction in data streams,” in Aggarwal [15], pp. 169–207.
- [17] D. Agrawal, A. E. Abbadi, A. Mostéfaoui, M. Raynal, and M. Roy, “The lord of the rings: Efficient maintenance of views at data warehouses,” in *Proceedings of the 2002 International Symposium on Distributed Computing*, pp. 33–47, Toulouse, France, October 2002.
- [18] D. Agrawal, A. E. Abbadi, A. K. Singh, and T. Yurek, “Efficient view maintenance at data warehouses,” in *Proceedings of the 1997 ACM SIGMOD International Conference on Management of Data*, pp. 417–427, Tucson, Arizona, USA, May 1997.
- [19] P. Agrawal, A. Silberstein, B. F. Cooper, U. Srivastava, and R. Ramakrishnan, “Asynchronous view maintenance for VLSD databases,” in *Proceedings of the 2009 ACM SIGMOD International Conference on Management of Data*, pp. 179–192, Providence, Rhode Island, USA, June 2009.
- [20] S. Agrawal, N. Bruno, S. Chaudhuri, and V. R. Narasayya, “AutoAdmin: Self-tuning database systems technology,” *IEEE Data Engineering Bulletin*, vol. 29, no. 3, pp. 7–15, 2006.
- [21] S. Agrawal, S. Chaudhuri, L. Kollár, A. P. Marathe, V. R. Narasayya, and M. Syamala, “Database tuning advisor for Microsoft SQL Server 2005,” in *Proceedings of the 2004 International Conference on Very Large Data Bases*, pp. 1110–1121, Toronto, Canada, August 2004.
- [22] S. Agrawal, S. Chaudhuri, L. Kollár, A. P. Marathe, V. R. Narasayya, and M. Syamala, “Database tuning advisor for Microsoft SQL Server 2005: demo,” in *Proceedings of the 2005 ACM SIGMOD International Conference on Management of Data*, pp. 930–932, Baltimore, Maryland, USA, June 2005.

- [23] S. Agrawal, S. Chaudhuri, and V. R. Narasayya, "Automated selection of materialized views and indexes in SQL databases," in *Proceedings of the 2000 International Conference on Very Large Data Bases*, pp. 496–505, Cairo, Egypt, September 2000.
- [24] S. Agrawal, S. Chaudhuri, and V. R. Narasayya, "Materialized view and index selection tool for Microsoft SQL Server 2000," in *Proceedings of the 2001 ACM SIGMOD International Conference on Management of Data*, p. 608, Santa Barbara, California, USA, June 2001.
- [25] M. K. Aguilera, R. E. Strom, D. C. Sturman, M. Astley, and T. D. Chandra, "Matching events in a content-based subscription system," in *Proceedings of the 1999 ACM Symposium on Principles of Distributed Computing*, pp. 53–61, Atlanta, Georgia, USA, May 1999.
- [26] Y. Ahmad, O. Kennedy, C. Koch, and M. Nikolic, "DBToaster: Higher-order delta processing for dynamic, frequently fresh views," *Proceedings of the VLDB Endowment*, vol. 5, no. 10, pp. 968–979, 2012.
- [27] A. Ailamaki, S. Babu, P. Furtado, S. Lightstone, G. M. Lohman, P. Martin, V. R. Narasayya, G. Pauley, K. Salem, K.-U. Sattler, and G. Weikum, "Report: 3rd International Workshop on Self-Managing Database Systems (SMDB 2008)," *IEEE Data Engineering Bulletin*, vol. 31, no. 4, pp. 2–5, 2008.
- [28] A. Ailamaki, S. Chaudhuri, S. Lightstone, G. M. Lohman, P. Martin, K. Salem, and G. Weikum, "Report on the Second International Workshop on Self-Managing Database Systems (SMDB 2007)," *IEEE Data Engineering Bulletin*, vol. 30, no. 2, pp. 2–4, 2007.
- [29] M. O. Akinde, O. G. Jensen, and M. H. Böhlen, "Minimizing detail data in data warehouses," in *Proceedings of the 1998 International Conference on Extending Database Technology*, pp. 293–307, Valencia, Spain, March 1998.
- [30] M. Arenas, P. Barceló, L. Libkin, and F. Murlak, *Relational and XML Data Exchange*. Synthesis Lectures on Data Management. Morgan & Claypool Publishers, 2010.
- [31] A. Arion, V. Benzaken, I. Manolescu, and Y. Papakonstantinou, "Structured materialized views for XML queries," in *Proceedings of the 2007 International Conference on Very Large Data Bases*, pp. 87–98, Vienna, Austria, September 2007.
- [32] Z. Asgharzadeh Talebi, R. Chirkova, and Y. Fathi, "Exact and inexact methods for solving the problem of view selection for aggregate queries," *International Journal of Business Intelligence and Data Mining*, vol. 4, no. 3/4, pp. 391–415, 2009.
- [33] Z. Asgharzadeh Talebi, R. Chirkova, Y. Fathi, and M. Stallmann, "Exact and inexact methods for selecting views and indexes for OLAP performance improvement," in *Proceedings of the 2008 International Conference on Extending Database Technology*, pp. 311–322, Nantes, France, March 2008.
- [34] B. Babcock, M. Datar, and R. Motwani, "Load shedding in data stream systems," in Aggarwal [15], pp. 127–147.
- [35] S. Babu, K. Munagala, J. Widom, and R. Motwani, "Adaptive caching for continuous queries," in *Proceedings of the 2005 International Conference on Data Engineering*, Tokyo, Japan, April 2005.

## 90 References

- [36] S. Babu and K.-U. Sattler, "Report: 5th international workshop on self-managing database systems (SMDB 2010)," *IEEE Data Engineering Bulletin*, vol. 33, no. 3, pp. 4–7, 2010.
- [37] S. Babu, U. Srivastava, and J. Widom, "Exploiting  $k$ -constraints to reduce memory overhead in continuous queries over data streams," *ACM Transactions on Database Systems*, vol. 29, no. 3, pp. 545–580, 2004.
- [38] A. Balmin, F. Özcan, K. S. Beyer, R. Cochrane, and H. Pirahesh, "A framework for using materialized XPath views in XML query processing," in *Proceedings of the 2004 International Conference on Very Large Data Bases*, pp. 60–71, Toronto, Canada, August 2004.
- [39] M. Bamha, F. Bentayeb, and G. Hains, "An efficient scalable parallel view maintenance algorithm for shared nothing multi-processor machines," in *Proceedings of the 1999 International Conference on Database and Expert Systems Applications*, pp. 616–625, Florence, Italy, August 1999.
- [40] E. Baralis, S. Paraboschi, and E. Teniente, "Materialized views selection in a multidimensional database," in *Proceedings of the 1997 International Conference on Very Large Data Bases*, pp. 156–165, Athens, Greece, August 1997.
- [41] P. Barceló, "Logical foundations of relational data exchange," *ACM SIGMOD Record*, vol. 38, no. 1, pp. 49–58, 2009.
- [42] P. Belknap, B. Dageville, K. Dias, and K. Yagoub, "Self-tuning for SQL performance in Oracle Database 11g," in *Proceedings of the 2009 International Conference on Data Engineering*, pp. 1694–1700, Shanghai, China, March 2009.
- [43] R. G. Bello, K. Dias, A. Downing, J. J. F. Jr., J. L. Finnerty, W. D. Norcott, H. Sun, A. Witkowski, and M. Ziauddin, "Materialized views in Oracle," in *Proceedings of the 1998 International Conference on Very Large Data Bases*, pp. 659–664, New York City, New York, USA, August 1998.
- [44] M. Benedikt and G. Gottlob, "The impact of virtual views on containment," *Proceedings of the VLDB Endowment*, vol. 3, no. 1, pp. 297–308, 2010.
- [45] P. A. Bernstein and L. M. Haas, "Information integration in the enterprise," *Communications of the ACM*, vol. 51, no. 9, pp. 72–79, 2008.
- [46] P. Bhatotia, A. Wieder, R. Rodrigues, U. A. Acar, and R. Pasquin, "Incoop: MapReduce for incremental computations," in *Proceedings of the 2011 ACM Symposium on Cloud Computing*, pp. 7:1–7:14, Cascais, Portugal, October 2011.
- [47] P. Bizarro, S. Babu, D. DeWitt, and J. Widom, "Content-based routing: Different plans for different data," in *Proceedings of the 2005 International Conference on Very Large Data Bases*, Trondheim, Norway, August 2005.
- [48] J. A. Blakeley, N. Coburn, and P.-Å. Larson, "Updating derived relations: Detecting irrelevant and autonomously computable updates," in *Proceedings of the 1986 International Conference on Very Large Data Bases*, pp. 457–466, Kyoto, Japan, August 1986.
- [49] J. A. Blakeley, N. Coburn, and P.-V. Larson, "Updating derived relations: Detecting irrelevant and autonomously computable updates," *ACM Transactions on Database Systems*, vol. 14, no. 3, pp. 369–400, 1989.
- [50] J. A. Blakeley, P.-Å. Larson, and F. W. Tompa, "Efficiently updating materialized views," in *Proceedings of the 1986 ACM SIGMOD International Conference on Management of Data*, pp. 61–71, Washington DC, USA, May 1986.

- [51] J. A. Blakeley and N. L. Martin, "Join index, materialized view, and hybrid-hash join: A performance analysis," in *Proceedings of the 1990 International Conference on Data Engineering*, pp. 256–263, Los Angeles, California, USA, February 1990.
- [52] A. Bonifati, M. H. Goodfellow, I. Manolescu, and D. Sileo, "Algebraic incremental maintenance of XML views," in *Proceedings of the 2011 International Conference on Extending Database Technology*, pp. 177–188, Uppsala, Sweden, March 2011.
- [53] P. Bonnet and D. Shasha, "Index tuning," in Liu and Özsu [291], pp. 1433–1435.
- [54] P. Bonnet and D. Shasha, "Schema tuning," in Liu and Özsu [291], pp. 2497–2499.
- [55] N. Bruno and S. Chaudhuri, "Automatic physical database tuning: A relaxation-based approach," in *Proceedings of the 2005 ACM SIGMOD International Conference on Management of Data*, pp. 227–238, Baltimore, Maryland, USA, June 2005.
- [56] N. Bruno and S. Chaudhuri, "Physical design refinement: The "merge-reduce" approach," in *Proceedings of the 2006 International Conference on Extending Database Technology*, pp. 386–404, Munich, Germany, March 2006.
- [57] N. Bruno and S. Chaudhuri, "To tune or not to tune? A lightweight physical design alerter," in *Proceedings of the 2006 International Conference on Very Large Data Bases*, pp. 499–510, Seoul, Korea, September 2006.
- [58] N. Bruno and S. Chaudhuri, "Online approach to physical design tuning," in *Proceedings of the 2007 International Conference on Data Engineering*, pp. 826–835, Istanbul, Turkey, April 2007.
- [59] N. Bruno and S. Chaudhuri, "Online AutoAdmin (physical design tuning)," in *Proceedings of the 2007 ACM SIGMOD International Conference on Management of Data*, pp. 1067–1069, Beijing, China, June 2007.
- [60] N. Bruno and S. Chaudhuri, "Physical design refinement: The merge-reduce approach," *ACM Transactions on Database Systems*, vol. 32, no. 4, pp. 28–43, 2007.
- [61] N. Bruno and S. Chaudhuri, "Constrained physical design tuning," *Proceedings of the VLDB Endowment*, vol. 1, pp. 4–15, 2008.
- [62] N. Bruno and S. Chaudhuri, "Constrained physical design tuning," *The VLDB Journal*, vol. 19, no. 1, pp. 21–44, 2010.
- [63] N. Bruno, S. Chaudhuri, and G. Weikum, "Database tuning using online algorithms," in Liu and Özsu [291], pp. 741–744.
- [64] P. Buneman and E. K. Clemons, "Efficiently monitoring relational databases," *ACM Transactions on Database Systems*, vol. 4, no. 3, pp. 368–382, 1979.
- [65] C. J. Bunger, L. S. Colby, R. L. Cole, W. J. McKenna, G. Mulagund, and D. Wilhite, "Aggregate maintenance for data warehousing in Informix Red Brick Vista," in *Proceedings of the 2001 International Conference on Very Large Data Bases*, pp. 659–662, Roma, Italy, September 2001.
- [66] A. Cali, D. Calvanese, G. D. Giacomo, and M. Lenzerini, "Data integration under integrity constraints," *Information Systems*, vol. 29, no. 2, pp. 147–163, 2004.

## 92 References

- [67] D. Calvanese and G. D. Giacomo, "Data integration: A logic-based perspective," *AI Magazine*, vol. 26, no. 1, pp. 59–70, 2005.
- [68] D. Calvanese, G. D. Giacomo, M. Lenzerini, D. Nardi, and R. Rosati, "Data integration in data warehousing," *International Journal of Cooperative Information Systems*, vol. 10, no. 3, pp. 237–271, 2001.
- [69] D. Calvanese, G. D. Giacomo, M. Lenzerini, and R. Rosati, "Logical foundations of peer-to-peer data integration," in *Proceedings of the 2004 ACM Symposium on Principles of Database Systems*, pp. 241–251, Paris, France, June 2004.
- [70] D. Calvanese, G. D. Giacomo, M. Lenzerini, and R. Rosati, "View-based query answering over description logic ontologies," in *Proceedings of the 2008 International Conference on Principles of Knowledge Representation and Reasoning*, pp. 242–251, Sydney, Australia, September 2008.
- [71] D. Calvanese, G. D. Giacomo, M. Lenzerini, and M. Y. Vardi, "Answering regular path queries using views," in *Proceedings of the 2000 International Conference on Data Engineering*, pp. 389–398, Los Angeles, California, USA, February 2000.
- [72] D. Calvanese, G. D. Giacomo, M. Lenzerini, and M. Y. Vardi, "Lossless regular views," in *Proceedings of the 2002 ACM Symposium on Principles of Database Systems*, pp. 247–258, Madison, Wisconsin, USA, June 2002.
- [73] D. Calvanese, G. D. Giacomo, M. Lenzerini, and M. Y. Vardi, "Query containment using views," in *Proceedings of the 2003 Italian Symposium on Advanced Database Systems*, pp. 467–474, Cetraro (CS), Italy, June 2003.
- [74] D. Calvanese, G. D. Giacomo, M. Lenzerini, and M. Y. Vardi, "View-based query containment," in *Proceedings of the 2003 ACM Symposium on Principles of Database Systems*, pp. 56–67, San Diego, California, USA, June 2003.
- [75] D. Calvanese, G. D. Giacomo, M. Lenzerini, and M. Y. Vardi, "View-based query processing: On the relationship between rewriting, answering and losslessness," in *Proceedings of the 2005 International Conference on Database Theory*, pp. 321–336, Edinburgh, UK, January 2005.
- [76] D. Calvanese, G. D. Giacomo, M. Lenzerini, and M. Y. Vardi, "View-based query processing: On the relationship between rewriting, answering and losslessness," *Theoretical Computer Science*, vol. 371, no. 3, pp. 169–182, 2007.
- [77] K. S. Candan, D. Agrawal, W.-S. Li, O. Po, and W.-P. Hsiung, "View invalidation for dynamic content caching in multitiered architectures," in *Proceedings of the 2002 International Conference on Very Large Data Bases*, pp. 562–573, Hong Kong, China, September 2002.
- [78] S. Castano, V. D. Antonellis, and S. D. C. di Vimercati, "Global viewing of heterogeneous data sources," *IEEE Transactions on Knowledge and Data Engineering*, vol. 13, no. 2, pp. 277–297, 2001.
- [79] B. Cautis, A. Deutsch, and N. Onose, "XPath rewriting using multiple views: Achieving completeness and efficiency," in *Proceedings of the 2008 International Workshop on the Web and Databases*, Vancouver, Canada, June 2008.
- [80] B. Cautis, A. Deutsch, and N. Onose, "Querying data sources that export infinite sets of views," in *Proceedings of the 2009 International Conference on Database Theory*, pp. 84–97, Saint-Petersburg, Russia, March 2009.

- [81] B. Cautis, A. Deutsch, N. Onose, and V. Vassalos, "Efficient rewriting of XPath queries using query set specifications," *Proceedings of the VLDB Endowment*, vol. 2, no. 1, pp. 301–312, 2009.
- [82] B. Cautis, A. Deutsch, N. Onose, and V. Vassalos, "Querying XML data sources that export very large sets of views," *ACM Transactions on Database Systems*, vol. 36, no. 1, p. 5, 2011.
- [83] S. Ceri and J. Widom, "Deriving production rules for incremental view maintenance," in *Proceedings of the 1991 International Conference on Very Large Data Bases*, pp. 577–589, Barcelona, Catalonia, Spain, 1991.
- [84] A. Chandra and P. Merlin, "Optimal implementation of conjunctive queries in relational data bases," in *Proceedings of the 1977 ACM Symposium on Theory of Computing*, pp. 77–90, Boulder, Colorado, USA, May 1977.
- [85] S. Chandrasekaran and M. J. Franklin, "PSoup: A system for streaming queries over streaming data," *The VLDB Journal*, vol. 12, no. 2, pp. 140–156, 2003.
- [86] S. Chaudhuri, "An overview of query optimization in relational systems," in *Proceedings of the 1998 ACM Symposium on Principles of Database Systems*, pp. 34–43, Seattle, Washington, USA, June 1998.
- [87] S. Chaudhuri, E. Christensen, G. Graefe, V. R. Narasayya, and M. J. Zwilling, "Self-tuning technology in Microsoft SQL Server," *IEEE Data Engineering Bulletin*, vol. 22, no. 2, pp. 20–26, 1999.
- [88] S. Chaudhuri, M. Datar, and V. R. Narasayya, "Index selection for databases: A hardness study and principled heuristic solution," *IEEE Transactions on Knowledge and Data Engineering*, vol. 16, pp. 1313–1323, 2004.
- [89] S. Chaudhuri and U. Dayal, "An overview of data warehousing and OLAP technology," *ACM SIGMOD Record*, vol. 26, no. 1, pp. 65–74, 1997.
- [90] S. Chaudhuri, R. Krishnamurthy, S. Potamianos, and K. Shim, "Optimizing queries with materialized views," in *Proceedings of the 1995 International Conference on Data Engineering*, pp. 190–200, Taipei, Taiwan, March 1995.
- [91] S. Chaudhuri and V. R. Narasayya, "An efficient cost-driven index selection tool for Microsoft SQL server," in *Proceedings of the 1997 International Conference on Very Large Data Bases*, pp. 146–155, Athens, Greece, August 1997.
- [92] S. Chaudhuri and V. R. Narasayya, "AutoAdmin 'what-if' index analysis utility," in *Proceedings of the 1998 ACM SIGMOD International Conference on Management of Data*, pp. 367–378, Seattle, Washington, USA, May 1998.
- [93] S. Chaudhuri and V. R. Narasayya, "Self-tuning database systems: A decade of progress," in *Proceedings of the 2007 International Conference on Very Large Data Bases*, pp. 3–14, Vienna, Austria, September 2007.
- [94] S. Chaudhuri, V. R. Narasayya, and G. Weikum, "Database tuning using combinatorial search," in Liu and Özsu [291], pp. 738–741.
- [95] S. Chaudhuri and M. Y. Vardi, "Optimization of real conjunctive queries," in *Proceedings of the 1993 ACM Symposium on Principles of Database Systems*, pp. 59–70, Washington DC, USA, May 1993.
- [96] S. Chaudhuri and G. Weikum, "Self-management technology in databases," in Liu and Özsu [291], pp. 2550–2555.
- [97] S. Chaudhuri and G. Weikum, "Rethinking database system architecture: Towards a self-tuning risc-style database system," in *Proceedings of the 2000*

## 94 References

- International Conference on Very Large Data Bases*, pp. 1–10, Cairo, Egypt, September 2000.
- [98] L. W. F. Chaves, E. Buchmann, F. Hueske, and K. Böhm, “Towards materialized view selection for distributed databases,” in *Proceedings of the 2009 International Conference on Extending Database Technology*, pp. 1088–1099, Saint Petersburg, Russia, March 2009.
- [99] D. Chen and C.-Y. Chan, “ViewJoin: Efficient view-based evaluation of tree pattern queries,” in *Proceedings of the 2010 International Conference on Data Engineering*, pp. 816–827, Long Beach, California, USA, March 2010.
- [100] J. Chen, S. Chen, and E. A. Rundensteiner, “A transactional model for data warehouse maintenance,” in *Proceedings of the 2002 International Conference on Conceptual Modeling*, pp. 247–262, Tampere, Finland, October 2002.
- [101] J. Chen, D. J. DeWitt, F. Tian, and Y. Wang, “NiagaraCQ: A scalable continuous query system for internet databases,” in *Proceedings of the 2000 ACM SIGMOD International Conference on Management of Data*, pp. 379–390, Dallas, Texas, USA, May 2000.
- [102] J. Chen, X. Zhang, S. Chen, A. Koeller, and E. A. Rundensteiner, “DyDa: Data warehouse maintenance in fully concurrent environments,” in *Proceedings of the 2001 ACM SIGMOD International Conference on Management of Data*, p. 619, Santa Barbara, California, USA, June 2001.
- [103] L. Chen and E. A. Rundensteiner, “XCache: XQuery-based caching system,” in *Proceedings of the 2002 International Workshop on the Web and Databases*, pp. 31–36, Madison, Wisconsin, USA, June 2002.
- [104] S. Chen, B. Liu, and E. A. Rundensteiner, “Multiversion-based view maintenance over distributed data sources,” *ACM Transactions on Database Systems*, vol. 29, no. 4, pp. 675–709, 2004.
- [105] J. Cheney, L. Chiticariu, and W. C. Tan, “Provenance in databases: Why, how and where,” *Foundations and Trends in Databases*, vol. 1, no. 4, pp. 379–474, 2009.
- [106] R. Chirkova, “Query containment,” in Liu and Özsu [291], pp. 2249–2253.
- [107] R. Chirkova, “The view-selection problem has an exponential-time lower bound for conjunctive queries and views,” in *Proceedings of the 2002 ACM Symposium on Principles of Database Systems*, pp. 159–168, Madison, Wisconsin, USA, June 2002.
- [108] R. Chirkova, A. Y. Halevy, and D. Suciu, “A formal perspective on the view selection problem,” *The VLDB Journal*, vol. 11, no. 3, pp. 216–237, 2002.
- [109] R. Chirkova and C. Li, “Materializing views with minimal size to answer queries,” in *Proceedings of the 2003 ACM Symposium on Principles of Database Systems*, pp. 38–48, San Diego, California, USA, June 2003.
- [110] R. Chirkova, C. Li, and J. Li, “Answering queries using materialized views with minimum size,” *The VLDB Journal*, vol. 15, no. 3, pp. 191–210, 2006.
- [111] S. Cohen, “Aggregation: Expressiveness and containment,” in Liu and Özsu [291], pp. 59–63.
- [112] S. Cohen, “Equivalence of queries combining set and bag-set semantics,” in *Proceedings of the 2006 ACM Symposium on Principles of Database Systems*, pp. 70–79, Chicago, Illinois, USA, June 2006.

- [113] S. Cohen, “User-defined aggregate functions: Bridging theory and practice,” in *Proceedings of the 2006 ACM SIGMOD International Conference on Management of Data*, pp. 49–60, Chicago, Illinois, USA, June 2006.
- [114] S. Cohen, “Equivalence of queries that are sensitive to multiplicities,” *The VLDB Journal*, vol. 18, pp. 765–785, 2009.
- [115] S. Cohen, W. Nutt, and Y. Sagiv, “Rewriting queries with arbitrary aggregation functions using views,” *ACM Transactions on Database Systems*, vol. 31, no. 2, pp. 672–715, 2006.
- [116] S. Cohen, W. Nutt, and Y. Sagiv, “Deciding equivalences among conjunctive aggregate queries,” *Journal of the ACM*, vol. 54, no. 2, 2007.
- [117] S. Cohen, W. Nutt, and A. Serebrenik, “Rewriting aggregate queries using views,” in *Proceedings of the 1999 ACM Symposium on Principles of Database Systems*, pp. 155–166, Philadelphia, Pennsylvania, USA, June 1999.
- [118] S. Cohen, W. Nutt, and A. Serebrenik, “Algorithms for rewriting aggregate queries using views,” in *Proceedings of the 2000 East European Conference on Advances in Databases and Information Systems Held Jointly with the International Conference on Database Systems for Advanced Applications*, pp. 65–78, Prague, Czech Republic, September 2000.
- [119] L. S. Colby, T. Griffin, L. Libkin, I. S. Mumick, and H. Trickey, “Algorithms for deferred view maintenance,” in *Proceedings of the 1996 ACM SIGMOD International Conference on Management of Data*, pp. 469–480, Montreal, Quebec, Canada, June 1996.
- [120] L. S. Colby, A. Kawaguchi, D. F. Lieuwen, I. S. Mumick, and K. A. Ross, “Supporting multiple view maintenance policies,” in *Proceedings of the 1997 ACM SIGMOD International Conference on Management of Data*, pp. 405–416, Tucson, Arizona, USA, May 1997.
- [121] M. Compton, “Finding equivalent rewritings with exact views,” in *Proceedings of the 2009 International Conference on Data Engineering*, pp. 1243–1246, Shanghai, China, March 2009.
- [122] T. Condie, N. Conway, P. Alvaro, J. M. Hellerstein, K. Elmeleegy, and R. Sears, “MapReduce online,” in *Proceedings of the 2010 USENIX Symposium on Networked Systems Design and Implementation*, San Jose, California, USA, April 2010.
- [123] G. Cormode, M. N. Garofalakis, P. J. Haas, and C. Jermaine, “Synopses for massive data: Samples, histograms, wavelets, sketches,” *Foundations and Trends in Databases*, vol. 4, no. 1–3, pp. 1–294, 2012.
- [124] Y. Cui and J. Widom, “Storing auxiliary data for efficient maintenance and lineage tracing of complex views,” in *Proceedings of the 2000 International Workshop on Design and Management of Data Warehouses*, Stockholm, Sweden, June 2000.
- [125] B. Dageville, D. Das, K. Dias, K. Yagoub, M. Zait, and M. Ziauddin, “Automatic SQL tuning in Oracle 10g,” in *Proceedings of the 2004 International Conference on Very Large Data Bases*, pp. 1098–1109, Toronto, Canada, August 2004.
- [126] B. Dageville and K. Dias, “Oracle’s self-tuning architecture and solutions,” *IEEE Data Engineering Bulletin*, vol. 29, no. 3, pp. 24–31, 2006.

## 96 References

- [127] N. N. Dalvi, C. Re, and D. Suciu, “Queries and materialized views on probabilistic databases,” *Journal of Computer and System Sciences*, vol. 77, no. 3, pp. 473–490, 2011.
- [128] N. N. Dalvi and D. Suciu, “Answering queries from statistics and probabilistic views,” in *Proceedings of the 2005 International Conference on Very Large Data Bases*, pp. 805–816, Trondheim, Norway, August 2005.
- [129] S. Dar, M. J. Franklin, B. T. Jónsson, D. Srivastava, and M. Tan, “Semantic data caching and replacement,” in *Proceedings of the 1996 International Conference on Very Large Data Bases*, pp. 330–341, Mumbai (Bombay), India, September 1996.
- [130] J. Dean and S. Ghemawat, “MapReduce: A flexible data processing tool,” *Communications of the ACM*, vol. 53, no. 1, pp. 72–77, 2010.
- [131] D. DeHaan, P.-Å. Larson, and J. Zhou, “Stacked indexed views in Microsoft SQL Server,” in *Proceedings of the 2005 ACM SIGMOD International Conference on Management of Data*, pp. 179–190, Baltimore, Maryland, USA, June 2005.
- [132] A. Deligiannakis, “View maintenance aspects,” in Liu and Özsu [291], pp. 3328–3331.
- [133] A. J. Demers, J. Gehrke, B. Panda, M. Riedewald, V. Sharma, and W. M. White, “Cayuga: A general purpose event monitoring system,” in *Proceedings of the 2007 Conference on Innovative Data Systems Research*, pp. 412–422, Asilomar, California, USA, January 2007.
- [134] A. Deutsch, “FOL modeling of integrity constraints (dependencies),” in Liu and Özsu [291], pp. 1155–1161.
- [135] A. Deutsch, Y. Katsis, and Y. Papakonstantinou, “Determining source contribution in integration systems,” in *Proceedings of the 2005 ACM Symposium on Principles of Database Systems*, pp. 304–315, Baltimore, Maryland, USA, June 2005.
- [136] A. Deutsch, B. Ludäscher, and A. Nash, “Rewriting queries using views with access patterns under integrity constraints,” in *Proceedings of the 2005 International Conference on Database Theory*, pp. 352–367, Edinburgh, UK, January 2005.
- [137] A. Deutsch and A. Nash, “Chase,” in Liu and Özsu [291], pp. 323–327.
- [138] A. Deutsch, L. Popa, and V. Tannen, “Query reformulation with constraints,” *ACM SIGMOD Record*, vol. 35, no. 1, pp. 65–73, 2006.
- [139] A. Doan, A. Halevy, and Z. Ives, *Principles of Data Integration*. Morgan Kaufmann, 1st ed., July 2012.
- [140] A. Doan and A. Y. Halevy, “Semantic integration research in the database community: A brief survey,” *AI Magazine*, vol. 26, no. 1, pp. 83–94, 2005.
- [141] G. Dong and J. Su, “Incremental computation of queries,” in Liu and Özsu [291], pp. 1414–1417.
- [142] A. El-Helw, I. F. Ilyas, and C. Zuzarte, “Statadvisor: Recommending statistical views,” *Proceedings of the VLDB Endowment*, vol. 2, no. 2, pp. 1306–1317, 2009.
- [143] M. El-Sayed, E. A. Rundensteiner, and M. Mani, “Incremental maintenance of materialized XQuery views,” in *Proceedings of the 2006 International Conference on Data Engineering*, p. 129, Atlanta, Georgia, USA, April 2006.

- [144] C. Elkan, “Independence of logic database queries and updates,” in *Proceedings of the 1990 ACM Symposium on Principles of Database Systems*, pp. 154–160, Nashville, Tennessee, USA, April 1990.
- [145] H. Engström, S. Chakravarthy, and B. Lings, “A systematic approach to selecting maintenance policies in a data warehouse environment,” in *Proceedings of the 2002 International Conference on Extending Database Technology*, pp. 317–335, Prague, Czech Republic, March 2002.
- [146] P. T. Eugster, P. Felber, R. Guerraoui, and A.-M. Kermarrec, “The many faces of publish/subscribe,” *ACM Computing Surveys*, vol. 35, no. 2, pp. 114–131, 2003.
- [147] F. Fabret, H.-A. Jacobsen, F. Llirbat, J. Pereira, K. A. Ross, and D. Shasha, “Filtering algorithms and implementation for very fast publish/subscribe,” in *Proceedings of the 2001 ACM SIGMOD International Conference on Management of Data*, Santa Barbara, California, USA, June 2001.
- [148] W. Fan, F. Geerts, X. Jia, and A. Kementsietsidis, “Rewriting regular XPath queries on XML views,” in *Proceedings of the 2007 International Conference on Data Engineering*, pp. 666–675, Istanbul, Turkey, April 2007.
- [149] S. J. Finkelstein, M. Schkolnick, and P. Tiberio, “Physical database design for relational databases,” *ACM Transactions on Database Systems*, vol. 13, no. 1, 1988.
- [150] S. Flesca and S. Greco, “Rewriting queries using views,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 13, no. 6, pp. 980–995, 2001.
- [151] N. Folkert, A. Gupta, A. Witkowski, S. Subramanian, S. Bellamkonda, S. Shankar, T. Bozkaya, and L. Sheng, “Optimizing refresh of a set of materialized views,” in *Proceedings of the 2005 International Conference on Very Large Data Bases*, pp. 1043–1054, Trondheim, Norway, August 2005.
- [152] Y. Fu, K. Kowalczykowski, K. W. Ong, Y. Papakonstantinou, and K. K. Zhao, “Ajax-based report pages as incrementally rendered views,” in *Proceedings of the 2010 ACM SIGMOD International Conference on Management of Data*, pp. 567–578, Indianapolis, Indiana, USA, June 2010.
- [153] Y. Fu, K. W. Ong, Y. Papakonstantinou, and M. Petropoulos, “The SQL-based all-declarative FORWARD web application development framework,” in *Proceedings of the 2011 Conference on Innovative Data Systems Research*, pp. 69–78, Asilomar, California, USA, January 2011.
- [154] A. Fuxman, P. G. Kolaitis, R. J. Miller, and W. C. Tan, “Peer data exchange,” in *Proceedings of the 2005 ACM Symposium on Principles of Database Systems*, pp. 160–171, Baltimore, Maryland, USA, June 2005.
- [155] A. Fuxman, P. G. Kolaitis, R. J. Miller, and W. C. Tan, “Peer data exchange,” *ACM Transactions on Database Systems*, vol. 31, no. 4, pp. 1454–1498, 2006.
- [156] A. Fuxman and R. J. Miller, “First-order query rewriting for inconsistent databases,” in *Proceedings of the 2005 International Conference on Database Theory*, pp. 337–351, Edinburgh, UK, January 2005.
- [157] H. Garcia-Molina, W. J. Labio, and J. Yang, “Expiring data in a warehouse,” in *Proceedings of the 1998 International Conference on Very Large Data Bases*, pp. 500–511, New York City, New York, USA, August 1998.

## 98 References

- [158] H. Garcia-Molina, J. Ullman, and J. Widom, *Database Systems: The Complete Book*. Pearson Prentice Hall, 2009.
- [159] H. Garcia-Molina, J. D. Ullman, and J. Widom, *Database Systems: The Complete Book*. Pearson Education, 2nd ed., 2009.
- [160] K. E. Gebaly and A. Aboulnaga, “Robustness in automatic physical database design,” in *Proceedings of the 2008 International Conference on Extending Database Technology*, pp. 145–156, Nantes, France, March 2008.
- [161] R. Gemulla and W. Lehner, “Deferred maintenance of disk-based random samples,” in *Proceedings of the 2006 International Conference on Extending Database Technology*, pp. 423–441, Munich, Germany, March 2006.
- [162] M. R. Genesereth, *Data Integration: The Relational Logic Approach. Synthesis Lectures on Artificial Intelligence and Machine Learning*. Morgan & Claypool Publishers, 2010.
- [163] T. M. Ghanem, A. K. Elmagarmid, P.-Å. Larson, and W. G. Aref, “Supporting views in data stream management systems,” *ACM Transactions on Database Systems*, vol. 35, no. 1, 2010.
- [164] G. D. Giacomo, D. Lembo, M. Lenzerini, and R. Rosati, “On reconciling data exchange, data integration, and peer data management,” in *Proceedings of the 2007 ACM Symposium on Principles of Database Systems*, pp. 133–142, Beijing, China, June 2007.
- [165] P. Godfrey and J. Gryz, “View disassembly: A rewrite that extracts portions of views,” *Journal of Computer and System Sciences*, vol. 73, no. 6, pp. 941–961, 2007.
- [166] P. Godfrey, J. Gryz, A. Hoppe, W. Ma, and C. Zuzarte, “Query rewrites with views for XML in DB2,” in *Proceedings of the 2009 International Conference on Data Engineering*, pp. 1339–1350, Shanghai, China, March 2009.
- [167] J. Goldstein and P.-Å. Larson, “Optimizing queries using materialized views: A practical, scalable solution,” in *Proceedings of the 2001 ACM SIGMOD International Conference on Management of Data*, pp. 331–342, Santa Barbara, California, USA, June 2001.
- [168] G. Gou, M. Kormilitsin, and R. Chirkova, “Query evaluation using overlapping views: Completeness and efficiency,” in *Proceedings of the 2006 ACM SIGMOD International Conference on Management of Data*, pp. 37–48, Chicago, Illinois, USA, June 2006.
- [169] G. Graefe and H. A. Kuno, “Self-selecting, self-tuning, incrementally optimized indexes,” in *Proceedings of the 2010 International Conference on Extending Database Technology*, pp. 371–381, Lausanne, Switzerland, March 2010.
- [170] G. Graefe and M. J. Zwillig, “Transaction support for indexed views,” in *Proceedings of the 2004 ACM SIGMOD International Conference on Management of Data*, Paris, France, June 2004.
- [171] G. Grahne and A. Thomo, “Query containment and rewriting using views for regular path queries under constraints,” in *Proceedings of the 2003 ACM Symposium on Principles of Database Systems*, pp. 111–122, San Diego, California, USA, June 2003.

- [172] J. Gray, A. Bosworth, A. Layman, and H. Pirahesh, "Data cube: A relational aggregation operator generalizing group-by, cross-tab, and sub-total," in *Proceedings of the 1996 International Conference on Data Engineering*, pp. 152–159, New Orleans, Louisiana, USA, February 1996.
- [173] J. Gray, S. Chaudhuri, A. Bosworth, A. Layman, D. Reichart, and M. Venkatrao, "Data cube: A relational aggregation operator generalizing group-by, cross-tab, and sub-totals," *Data Mining and Knowledge Discovery*, vol. 1, no. 1, pp. 29–53, 1997.
- [174] T. J. Green, "Bag semantics," in Liu and Özsu [291], pp. 201–206.
- [175] T. J. Green and Z. G. Ives, "Recomputing materialized instances after changes to mappings and data," in *Proceedings of the 2012 International Conference on Data Engineering*, pp. 330–341, Washington DC, USA, April 2012.
- [176] T. J. Green, G. Karvounarakis, Z. G. Ives, and V. Tannen, "Update exchange with mappings and provenance," in *Proceedings of the 2007 International Conference on Very Large Data Bases*, pp. 675–686, Vienna, Austria, September 2007.
- [177] T. J. Green, G. Karvounarakis, Z. G. Ives, and V. Tannen, "Provenance in ORCHESTRA," *IEEE Data Engineering Bulletin*, vol. 33, no. 3, pp. 9–16, 2010.
- [178] T. J. Green, G. Karvounarakis, N. E. Taylor, O. Biton, Z. G. Ives, and V. Tannen, "ORCHESTRA: Facilitating collaborative data sharing," in *Proceedings of the 2007 ACM SIGMOD International Conference on Management of Data*, pp. 1131–1133, Beijing, China, June 2007.
- [179] T. Griffin and B. Kumar, "Algebraic change propagation for semijoin and outerjoin queries," *ACM SIGMOD Record*, vol. 27, no. 3, pp. 22–27, 1998.
- [180] T. Griffin and L. Libkin, "Incremental maintenance of views with duplicates," in *Proceedings of the 1995 ACM SIGMOD International Conference on Management of Data*, pp. 328–339, San Jose, California, USA, May 1995.
- [181] T. Griffin, L. Libkin, and H. Trickey, "An improved algorithm for the incremental recomputation of active relational expressions," *IEEE Transactions on Knowledge and Data Engineering*, vol. 9, no. 3, pp. 508–511, 1997.
- [182] S. Grumbach, M. Rafanelli, and L. Tininini, "On the equivalence and rewriting of aggregate queries," *Acta Informatica*, vol. 40, no. 8, pp. 529–584, 2004.
- [183] S. Grumbach and L. Tininini, "On the content of materialized aggregate views," *Journal of Computer and System Sciences*, vol. 66, no. 1, pp. 133–168, 2003.
- [184] A. Gupta and J. A. Blakeley, "Using partial information to update materialized views," *Information Systems*, vol. 20, no. 8, pp. 641–662, 1995.
- [185] A. Gupta, V. Harinarayan, and D. Quass, "Aggregate-query processing in data warehousing environments," in *Proceedings of the 1995 International Conference on Very Large Data Bases*, pp. 358–369, Zurich, Switzerland, September 1995.
- [186] A. Gupta, H. V. Jagadish, and I. S. Mumick, "Data integration using self-maintainable views," in *Proceedings of the 1996 International Conference on Extending Database Technology*, pp. 140–144, Avignon, France, March 1996.

100 *References*

- [187] A. Gupta and I. S. Mumick, "Maintenance of materialized views: Problems, techniques, and applications," *IEEE Data Engineering Bulletin*, vol. 18, no. 2, pp. 3–18, 1995.
- [188] A. Gupta and I. S. Mumick, eds., *Materialized Views: Techniques, Implementations, and Applications*. MIT Press, 1999.
- [189] A. Gupta, I. S. Mumick, and V. S. Subrahmanian, "Maintaining views incrementally," in *Proceedings of the 1993 ACM SIGMOD International Conference on Management of Data*, pp. 157–166, Washington DC, USA, May 1993.
- [190] A. K. Gupta, A. Y. Halevy, and D. Suciu, "View selection for stream processing," in *Proceedings of the 2002 International Workshop on the Web and Databases*, pp. 83–88, Madison, Wisconsin, USA, June 2002.
- [191] A. K. Gupta, D. Suciu, and A. Y. Halevy, "The view selection problem for XML content based routing," in *Proceedings of the 2003 ACM Symposium on Principles of Database Systems*, pp. 68–77, San Diego, California, USA, June 2003.
- [192] H. Gupta, "Selection of views to materialize in a data warehouse," in *Proceedings of the 1997 International Conference on Database Theory*, pp. 98–112, Delphi, Greece, January 1997.
- [193] H. Gupta, "Selection and maintenance of views in a data warehouse," PhD thesis, Department of Computer Science, Stanford University, 1999.
- [194] H. Gupta, V. Harinarayan, A. Rajaraman, and J. D. Ullman, "Index selection for OLAP," in *Proceedings of the 1997 International Conference on Data Engineering*, pp. 208–219, Birmingham, UK, April 1997.
- [195] H. Gupta and I. S. Mumick, "Selection of views to materialize under a maintenance cost constraint," in *Proceedings of the 1999 International Conference on Database Theory*, pp. 453–470, Jerusalem, Israel, January 1999.
- [196] H. Gupta and I. S. Mumick, "Selection of views to materialize in a data warehouse," *IEEE Transactions on Knowledge and Data Engineering*, vol. 17, no. 1, pp. 24–43, 2005.
- [197] H. Gupta and I. S. Mumick, "Incremental maintenance of aggregate and outerjoin expressions," *Information Systems*, vol. 31, no. 6, pp. 435–464, 2006.
- [198] N. Gupta, L. Kot, G. Bender, S. Roy, J. Gehrke, and C. Koch, "Coordination through querying in the Youtopia system," in *Proceedings of the 2011 ACM SIGMOD International Conference on Management of Data*, pp. 1331–1334, Athens, Greece, June 2011.
- [199] L. M. Haas, "Beauty and the beast: The theory and practice of information integration," in *Proceedings of the 2007 International Conference on Database Theory*, pp. 28–43, Barcelona, Spain, January 2007.
- [200] A. Halevy, "Data integration: A status report," in *Datenbanksysteme für Business, Technologie und Web*, pp. 24–29, Leipzig, Germany, February 2003.
- [201] A. Y. Halevy, "Information integration," in Liu and Özsu [291], pp. 1490–1496.
- [202] A. Y. Halevy, "Theory of answering queries using views," *ACM SIGMOD Record*, vol. 29, no. 4, pp. 40–47, 2000.
- [203] A. Y. Halevy, "Answering queries using views: A survey," *The VLDB Journal*, vol. 10, no. 4, pp. 270–294, 2001.

- [204] A. Y. Halevy, N. Ashish, D. Bitton, M. J. Carey, D. Draper, J. Pollock, A. Rosenthal, and V. Sikka, "Enterprise information integration: Successes, challenges and controversies," in *Proceedings of the 2005 ACM SIGMOD International Conference on Management of Data*, pp. 778–787, Baltimore, Maryland, USA, June 2005.
- [205] A. Y. Halevy, Z. G. Ives, J. Madhavan, P. Mork, D. Suciu, and I. Tatarinov, "The Piazza peer data management system," *IEEE Transactions on Knowledge and Data Engineering*, vol. 16, no. 7, pp. 787–798, 2004.
- [206] A. Y. Halevy, Z. G. Ives, P. Mork, and I. Tatarinov, "Piazza: Data management infrastructure for semantic web applications," in *Proceedings of the 2003 International Conference on World Wide Web*, pp. 556–567, Budapest, Hungary, May 2003.
- [207] A. Y. Halevy, A. Rajaraman, and J. J. Ordille, "Data integration: The teenage years," in *Proceedings of the 2006 International Conference on Very Large Data Bases*, pp. 9–16, Seoul, Korea, September 2006.
- [208] J. Han, M. Kamber, and J. Pei, *Data Mining: Concepts and Techniques*. Morgan Kaufmann, 3rd ed., 2005.
- [209] E. N. Hanson, "A performance analysis of view materialization strategies," in *Proceedings of the 1987 ACM SIGMOD International Conference on Management of Data*, pp. 440–453, San Francisco, California, USA, May 1987.
- [210] E. N. Hanson, C. Carnes, L. Huang, M. Konyala, L. Noronha, S. Parthasarathy, J. B. Park, and A. Vernon, "Scalable trigger processing," in *Proceedings of the 1999 International Conference on Data Engineering*, pp. 266–275, Sydney, Australia, March 1999.
- [211] N. Hanusse, S. Maabout, and R. Tofan, "A view selection algorithm with performance guarantee," in *Proceedings of the 2009 International Conference on Extending Database Technology*, pp. 946–957, Saint Petersburg, Russia, March 2009.
- [212] V. Harinarayan, A. Rajaraman, and J. D. Ullman, "Implementing data cubes efficiently," in *Proceedings of the 1996 ACM SIGMOD International Conference on Management of Data*, pp. 205–216, Montreal, Quebec, Canada, June 1996.
- [213] H. He, J. Xie, J. Yang, and H. Yu, "Asymmetric batch incremental view maintenance," in *Proceedings of the 2005 International Conference on Data Engineering*, pp. 106–117, Tokyo, Japan, April 2005.
- [214] J. M. Hellerstein, M. Stonebraker, and J. R. Hamilton, "Architecture of a database system," *Foundations and Trends in Databases*, vol. 1, no. 2, pp. 141–259, 2007.
- [215] K. Hose, D. Klan, and K.-U. Sattler, "Online tuning of aggregation tables for OLAP," in *Proceedings of the 2009 International Conference on Data Engineering*, pp. 1679–1686, Shanghai, China, March 2009.
- [216] V. Hristidis and M. Petropoulos, "Semantic caching of XML databases," in *Proceedings of the 2002 International Workshop on the Web and Databases*, pp. 25–30, Madison, Wisconsin, USA, June 2002.
- [217] R. Hull and G. Zhou, "A framework for supporting data integration using the materialized and virtual approaches," in *Proceedings of the 1996 ACM*

102 *References*

- SIGMOD International Conference on Management of Data*, pp. 481–492, Montreal, Quebec, Canada, June 1996.
- [218] E. Hung, Y. Deng, and V. S. Subrahmanian, “RDF aggregate queries and views,” in *Proceedings of the 2005 International Conference on Data Engineering*, pp. 717–728, Tokyo, Japan, April 2005.
- [219] C. A. Hurtado, C. Gutiérrez, and A. O. Mendelzon, “Capturing summarizability with integrity constraints in OLAP,” *ACM Transactions on Database Systems*, vol. 30, no. 3, pp. 854–886, 2005.
- [220] N. Huyn, “Efficient view self-maintenance,” in *Proceedings of the 1996 Workshop on Materialized Views*, pp. 17–25, 1996.
- [221] N. Huyn, “Multiple-view self-maintenance in data warehousing environments,” in *Proceedings of the 1997 International Conference on Very Large Data Bases*, pp. 26–35, Athens, Greece, August 1997.
- [222] N. Huyn, “Speeding up view maintenance using cheap filters at the warehouse,” in *Proceedings of the 2000 International Conference on Data Engineering*, p. 308, Los Angeles, California, USA, February 2000.
- [223] Y. Ioannidis and R. Ramakrishnan, “Containment of conjunctive queries: Beyond relations as sets,” *ACM Transactions on Database Systems*, vol. 20, no. 3, pp. 288–324, 1995.
- [224] Y. E. Ioannidis, “Query optimization,” in *The Computer Science and Engineering Handbook*, (A. B. Tucker, ed.), pp. 1038–1057, CRC Press, 1997.
- [225] Z. G. Ives, T. J. Green, G. Karvounarakis, N. E. Taylor, V. Tannen, P. P. Talukdar, M. Jacob, and F. Pereira, “The ORCHESTRA collaborative data sharing system,” *ACM SIGMOD Record*, vol. 37, no. 3, pp. 26–32, 2008.
- [226] Z. G. Ives, A. Y. Halevy, P. Mork, and I. Tatarinov, “Piazza: Mediation and integration infrastructure for semantic web data,” *The Journal of Web Semantics*, vol. 1, no. 2, pp. 155–175, 2004.
- [227] Z. G. Ives, N. Khandelwal, A. Kapur, and M. Cakir, “ORCHESTRA: Rapid, collaborative sharing of dynamic data,” in *Proceedings of the 2005 Conference on Innovative Data Systems Research*, pp. 107–118, Asilomar, California, USA, January 2005.
- [228] H. V. Jagadish, I. S. Mumick, and A. Silberschatz, “View maintenance issues for the chronicle data model,” in *Proceedings of the 1995 ACM Symposium on Principles of Database Systems*, pp. 113–124, San Jose, California, USA, June 1995.
- [229] M. Jarke and J. Koch, “Query optimization in database systems,” *ACM Computing Surveys*, vol. 16, no. 2, pp. 111–152, 1984.
- [230] C. Jermaine, A. Pol, and S. Arumugam, “Online maintenance of very large random samples,” in *Proceedings of the 2004 ACM SIGMOD International Conference on Management of Data*, pp. 299–310, Paris, France, June 2004.
- [231] H. Jiang, D. Gao, and W.-S. Li, “Exploiting correlation and parallelism of materialized-view recommendation for distributed data warehouses,” in *Proceedings of the 2007 International Conference on Data Engineering*, pp. 276–285, Istanbul, Turkey, April 2007.
- [232] S. Joshi and C. Jermaine, “Materialized sample views for database approximation,” in *Proceedings of the 2006 International Conference on Data Engineering*, p. 151, Atlanta, Georgia, USA, April 2006.

- [233] S. Joshi and C. M. Jermaine, “Materialized sample views for database approximation,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 20, no. 3, pp. 337–351, 2008.
- [234] B. Kähler and O. Risnes, “Extending logging for database snapshot refresh,” in *Proceedings of the 1987 International Conference on Very Large Data Bases*, pp. 389–398, Brighton, England, September 1987.
- [235] H.-G. Kang and C.-W. Chung, “Exploiting versions for on-line data warehouse maintenance in MOLAP servers,” in *Proceedings of the 2002 International Conference on Very Large Data Bases*, pp. 742–753, Hong Kong, China, September 2002.
- [236] J. Kang, J. F. Naughton, and S. Viglas, “Evaluating window joins over unbounded streams,” in *Proceedings of the 2003 International Conference on Data Engineering*, pp. 341–352, Bangalore, India, March 2003.
- [237] H. J. Karloff and M. Mihail, “On the complexity of the view-selection problem,” in *Proceedings of the 1999 ACM Symposium on Principles of Database Systems*, pp. 167–173, Philadelphia, Pennsylvania, USA, June 1999.
- [238] G. Karvounarakis and Z. G. Ives, “Bidirectional mappings for data and update exchange,” in *Proceedings of the 2008 International Workshop on the Web and Databases*, Vancouver, Canada, June 2008.
- [239] Y. Katsis and Y. Papakonstantinou, “View-based data integration,” in Liu and Özsu [291], pp. 3332–3339.
- [240] A. Kawaguchi, D. F. Lieuwen, I. S. Mumick, D. Quass, and K. A. Ross, “Concurrency control theory for deferred materialized views,” in *Proceedings of the 1997 International Conference on Database Theory*, pp. 306–320, Delphi, Greece, January 1997.
- [241] O. Kennedy, Y. Ahmad, and C. Koch, “DBToaster: Agile views for a dynamic data management system,” in *Proceedings of the 2011 Conference on Innovative Data Systems Research*, pp. 284–295, Asilomar, California, USA, January 2011.
- [242] R. Kimball and M. Ross, *The Data Warehouse Toolkit: Practical Techniques for Building Dimensional Data Warehouses*. John Wiley, 2nd ed., 2002.
- [243] H. Kimura, G. Huo, A. Rasin, S. Madden, and S. B. Zdonik, “CORADD: Correlation aware database designer for materialized views and indexes,” *Proceedings of the VLDB Endowment*, vol. 3, no. 1, pp. 1103–1113, 2010.
- [244] A. Klug, “On conjunctive queries containing inequalities,” *Journal of the ACM*, vol. 35, no. 1, pp. 146–160, 1988.
- [245] P. G. Kolaitis, D. L. Martin, and M. N. Thakur, “On the complexity of the containment problem for conjunctive queries with built-in predicates,” in *Proceedings of the 1998 ACM Symposium on Principles of Database Systems*, pp. 197–204, Seattle, Washington, USA, June 1998.
- [246] M. Kormilitsin, R. Chirkova, Y. Fathi, and M. Stallmann, “View and index selection for query-performance improvement: Quality-centered algorithms and heuristics,” in *Proceedings of the 2008 International Conference on Information and Knowledge Management*, pp. 1329–1330, Napa Valley, California, USA, October 2008.
- [247] M. Kormilitsin, R. Chirkova, Y. Fathi, and M. Stallmann, “Systematic exploration of efficient query plans for automated database restructuring,” in

- Proceedings of the 2009 East European Conference on Advances in Databases and Information Systems*, pp. 133–148, Riga, Latvia, September 2009.
- [248] L. Kot and C. Koch, “Cooperative update exchange in the Youtopia system,” *Proceedings of the VLDB Endowment*, vol. 2, no. 1, pp. 193–204, 2009.
- [249] Y. Kotidis and N. Roussopoulos, “DynaMat: A dynamic view management system for data warehouses,” in *Proceedings of the 1999 ACM SIGMOD International Conference on Management of Data*, pp. 371–382, Philadelphia, Pennsylvania, USA, May 1999.
- [250] Y. Kotidis and N. Roussopoulos, “A case for dynamic view management,” *ACM Transactions on Database Systems*, vol. 26, no. 4, pp. 388–423, 2001.
- [251] S. Kulkarni and M. K. Mohania, “Concurrent maintenance of views using multiple versions,” in *Proceedings of the 1999 International Database Engineering and Applications Symposium*, pp. 254–259, Montreal, Canada, August 1999.
- [252] W. Labio, D. Quass, and B. Adelberg, “Physical database design for data warehouses,” in *Proceedings of the 1997 International Conference on Data Engineering*, pp. 277–288, Birmingham, UK, April 1997.
- [253] W. J. Labio, J. Yang, Y. Cui, H. Garcia-Molina, and J. Widom, “Performance issues in incremental warehouse maintenance,” in *Proceedings of the 2000 International Conference on Very Large Data Bases*, pp. 461–472, Cairo, Egypt, September 2000.
- [254] W. J. Labio, R. Yerneni, and H. Garcia-Molina, “Shrinking the warehouse update window,” in *Proceedings of the 1999 ACM SIGMOD International Conference on Management of Data*, pp. 383–394, Philadelphia, Pennsylvania, USA, May 1999.
- [255] A. Labrinidis, Q. Luo, J. Xu, and W. Xue, “Caching and materialization for web databases,” *Foundations and Trends in Databases*, vol. 2, no. 3, pp. 169–266, 2009.
- [256] L. V. S. Lakshmanan, J. Pei, and Y. Zhao, “QC-trees: An efficient summary structure for semantic OLAP,” in *Proceedings of the 2003 ACM SIGMOD International Conference on Management of Data*, pp. 64–75, San Diego, California, USA, June 2003.
- [257] L. V. S. Lakshmanan, W. H. Wang, and Z. J. Zhao, “Answering tree pattern queries using views,” in *Proceedings of the 2006 International Conference on Very Large Data Bases*, pp. 571–582, Seoul, Korea, September 2006.
- [258] P.-Å. Larson, W. Lehner, J. Zhou, and P. Zabback, “Cardinality estimation using sample views with quality assurance,” in *Proceedings of the 2007 ACM SIGMOD International Conference on Management of Data*, pp. 175–186, Beijing, China, June 2007.
- [259] P.-Å. Larson, W. Lehner, J. Zhou, and P. Zabback, “Exploiting self-monitoring sample views for cardinality estimation,” in *Proceedings of the 2007 ACM SIGMOD International Conference on Management of Data*, pp. 1073–1075, Beijing, China, June 2007.
- [260] P.-Å. Larson and H. Z. Yang, “Computing queries from derived relations,” in *Proceedings of the 1985 International Conference on Very Large Data Bases*, pp. 259–269, Stockholm, Sweden, August 1985.

- [261] P.-Å. Larson and J. Zhou, “View matching for outer-join views,” in *Proceedings of the 2005 International Conference on Very Large Data Bases*, pp. 445–456, Trondheim, Norway, August 2005.
- [262] P.-Å. Larson and J. Zhou, “Efficient maintenance of materialized outer-join views,” in *Proceedings of the 2007 International Conference on Data Engineering*, pp. 56–65, Istanbul, Turkey, April 2007.
- [263] P.-Å. Larson and J. Zhou, “View matching for outer-join views,” *The VLDB Journal*, vol. 16, no. 1, pp. 29–53, 2007.
- [264] J. Lechtenbörger and G. Vossen, “On the computation of relational view complements,” in *Proceedings of the 2002 ACM Symposium on Principles of Database Systems*, pp. 142–149, Madison, Wisconsin, USA, June 2002.
- [265] J. Lechtenbörger and G. Vossen, “On the computation of relational view complements,” *ACM Transactions on Database Systems*, vol. 28, no. 2, pp. 175–208, 2003.
- [266] M. Lee and J. Hammer, “Speeding up materialized view selection in data warehouses using a randomized algorithm,” *International Journal of Cooperative Information Systems*, vol. 10, no. 3, pp. 327–353, 2001.
- [267] W. Lehner, “Query processing in data warehouses,” in Liu and Özsu [291], pp. 2297–2301.
- [268] W. Lehner, R. Cochrane, H. Pirahesh, and M. Zaharioudakis, “fAST refresh using mass query optimization,” in *Proceedings of the 2001 International Conference on Data Engineering*, pp. 391–398, Heidelberg, Germany, April 2001.
- [269] W. Lehner, R. Sidle, H. Pirahesh, and R. Cochrane, “Maintenance of automatic summary tables,” in *Proceedings of the 2000 ACM SIGMOD International Conference on Management of Data*, pp. 512–513, Dallas, Texas, USA, May 2000.
- [270] M. Lenzerini, “Data integration: A theoretical perspective,” in *Proceedings of the 2002 ACM Symposium on Principles of Database Systems*, pp. 233–246, Madison, Wisconsin, USA, June 2002.
- [271] M. Lenzerini, “Data integration: A theoretical perspective,” in *Proceedings of the 2002 ACM Symposium on Principles of Database Systems*, pp. 233–246, Madison, Wisconsin, USA, June 2002.
- [272] A. Y. Levy, A. O. Mendelzon, Y. Sagiv, and D. Srivastava, “Answering queries using views,” in *Proceedings of the 1995 ACM Symposium on Principles of Database Systems*, pp. 95–104, San Jose, California, USA, June 1995.
- [273] A. Y. Levy and Y. Sagiv, “Queries independent of updates,” in *Proceedings of the 1993 International Conference on Very Large Data Bases*, pp. 171–181, Dublin, Ireland, August 1993.
- [274] C. Li, “Rewriting queries using views,” in Liu and Özsu [291], pp. 2438–2441.
- [275] C. Li, M. Bawa, and J. D. Ullman, “Minimizing view sets without losing query-answering power,” in *Proceedings of the 2001 International Conference on Database Theory*, pp. 99–113, London, UK, January 2001.
- [276] J. Li, Z. Asgharzadeh Talebi, R. Chirkova, and Y. Fathi, “A formal model for the problem of view selection for aggregate queries,” in *Proceedings of the 2005 East European Conference on Advances in Databases and Information Systems*, pp. 125–138, Tallinn, Estonia, September 2005.

106 *References*

- [277] W.-S. Li, D. C. Zilio, V. S. Batra, M. Subramanian, C. Zuzarte, and I. Narang, "Load balancing for multi-tiered database systems through autonomic placement of materialized views," in *Proceedings of the 2006 International Conference on Data Engineering*, p. 102, Atlanta, Georgia, USA, April 2006.
- [278] W. Liang, H. Li, H. Wang, and M. E. Orłowska, "Making multiple views self-maintainable in a data warehouse," *Data and Knowledge Engineering*, vol. 30, no. 2, pp. 121–134, 1999.
- [279] W. Liang, H. Wang, and M. Orłowska, "Materialized view selection under the maintenance time constraint," *Data and Knowledge Engineering*, vol. 37, pp. 203–216, 2001.
- [280] S. Lifschitz and M. A. V. Salles, "Autonomic index management," in *Proceedings of the 2005 International Conference on Autonomic Computing*, pp. 304–305, Seattle, Washington, USA, June 2005.
- [281] S. Lightstone, "Seven software engineering principles for autonomic computing development," *Innovations in Systems and Software Engineering*, vol. 3, no. 1, pp. 71–74, 2007.
- [282] S. Lightstone, G. M. Lohman, P. J. Haas, V. Markl, J. Rao, A. J. Storm, M. Surendra, and D. C. Zilio, "Making DB2 products self-managing: Strategies and experiences," *IEEE Data Engineering Bulletin*, vol. 29, no. 3, pp. 16–23, 2006.
- [283] S. Lightstone, G. M. Lohman, and D. C. Zilio, "Toward autonomic computing with DB2 universal database," *ACM SIGMOD Record*, vol. 31, no. 3, pp. 55–61, 2002.
- [284] S. Lightstone, M. Surendra, Y. Diao, S. S. Parekh, J. L. Hellerstein, K. Rose, A. J. Storm, and C. Garcia-Arellano, "Control theory: A foundational technique for self managing databases," in *ICDE Workshops*, pp. 395–403, 2007.
- [285] S. Lightstone, T. Teorey, and T. Nadeau, *Physical Database Design: The Database Professional's Guide to Exploiting Indexes, Views, Storage, and More*. Morgan Kaufmann, 2007.
- [286] B. G. Lindsay, L. M. Haas, C. Mohan, H. Pirahesh, and P. F. Wilms, "A snapshot differential refresh algorithm," in *Proceedings of the 1986 ACM SIGMOD International Conference on Management of Data*, pp. 53–60, Washington DC, USA, May 1986.
- [287] B. Liu, S. Chen, and E. A. Rundensteiner, "Batch data warehouse maintenance in dynamic environments," in *Proceedings of the 2002 International Conference on Information and Knowledge Management*, pp. 68–75, McLean, Virginia, USA, November 2002.
- [288] B. Liu, S. Chen, and E. A. Rundensteiner, "A transactional approach to parallel data warehouse maintenance," in *Proceedings of the 2002 International Conference on Data Warehousing and Knowledge Discovery*, pp. 307–316, Aix-en-Provence, France, September 2002.
- [289] B. Liu and E. A. Rundensteiner, "Cost-driven general join view maintenance over distributed data sources," in *Proceedings of the 2005 International Conference on Data Engineering*, pp. 578–579, Tokyo, Japan, April 2005.
- [290] B. Liu, E. A. Rundensteiner, and D. Finkel, "Restructuring batch view maintenance efficiently," in *Proceedings of the 2004 International Conference on*

- Information and Knowledge Management*, pp. 228–229, Washington DC, USA, November 2004.
- [291] L. Liu and M. T. Özsu, eds., *Encyclopedia of Database Systems*. Springer, 2009.
- [292] Z. Liu and Y. Chen, “Answering keyword queries on XML using materialized views,” in *Proceedings of the 2008 International Conference on Data Engineering*, pp. 1501–1503, Cancun, Mexico, April 2008.
- [293] G. M. Lohman and S. Lightstone, “SMART: Making DB2 (more) autonomic,” in *Proceedings of the 2002 International Conference on Very Large Data Bases*, pp. 877–879, Hong Kong, China, September 2002.
- [294] B. T. Loo, T. Condie, M. N. Garofalakis, D. E. Gay, J. M. Hellerstein, P. Maniatis, R. Ramakrishnan, T. Roscoe, and I. Stoica, “Declarative networking,” *Communications of the ACM*, vol. 52, no. 11, pp. 87–95, 2009.
- [295] G. Luo, “V locking protocol for materialized aggregate join views on B-tree indices,” in *Proceedings of the 2010 International Conference on Web-Age Information Management*, vol. 6184 of Lecture Notes in Computer Science, (L. Chen, C. Tang, J. Yang, and Y. Gao, eds.), pp. 768–780, Jiuzhaigou, Sichuan, China: Springer, July 2010. ISBN 978-3-642-14245-1.
- [296] G. Luo, “Partial materialized views,” in *Proceedings of the 2007 International Conference on Data Engineering*, pp. 756–765, Istanbul, Turkey, April 2007.
- [297] G. Luo, J. F. Naughton, C. J. Ellmann, and M. Watzke, “A comparison of three methods for join view maintenance in parallel RDBMS,” in *Proceedings of the 2003 International Conference on Data Engineering*, pp. 177–188, Bangalore, India, March 2003.
- [298] G. Luo, J. F. Naughton, C. J. Ellmann, and M. Watzke, “Locking protocols for materialized aggregate join views,” in *Proceedings of the 2003 International Conference on Very Large Data Bases*, pp. 596–607, Berlin, Germany, September 2003.
- [299] G. Luo, J. F. Naughton, C. J. Ellmann, and M. Watzke, “Locking protocols for materialized aggregate join views,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 17, no. 6, pp. 796–807, 2005.
- [300] G. Luo and P. S. Yu, “Content-based filtering for efficient online materialized view maintenance,” in *Proceedings of the 2008 International Conference on Information and Knowledge Management*, pp. 163–172, Napa Valley, California, USA, October 2008.
- [301] S. Madden, M. A. Shah, J. M. Hellerstein, and V. Raman, “Continuously adaptive continuous queries over streams,” in *Proceedings of the 2002 ACM SIGMOD International Conference on Management of Data*, Madison, Wisconsin, USA, June 2002.
- [302] B. Mandhani and D. Suciu, “Query caching and view selection for XML databases,” in *Proceedings of the 2005 International Conference on Very Large Data Bases*, pp. 469–480, Trondheim, Norway, August 2005.
- [303] M. Marx, “Queries determined by views: Pack your views,” in *Proceedings of the 2007 ACM Symposium on Principles of Database Systems*, pp. 23–30, Beijing, China, June 2007.

108 *References*

- [304] G. Mecca, A. O. Mendelzon, and P. Merialdo, “Efficient queries over web views,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 14, no. 6, pp. 1280–1298, 2002.
- [305] H. Mistry, P. Roy, S. Sudarshan, and K. Ramamritham, “Materialized view selection and maintenance using multi-query optimization,” in *Proceedings of the 2001 ACM SIGMOD International Conference on Management of Data*, pp. 307–318, Santa Barbara, California, USA, June 2001.
- [306] M. K. Mohania and Y. Kambayashi, “Making aggregate views self-maintainable,” *Data and Knowledge Engineering*, vol. 32, no. 1, pp. 87–109, 2000.
- [307] I. S. Mumick, D. Quass, and B. S. Mumick, “Maintenance of data cubes and summary tables in a warehouse,” in *Proceedings of the 1997 ACM SIGMOD International Conference on Management of Data*, pp. 100–111, Tucson, Arizona, USA, May 1997.
- [308] K. Munagala, J. Yang, and H. Yu, “Online view maintenance under a response-time constraint,” in *Proceedings of the 2005 European Symposium on Algorithms*, pp. 677–688, Palma de Mallorca, Spain, October 2005.
- [309] S. Muthukrishnan, “Data streams: Algorithms and applications,” *Theoretical Computer Science*, vol. 1, 2006.
- [310] A. Nash, L. Segoufin, and V. Vianu, “Views and queries: Determinacy and rewriting,” *ACM Transactions on Database Systems*, vol. 35, no. 3, 2010.
- [311] A. Nica, A. J. Lee, and E. A. Rundensteiner, “The CVS algorithm for view synchronization in evolvable large-scale information systems,” in *Proceedings of the 1998 International Conference on Extending Database Technology*, pp. 359–373, Valencia, Spain, March 1998.
- [312] N. F. Noy, A. Doan, and A. Y. Halevy, “Semantic integration,” *AI Magazine*, vol. 26, no. 1, pp. 7–10, 2005.
- [313] K. O’Gorman, D. Agrawal, and A. E. Abbadi, “Posse: A framework for optimizing incremental view maintenance at data warehouse,” in *Proceedings of the 1999 International Conference on Data Warehousing and Knowledge Discovery*, pp. 106–115, Florence, Italy, September 1999.
- [314] K. O’Gorman, D. Agrawal, and A. E. Abbadi, “On the importance of tuning in incremental view maintenance: An experience case study,” in *Proceedings of the 2000 International Conference on Data Warehousing and Knowledge Discovery*, pp. 77–82, London, UK, September 2000.
- [315] F. Olken and D. Rotem, “Random sampling from database files: A survey,” in *Proceedings of the 1990 International Conference on Scientific and Statistical Database Management*, pp. 92–111, Charlotte, North Carolina, USA, April 1990.
- [316] C. Olston, G. Chiou, L. Chitnis, F. Liu, Y. Han, M. Larsson, A. Neumann, V. B. N. Rao, V. Sankarasubramanian, S. Seth, C. Tian, T. ZiCornell, and X. Wang, “Nova: Continuous Pig/Hadoop workflows,” in *Proceedings of the 2011 ACM SIGMOD International Conference on Management of Data*, pp. 1081–1090, Athens, Greece, June 2011.
- [317] C. Olston, B. Reed, U. Srivastava, R. Kumar, and A. Tomkins, “Pig latin: A not-so-foreign language for data processing,” in *Proceedings of the 2008 ACM*

- SIGMOD International Conference on Management of Data*, pp. 1099–1110, Vancouver, Canada, June 2008.
- [318] P. E. O’Neil and D. Quass, “Improved query performance with variant indexes,” in *Proceedings of the 1997 ACM SIGMOD International Conference on Management of Data*, pp. 38–49, Tucson, Arizona, USA, May 1997.
- [319] N. Onose, A. Deutsch, Y. Papakonstantinou, and E. Curtmola, “Rewriting nested XML queries using nested views,” in *Proceedings of the 2006 ACM SIGMOD International Conference on Management of Data*, pp. 443–454, Chicago, Illinois, USA, June 2006.
- [320] T. Palpanas, R. Sidle, R. Cochrane, and H. Pirahesh, “Incremental maintenance for non-distributive aggregate functions,” in *Proceedings of the 2002 International Conference on Very Large Data Bases*, pp. 802–813, Hong Kong, China, September 2002.
- [321] S. Papadomanolakis and A. Ailamaki, “An integer linear programming approach to database design,” in *ICDE Workshops*, pp. 442–449, 2007.
- [322] S. Paraboschi, G. Sindoni, E. Baralis, and E. Teniente, “Materialized views in multidimensional databases,” in *Multidimensional Databases*, pp. 222–251, Idea Group, 2003.
- [323] C.-S. Park, M. Kim, and Y.-J. Lee, “Rewriting OLAP queries using materialized views and dimension hierarchies in data warehouses,” in *Proceedings of the 2001 International Conference on Data Engineering*, pp. 515–523, Heidelberg, Germany, April 2001.
- [324] C.-S. Park, M. Kim, and Y.-J. Lee, “Finding an efficient rewriting of OLAP queries using materialized views in data warehouses,” *Decision Support Systems*, vol. 32, no. 4, pp. 379–399, 2002.
- [325] N. W. Paton and O. Díaz, “Active database systems,” *ACM Computing Surveys*, vol. 31, no. 1, pp. 63–103, 1999.
- [326] T. Phan and W.-S. Li, “Dynamic materialization of query views for data warehouse workloads,” in *Proceedings of the 2008 International Conference on Data Engineering*, pp. 436–445, Cancun, Mexico, April 2008.
- [327] B. C. Pierce, “Linguistic foundations for bidirectional transformations: Invited tutorial,” in *Proceedings of the 2012 ACM Symposium on Principles of Database Systems*, pp. 61–64, Scottsdale, Arizona, USA, May 2012.
- [328] V. Poe, *Building a Data Warehouse for Decision Support*. Prentice Hall, 1996.
- [329] A. Pol, C. M. Jermaine, and S. Arumugam, “Maintaining very large random samples using the geometric file,” *The VLDB Journal*, vol. 17, no. 5, pp. 997–1018, 2008.
- [330] L. Popa, M. Budiu, Y. Yu, and M. Isard, “DryadInc: Reusing work in large-scale computations,” in *Proceedings of the 2009 Workshop on Hot Topics on Cloud Computing*, Boston, Massachusetts, USA, June 2009.
- [331] X. Qian and G. Wiederhold, “Incremental recomputation of active relational expressions,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 3, no. 3, pp. 337–341, 1991.
- [332] D. Quass, “Maintenance expressions for views with aggregation,” in *Proceedings of the 1996 Workshop on Materialized Views*, pp. 110–118, 1996.
- [333] D. Quass, A. Gupta, I. S. Mumick, and J. Widom, “Making views self-maintainable for data warehousing,” in *Proceedings of the 1996 International*

110 *References*

- Conference on Parallel and Distributed Information Systems*, pp. 158–169, Miami Beach, Florida, USA, December 1996.
- [334] D. Quass and J. Widom, “On-line warehouse view maintenance,” in *Proceedings of the 1997 ACM SIGMOD International Conference on Management of Data*, pp. 393–404, Tucson, Arizona, USA, May 1997.
- [335] R. Ramakrishnan and J. Gehrke, *Database Management Systems*. McGraw-Hill, 3rd ed., 2009.
- [336] C. Re and D. Suciu, “Materialized views in probabilistic databases for information exchange and query optimization,” in *Proceedings of the 2007 International Conference on Very Large Data Bases*, pp. 51–62, Vienna, Austria, September 2007.
- [337] S. Rizvi, A. Mendelzon, S. Sudarshan, and P. Roy, “Extending query rewriting techniques for fine-grained access control,” in *Proceedings of the 2004 ACM SIGMOD International Conference on Management of Data*, pp. 551–562, Paris, France, June 2004.
- [338] K. A. Ross, “View adaptation,” in Liu and Özsu [291], pp. 3324–3325.
- [339] K. A. Ross, D. Srivastava, and S. Sudarshan, “Materialized view maintenance and integrity constraint checking: Trading space for time,” in *Proceedings of the 1996 ACM SIGMOD International Conference on Management of Data*, pp. 447–458, Montreal, Quebec, Canada, June 1996.
- [340] N. Roussopoulos, “An incremental access method for ViewCache: Concept, algorithms, and cost analysis,” *ACM Transactions on Database Systems*, vol. 16, no. 3, pp. 535–563, 1991.
- [341] N. Roussopoulos, C.-M. Chen, S. Kelley, A. Delis, and Y. Papakonstantinou, “The ADMS project: Views “R” us,” *IEEE Data Engineering Bulletin*, vol. 18, no. 2, pp. 19–28, 1995.
- [342] N. Roussopoulos and H. Kang, “Principles and techniques in the design of ADMS±,” *IEEE Computer*, vol. 19, no. 12, pp. 19–25, 1986.
- [343] S. Rozen and D. Shasha, “A framework for automating physical database design,” in *Proceedings of the 1991 International Conference on Very Large Data Bases*, pp. 401–411, Barcelona, Catalonia, Spain, 1991.
- [344] G. Ruberg and M. Mattoso, “XCraft: Boosting the performance of active XML materialization,” in *Proceedings of the 2008 International Conference on Extending Database Technology*, pp. 299–310, Nantes, France, March 2008.
- [345] Y. Sagiv and M. Yannakakis, “Equivalences among relational expressions with the union and difference operators,” *Journal of the ACM*, vol. 27, no. 4, pp. 633–655, 1980.
- [346] K. Salem, K. S. Beyer, R. Cochrane, and B. G. Lindsay, “How to roll a join: Asynchronous incremental view maintenance,” in *Proceedings of the 2000 ACM SIGMOD International Conference on Management of Data*, pp. 129–140, Dallas, Texas, USA, May 2000.
- [347] S. Samtani, V. Kumar, and M. K. Mohania, “Self maintenance of multiple views in data warehousing,” in *Proceedings of the 1999 International Conference on Information and Knowledge Management*, pp. 292–299, Kansas City, Missouri, USA, November 1999.

- [348] A. D. Sarma, M. Theobald, and J. Widom, "LIVE: A lineage-supported versioned DBMS," in *Proceedings of the 2010 International Conference on Scientific and Statistical Database Management*, pp. 416–433, Heidelberg, Germany, June 2010.
- [349] A. Sawires, J. Tatemura, O. Po, D. Agrawal, A. E. Abbadi, and K. S. Candan, "Maintaining XPath views in loosely coupled systems," in *Proceedings of the 2006 International Conference on Very Large Data Bases*, pp. 583–594, Seoul, Korea, September 2006.
- [350] A. Sawires, J. Tatemura, O. Po, D. Agrawal, and K. S. Candan, "Incremental maintenance of path expression views," in *Proceedings of the 2005 ACM SIGMOD International Conference on Management of Data*, pp. 443–454, Baltimore, Maryland, USA, June 2005.
- [351] A. Segev and W. Fang, "Currency-based updates to distributed materialized views," in *Proceedings of the 1990 International Conference on Data Engineering*, pp. 512–520, Los Angeles, California, USA, February 1990.
- [352] A. Segev and J. Park, "Maintaining materialized views in distributed databases," in *Proceedings of the 1989 International Conference on Data Engineering*, pp. 262–270, Los Angeles, California, USA, February 1989.
- [353] A. Segev and J. Park, "Updating distributed materialized views," *IEEE Transactions on Knowledge and Data Engineering*, vol. 1, no. 2, pp. 173–184, 1989.
- [354] L. Segoufin and V. Vianu, "Views and queries: Determinacy and rewriting," in *Proceedings of the 2005 ACM Symposium on Principles of Database Systems*, pp. 49–60, Baltimore, Maryland, USA, June 2005.
- [355] P. Seshadri and A. N. Swami, "Generalized partial indexes," in *Proceedings of the 1995 International Conference on Data Engineering*, pp. 420–427, Taipei, Taiwan, March 1995.
- [356] D. Shasha, "Tuning database design for high performance," in *The Computer Science and Engineering Handbook*, pp. 995–1011, CRC Press, 1997.
- [357] D. Shasha, P. Bonnet, and J. Gray, *Database Tuning: Principles, Experiments and Troubleshooting Techniques*. Morgan Kaufmann, 2003.
- [358] A. Shukla, P. Deshpande, and J. F. Naughton, "Materialized view selection for multidimensional datasets," in *Proceedings of the 1998 International Conference on Very Large Data Bases*, pp. 488–499, New York City, New York, USA, August 1998.
- [359] A. Shukla, P. Deshpande, and J. F. Naughton, "Materialized view selection for multi-cube data models," in *Proceedings of the 2000 International Conference on Extending Database Technology*, pp. 269–284, Konstanz, Germany, March 2000.
- [360] A. Simitsis and D. Theodoratos, "Data warehouse back-end tools," in *Encyclopedia of Data Warehousing and Mining*, (J. Wang, ed.), pp. 572–579, IGI Global, 2009.
- [361] Y. Sismanis, A. Deligiannakis, N. Roussopoulos, and Y. Kotidis, "Dwarf: Shrinking the PetaCube," in *Proceedings of the 2002 ACM SIGMOD International Conference on Management of Data*, pp. 464–475, Madison, Wisconsin, USA, June 2002.

112 *References*

- [362] Y. Sismanis and N. Roussopoulos, "The complexity of fully materialized coalesced cubes," in *Proceedings of the 2004 International Conference on Very Large Data Bases*, pp. 540–551, Toronto, Canada, August 2004.
- [363] D. Srivastava, S. Dar, H. V. Jagadish, and A. Y. Levy, "Answering queries with aggregation using views," in *Proceedings of the 1996 International Conference on Very Large Data Bases*, pp. 318–329, Mumbai (Bombay), India, September 1996.
- [364] J. Srivastava and D. Rotem, "Analytical modeling of materialized view maintenance," in *Proceedings of the 1988 ACM Symposium on Principles of Database Systems*, pp. 126–134, Austin, Texas, USA, March 1988.
- [365] M. Stonebraker, "The case for partial indexes," *ACM SIGMOD Record*, vol. 18, no. 4, pp. 4–11, 1989.
- [366] K. Tajima and Y. Fukui, "Answering XPath queries over networks by sending minimal views," in *Proceedings of the 2004 International Conference on Very Large Data Bases*, pp. 48–59, Toronto, Canada, August 2004.
- [367] P.-N. Tan, M. Steinbach, and V. Kumar, *Introduction to Data Mining*. Addison-Wesley, 2005.
- [368] W. C. Tan, "Provenance in databases: Past, current, and future," *IEEE Data Engineering Bulletin*, vol. 30, no. 4, pp. 3–12, 2007.
- [369] N. Tang, J. X. Yu, M. T. Özsu, B. Choi, and K.-F. Wong, "Multiple materialized view selection for XPath query rewriting," in *Proceedings of the 2008 International Conference on Data Engineering*, pp. 873–882, Cancun, Mexico, April 2008.
- [370] V. Tannen, "Relational algebra," in Liu and Özsu [291], pp. 2369–2370.
- [371] I. Tatarinov, Z. G. Ives, J. Madhavan, A. Y. Halevy, D. Suciu, N. N. Dalvi, X. Dong, Y. Kadiyska, G. Miklau, and P. Mork, "The Piazza peer data management project," *ACM SIGMOD Record*, vol. 32, no. 3, pp. 47–52, 2003.
- [372] M. Teschke and A. Ulbrich, "Concurrent warehouse maintenance without compromising session consistency," in *Proceedings of the 1998 International Conference on Database and Expert Systems Applications*, pp. 776–785, Vienna, Austria, August 1998.
- [373] D. Theodoratos, "Detecting redundant materialized views in data warehouse evolution," *Information Systems*, vol. 26, no. 5, pp. 363–381, 2001.
- [374] D. Theodoratos and M. Bouzeghoub, "Data currency quality satisfaction in the design of a data warehouse," *International Journal of Cooperative Information Systems*, vol. 10, no. 3, pp. 299–326, 2001.
- [375] D. Theodoratos, S. Ligoudistianos, and T. K. Sellis, "View selection for designing the global data warehouse," *Data and Knowledge Engineering*, vol. 39, no. 3, pp. 219–240, 2001.
- [376] D. Theodoratos and T. Sellis, "Data warehouse configuration," in *Proceedings of the 1997 International Conference on Very Large Data Bases*, pp. 126–135, Athens, Greece, August 1997.
- [377] D. Theodoratos and T. Sellis, "Designing data warehouses," *Data and Knowledge Engineering*, vol. 31, pp. 279–301, 1999.
- [378] D. Theodoratos and T. K. Sellis, "Incremental design of a data warehouse," *Journal of Intelligent Information Systems*, vol. 15, no. 1, pp. 7–27, 2000.

- [379] D. Theodoratos, W. Xu, and A. Simitsis, “Materialized view selection for data warehouse design,” in *Encyclopedia of Data Warehousing and Mining*, (J. Wang, ed.), pp. 1182–1187, IGI Global, 2009.
- [380] A. Thiem and K.-U. Sattler, “An integrated approach to performance monitoring for autonomous tuning,” in *Proceedings of the 2009 International Conference on Data Engineering*, pp. 1671–1678, Shanghai, China, March 2009.
- [381] A. Thusoo, J. S. Sarma, N. Jain, Z. Shao, P. Chakka, S. Anthony, H. Liu, P. Wyckoff, and R. Murthy, “Hive — a warehousing solution over a map-reduce framework,” *Proceedings of the VLDB Endowment*, vol. 2, no. 2, pp. 1626–1629, 2009.
- [382] F. W. Tompa and J. A. Blakeley, “Maintaining materialized views without accessing base data,” *Information Systems*, vol. 13, no. 4, pp. 393–406, 1988.
- [383] O. G. Tsatalos, M. H. Solomon, and Y. E. Ioannidis, “The gmap: A versatile tool for physical data independence,” in *Proceedings of the 1994 International Conference on Very Large Data Bases*, pp. 367–378, Santiago de Chile, Chile, September 1994.
- [384] O. G. Tsatalos, M. H. Solomon, and Y. E. Ioannidis, “The GMAP: A versatile tool for physical data independence,” *The VLDB Journal*, vol. 5, no. 2, pp. 101–118, 1996.
- [385] J. D. Ullman, “Information integration using logical views,” *Theoretical Computer Science*, vol. 239, no. 2, pp. 189–210, 2000.
- [386] G. Valentin, M. Zuliani, D. C. Zilio, G. M. Lohman, and A. Skelley, “DB2 advisor: An optimizer smart enough to recommend its own indexes,” in *Proceedings of the 2000 International Conference on Data Engineering*, pp. 101–110, Los Angeles, California, USA, February 2000.
- [387] R. van der Meyden, “The complexity of querying indefinite data about linearly ordered domains,” in *Proceedings of the 1992 ACM Symposium on Principles of Database Systems*, pp. 331–345, San Diego, CA, USA, June 1992.
- [388] V. Vassalos, “Answering queries using views,” in Liu and Özsu [291], pp. 92–98.
- [389] Y. Velegrakis, “Side-effect-free view updates,” in Liu and Özsu [291], pp. 2639–2642.
- [390] Y. Velegrakis, “Updates through views,” in Liu and Özsu [291], pp. 3244–3247.
- [391] S. D. Viglas, J. F. Naughton, and J. Burger, “Maximizing the output rate of multi-way join queries over streaming information sources,” in *Proceedings of the 2003 International Conference on Very Large Data Bases*, pp. 285–296, Berlin, Germany, September 2003.
- [392] D. Vista, “Optimizing incremental view maintenance expressions in relational databases,” PhD thesis, University of Toronto, 1996.
- [393] D. Vista, “Integration of incremental view maintenance into query optimizers,” in *Proceedings of the 1998 International Conference on Extending Database Technology*, pp. 374–388, Valencia, Spain, March 1998.
- [394] J. Wang, M. J. Maher, and R. W. Topor, “Rewriting unions of general conjunctive queries using views,” in *Proceedings of the 2002 International Conference on Extending Database Technology*, pp. 52–69, Prague, Czech Republic, March 2002.

114 *References*

- [395] W. Wang, H. Lu, J. Feng, and J. X. Yu, “Condensed cube: An efficient approach to reducing data cube size,” in *Proceedings of the 2002 International Conference on Data Engineering*, pp. 155–165, San Jose, California, USA, February 2002.
- [396] G. Weikum, A. Mönkeberg, C. Hasse, and P. Zabback, “Self-tuning database technology and information services: From wishful thinking to viable engineering,” in *Proceedings of the 2002 International Conference on Very Large Data Bases*, pp. 20–31, Hong Kong, China, September 2002.
- [397] C. M. Wyss and E. L. Robertson, “Relational languages for metadata integration,” *ACM Transactions on Database Systems*, vol. 30, no. 2, pp. 624–660, 2005.
- [398] J. Xie and J. Yang, “A survey of join processing in data streams,” in Aggarwal [15], pp. 209–236.
- [399] M. Xu and C. I. Ezeife, “Maintaining horizontally partitioned warehouse views,” in *Proceedings of the 2000 International Conference on Data Warehousing and Knowledge Discovery*, pp. 126–133, London, UK, September 2000.
- [400] W. Xu, “The framework of an XML semantic caching system,” in *Proceedings of the 2005 International Workshop on the Web and Databases*, pp. 127–132, Baltimore, Maryland, USA, June 2005.
- [401] W. Xu and Z. M. Özsoyoglu, “Rewriting XPath queries using materialized views,” in *Proceedings of the 2005 International Conference on Very Large Data Bases*, pp. 121–132, Trondheim, Norway, August 2005.
- [402] W. Xu, C. Zuzarte, D. Theodoratos, and W. Ma, “Preprocessing for fast refreshing materialized views in DB2,” in *Proceedings of the 2006 International Conference on Data Warehousing and Knowledge Discovery*, pp. 55–64, Krakow, Poland, September 2006.
- [403] J. Yang, K. Karlapalem, and Q. Li, “Algorithms for materialized view design in data warehousing environment,” in *Proceedings of the 1997 International Conference on Very Large Data Bases*, pp. 136–145, Athens, Greece, August 1997.
- [404] J. Yang and J. Widom, “Incremental computation and maintenance of temporal aggregates,” in *Proceedings of the 2001 International Conference on Data Engineering*, pp. 51–60, Heidelberg, Germany, April 2001.
- [405] J. Yang and J. Widom, “Incremental computation and maintenance of temporal aggregates,” *The VLDB Journal*, vol. 12, no. 3, pp. 262–283, 2003.
- [406] K. Yi, H. Yu, J. Yang, G. Xia, and Y. Chen, “Efficient maintenance of materialized top-k views,” in *Proceedings of the 2003 International Conference on Data Engineering*, pp. 189–200, Bangalore, India, March 2003.
- [407] M. Zaharioudakis, R. Cochrane, G. Lapis, H. Pirahesh, and M. Urata, “Answering complex SQL queries using automatic summary tables,” in *Proceedings of the 2000 ACM SIGMOD International Conference on Management of Data*, pp. 105–116, Dallas, Texas, USA, May 2000.
- [408] C. Zhang, J. Yang, and K. Karlapalem, “Dynamic materialized view selection in data warehouse environment,” *Informatika (Slovenia)*, vol. 27, no. 4, pp. 451–460, 2003.

- [409] C. Zhang, X. Yao, and J. Yang, "An evolutionary approach to materialized views selection in a data warehouse environment," *IEEE Transactions on Systems, Man, and Cybernetics, Part C*, vol. 31, no. 3, pp. 282–294, 2001.
- [410] X. Zhang, K. Dimitrova, L. Wang, M. El-Sayed, B. Murphy, B. Pielech, M. Mulchandani, L. Ding, and E. A. Rundensteiner, "Rainbow: Multi-XQuery optimization using materialized XML views," in *Proceedings of the 2003 ACM SIGMOD International Conference on Management of Data*, p. 671, San Diego, California, USA, June 2003.
- [411] X. Zhang, L. Ding, and E. A. Rundensteiner, "PVM: Parallel view maintenance under concurrent data updates of distributed sources," in *Proceedings of the 2001 International Conference on Data Warehousing and Knowledge Discovery*, pp. 230–239, Munich, Germany, September 2001.
- [412] X. Zhang, L. Ding, and E. A. Rundensteiner, "Parallel multisource view maintenance," *The VLDB Journal*, vol. 13, no. 1, pp. 22–48, 2004.
- [413] X. Zhang and E. A. Rundensteiner, "DyDa: Dynamic data warehouse maintenance in a fully concurrent environment," in *Proceedings of the 2000 International Conference on Data Warehousing and Knowledge Discovery*, pp. 94–103, London, UK, September 2000.
- [414] Z. Zhang and A. O. Mendelzon, "Authorization views and conditional query containment," in *Proceedings of the 2005 International Conference on Database Theory*, pp. 259–273, Edinburgh, UK, January 2005.
- [415] J. Zhou, P.-Å. Larson, and H. G. Elmongui, "Lazy maintenance of materialized views," in *Proceedings of the 2007 International Conference on Very Large Data Bases*, pp. 231–242, Vienna, Austria, September 2007.
- [416] J. Zhou, P.-Å. Larson, J. C. Freytag, and W. Lehner, "Efficient exploitation of similar subexpressions for query processing," in *Proceedings of the 2007 ACM SIGMOD International Conference on Management of Data*, pp. 533–544, Beijing, China, June 2007.
- [417] J. Zhou, P.-Å. Larson, J. Goldstein, and L. Ding, "Dynamic materialized views," in *Proceedings of the 2007 International Conference on Data Engineering*, pp. 526–535, Istanbul, Turkey, April 2007.
- [418] Y. Zhuge, H. Garcia-Molina, J. Hammer, and J. Widom, "View maintenance in a warehousing environment," in *Proceedings of the 1995 ACM SIGMOD International Conference on Management of Data*, pp. 316–327, San Jose, California, USA, May 1995.
- [419] Y. Zhuge, H. Garcia-Molina, and J. L. Wiener, "The strobe algorithms for multi-source warehouse consistency," in *Proceedings of the 1996 International Conference on Parallel and Distributed Information Systems*, pp. 146–157, Miami Beach, Florida, USA, December 1996.
- [420] Y. Zhuge, H. Garcia-Molina, and J. L. Wiener, "Multiple view consistency for data warehousing," in *Proceedings of the 1997 International Conference on Data Engineering*, pp. 289–300, Birmingham, UK, April 1997.
- [421] Y. Zhuge, H. Garcia-Molina, and J. L. Wiener, "Consistency algorithms for multi-source warehouse view maintenance," *Distributed and Parallel Databases*, vol. 6, no. 1, pp. 7–40, 1998.

116 *References*

- [422] D. C. Zilio, J. Rao, S. Lightstone, G. M. Lohman, A. J. Storm, C. Garcia-Arellano, and S. Fadden, “DB2 Design Advisor: Integrated automatic physical database design,” in *Proceedings of the 2004 International Conference on Very Large Data Bases*, pp. 1087–1097, Toronto, Canada, August 2004.
- [423] D. C. Zilio, C. Zuzarte, S. Lightstone, W. Ma, G. M. Lohman, R. Cochrane, H. Pirahesh, L. S. Colby, J. Gryz, E. Alton, D. Liang, and G. Valentin, “Recommending views and indexes with IBM DB2 design advisor,” in *Proceedings of the 2004 International Conference on Autonomic Computing*, pp. 180–188, New York City, New York, USA, May 2004.