



Quick answers to common problems

MDX with Microsoft SQL Server 2008 R2 Analysis Services: Cookbook

80 recipes for enriching your Business Intelligence solutions with high-performance MDX calculations and flexible MDX queries

Tomislav Piasevoli

[PACKT] enterprise 
PUBLISHING professional expertise distilled

MDX with Microsoft SQL Server 2008 R2 Analysis Services: Cookbook

80 recipes for enriching your Business Intelligence solutions with high-performance MDX calculations and flexible MDX queries

Tomislav Piasevoli

[PACKT] enterprise 
PUBLISHING professional expertise distilled

BIRMINGHAM - MUMBAI

MDX with Microsoft SQL Server 2008 R2 Analysis Services: Cookbook

Copyright © 2011 Packt Publishing

All rights reserved. No part of this book may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, without the prior written permission of the publisher, except in the case of brief quotations embedded in critical articles or reviews.

Every effort has been made in the preparation of this book to ensure the accuracy of the information presented. However, the information contained in this book is sold without warranty, either express or implied. Neither the author, nor Packt Publishing, and its dealers and distributors will be held liable for any damages caused or alleged to be caused directly or indirectly by this book.

Packt Publishing has endeavored to provide trademark information about all of the companies and products mentioned in this book by the appropriate use of capitals. However, Packt Publishing cannot guarantee the accuracy of this information.

First published: July 2011

Production Reference: 2010811

Published by Packt Publishing Ltd.
Livery Place
35 Livery Street
Birmingham B3 2PB, UK.

ISBN 978-1-849681-30-8

www.packtpub.com

Cover Image by Artie Ng (artherng@yahoo.com.au)

Credits

Author

Tomislav Piasevoli

Project Coordinator

Zainab Bagasrawala

Reviewers

Greg Galloway

Darren Gosbell

Deepak Puri

Marco Russo

Chris Webb

Proofreader

Josh Toth

Indexer

Rekha Nair

Acquisition Editor

Kerry George

Production Coordinator

Arvindkumar Gupta

Development Editor

Chris Rodrigues

Cover Work

Arvindkumar Gupta

Technical Editors

Ajay Chamkeri

Merwine Machado

About the Author

Tomislav Piasevoli (tomislav@piasevoli.com) is a Business Intelligence Specialist with years of experience in Microsoft SQL Server Analysis Services (SSAS). He lives in Croatia and works for SoftPro Tetral d.o.o., a company specializing in development of SSAS frontends and implementation of BI solutions.

His main interests are dimensional modeling, cube design, and MDX about which he blogs at <http://tomislav.piasevoli.com>. Tomislav also writes articles and speaks at regional conferences and user groups. For his contribution to the community, Microsoft awarded him with the Microsoft SQL Server MVP title.

Acknowledgement

I wish to thank everyone who contributed to this book, in one way or another.

First and foremost, a big thank you goes to my wife Kristina and our three children—one of them was born while I was beginning to write this book—for having enormous patience throughout this period.

Secondly, I'd like to thank Chris Webb for vouching for me to the publisher and providing valuable information as one of the reviewers of this book. Greg Galloway, Marco Russo, Darren Gosbell, and Deepak Puri were the other precious reviewers of this book who corrected my mistakes and provided additional information relevant to the chapter topics.

Next, I'd like to thank all the people at Packt Publishing who worked on this book, to help make it better by assisting me during the book-writing process: Kerry George, Zainab Bagasrawala, Chris Rodrigues, and Merwine Machado.

There were also people who knowingly and unknowingly helped me with their ideas, articles, and helpful tips. My younger brother Hrvoje Piasevoli and Willfried Färber are representatives of the first group. The other group consists of bloggers and subject-matter experts whose articles I found useful for many recipes of this book. Here they are: Mosha Pasumansky, Teo Lachev, Jeffrey Wang, Jeremy Kashel, Vidas Matelis, Thomas Kejser (and other bloggers at SQLCAT site), Jeff Moden, Michael Coles, Itzik Ben-Gan, Irina Gorbach, Vincent Rainardi and in particular my reviewers again. Additionally, I acknowledge countless contributors to MSDN SQL Server Analysis Services forum—whom I had the luck to meet—for the solutions they shared with the rest of us there.

Last but not least, a thank you goes to my current and former colleagues at SoftPro Tetral company who played a part by exchanging ideas with me during the time spent together all these years.

About the Reviewers

Greg Galloway is a recipient of Microsoft's Most Valuable Professional, mainly for his work with Analysis Services. He has been a BI architect with Artis Consulting in Dallas, Texas, since 2003. He is a coordinator on several Codeplex projects including the well-known BIDS Helper, the Analysis Services Stored Procedure project, and OLAP PivotTable Extensions. Greg blogs at <http://www.artisconsulting.com/blogs/GregGalloway>.

Darren Gosbell is a Solution Architect with James & Monroe, a consulting firm whose specialties include the areas of Business Intelligence, Data Warehousing, and Business Performance Management (BPM). He has over ten years of practical experience in building and implementing data warehousing and business intelligence solutions on the Microsoft platform. Darren is a Microsoft MVP in SQL Server, having been a recipient of this award since 2006 in recognition of his contribution to the community specifically in the areas of Analysis Services and MDX. He lives in Melbourne, Australia with his wife and two children.

Deepak Puri is a Business Intelligence Consultant who has been working with SQL Server Analysis Services since 2000. He is an active contributor to the MSDN Analysis Services forum and a regular presenter at Ohio North SQL Server User Group meetings. Deepak has also presented at the recent Cleveland and Honolulu SQL Saturday events.

Marco Russo (marco.russo@sqlbi.com) is a consultant and trainer involved in several Business Intelligence projects, making data warehouse relational and multidimensional designs, with particular experience in sectors such as banking and financial services, manufacturing, and commercial distribution. He speaks at several international conferences and actively contributes to the SQL community. Marco is the founder of SQLBI (<http://www.sqlbi.com>) and his blog is available at http://sqlblog.com/blogs/marco_russo.

Marco has written several books: *Expert Cube Development with Microsoft SQL Server 2008 Analysis Services*, *Microsoft PowerPivot 2010 for Excel: Give Your Data Meaning*, and *The many-to-many revolution*, a mini-book about many-to-many dimension relationships in Analysis Services. He has also co-authored the *SQLBI Methodology* with Alberto Ferrari and has written several books about .NET and three books about LINQ published by Microsoft Press.

Chris Webb (chris@crossjoin.co.uk) is an independent consultant specializing in Microsoft SQL Server Analysis Services and the MDX query language. He is the co-author of the book *Expert Cube Development with SQL Server Analysis Services 2008* and blogs about Microsoft BI topics at <http://cwebbbi.wordpress.com/>.

www.PacktPub.com

Support files, eBooks, discount offers and more

You might want to visit www.PacktPub.com for support files and downloads related to your book.

Did you know that Packt offers eBook versions of every book published, with PDF and ePub files available? You can upgrade to the eBook version at www.PacktPub.com and as a print book customer, you are entitled to a discount on the eBook copy. Get in touch with us at service@packtpub.com for more details.

At www.PacktPub.com, you can also read a collection of free technical articles, sign up for a range of free newsletters and receive exclusive discounts and offers on Packt books and eBooks.



<http://PacktLib.PacktPub.com>

Do you need instant solutions to your IT questions? PacktLib is Packt's online digital book library. Here, you can access, read and search across Packt's entire library of books.

Why Subscribe?

- ▶ Fully searchable across every book published by Packt
- ▶ Copy and paste, print and bookmark content
- ▶ On demand and accessible via web browser

Free Access for Packt account holders

If you have an account with Packt at www.PacktPub.com, you can use this to access PacktLib today and view nine entirely free books. Simply use your login credentials for immediate access.

Instant Updates on New Packt Books

Get notified! Find out when new books are published by following [@PacktEnterprise](https://twitter.com/PacktEnterprise) on Twitter, or the *Packt Enterprise* Facebook page.

*I dedicate this book to my children Petra, Matko, and Nina, and to my wife Kristina.
Without their devoted support, I could not have completed this project.*

Table of Contents

Preface	1
Chapter 1: Elementary MDX Techniques	7
Introduction	8
Skipping axis	8
Handling division by zero errors	10
Setting special format for negative, zero and null values	13
Applying conditional formatting on calculations	17
Setting default member of a hierarchy in MDX script	20
Implementing NOT IN set logic	23
Implementing logical OR on members from different hierarchies	26
Iterating on a set in order to reduce it	30
Iterating on a set in order to create a new one	34
Iterating on a set using recursion	37
Dissecting and debugging MDX queries	41
Using NON_EMPTY_BEHAVIOR	45
Optimizing MDX queries using the NonEmpty() function	48
Implementing logical AND on members from the same hierarchy	51
Chapter 2: Working with Time	59
Introduction	59
Calculating the YTD (Year-To-Date) value	60
Calculating the YoY (Year-over-Year) growth (parallel periods)	66
Calculating moving averages	71
Finding the last date with data	74
Getting values on the last date with data	79
Hiding calculation values on future dates	83
Calculating today's date using the string functions	87
Calculating today's date using the MemberValue function	94
Calculating today's date using an attribute hierarchy	97

Calculating the difference between two dates	100
Calculating the difference between two times	103
Calculating parallel periods for multiple dates in a set	106
Calculating parallel periods for multiple dates in a slicer	110
Chapter 3: Concise Reporting	115
Introduction	115
Isolating the best N members in a set	116
Isolating the worst N members in a set	122
Identifying the best/worst members for each member of another hierarchy	125
Displaying few important members, others as a single row, and the total at the end	130
Combining two hierarchies into one	134
Finding the name of a child with the best/worst value	138
Highlighting siblings with the best/worst values	143
Implementing bubble-up exceptions	149
Chapter 4: Navigation	155
Introduction	155
Detecting a particular member in a hierarchy	156
Detecting the root member	160
Detecting members on the same branch	164
Finding related members in the same dimension	172
Finding related members in another dimension	178
Calculating various percentages	183
Calculating various averages	189
Calculating various ranks	194
Chapter 5: Business Analytics	205
Introduction	205
Forecasting using the linear regression	206
Forecasting using the periodic cycles	212
Allocating the non-allocated company expenses to departments	219
Calculating the number of days from the last sales to identify the slow-moving goods	227
Analyzing fluctuation of customers	233
Implementing the ABC analysis	241
Chapter 6: When MDX is Not Enough	247
Introduction	247
Using a new attribute to separate members on a level	249
Using a distinct count measure to implement histograms over existing hierarchies	254

Using a dummy dimension to implement histograms over non-existing hierarchies	257
Creating a physical measure as a placeholder for MDX assignments	266
Using a new dimension to calculate the most frequent price	271
Using a utility dimension to implement flexible display units	275
Using a utility dimension to implement time-based calculations	281
Chapter 7: Context-aware Calculations	295
Introduction	295
Identifying the number of columns and rows a query will return	297
Identifying the axis with measures	301
Identifying the axis without measures	304
Adjusting the number of columns and rows for OWC and Excel	307
Identifying the content of axes	310
Calculating row numbers	316
Calculating the bit-string for hierarchies on an axis	319
Preserving empty rows	324
Implementing utility dimension with context-aware calculations	327
Chapter 8: Advanced MDX Topics	337
Introduction	337
Displaying members without children (leaves)	338
Displaying members with data in parent-child hierarchies	343
Implementing the Tally table utility dimension	347
Displaying random values	352
Displaying a random sample of hierarchy members	356
Displaying a sample from a random hierarchy	359
Performing complex sorts	366
Using recursion to calculate cumulative values	373
Chapter 9: On the Edge	381
Introduction	381
Clearing the Analysis Services cache	382
Using Analysis Services stored procedures	387
Executing MDX queries in T-SQL environments	394
Using SSAS Dynamic Management Views (DMV) to fast-document a cube	400
Using SSAS Dynamic Management Views (DMVs) to monitor activity and usage	406
Capturing MDX queries generated by SSAS front-ends	411
Performing custom drillthrough	416
Conclusion	423

Appendix: Glossary of Terms	427
Parts of an MDX query	427
MDX query in action	433
Cube and dimension design	437
MDX script	440
Query optimization	442
Types of query	445
Index	449

Preface

MDX-related books often dedicate a significant part of their content to explaining the concepts of multidimensional cubes, the MDX language and its functions, and other specifics related to working with Analysis Services. And that's perfectly fine, there should be books like that, the tutorials that teach the concepts. However, that also means that when it comes to examples, there's usually not enough space to provide all the details about them and their variations, otherwise the book would become huge or oftentimes lose its focus. The result of that is that making a step further from the provided calculations and queries might not be an easy task for an average reader.

The other problem with tutorials is that the solution to a particular problem might be scattered throughout the book, which is where the cookbook style of books like this one come into play. Similar to data warehouses where we consolidate many tables of the relational database into a few and then organize those dimension tables in a topic-based star schema, in cookbooks we aggregate the information about a particular problem in form of one or more recipes and present that topic-based knowledge in full detail.

Both the relational databases and the data warehouses have their purpose; it's not uncommon to see them together in one place. The same is true about books and their approaches. What we also know is that there are far too few data warehouses than relational databases. Again, the same is with MDX-related cookbooks in respect to MDX tutorials, particularly those dealing with advanced topics.

As a writer, I hope you recognize my intention and the value this book could bring you. As a reader, we rarely have enough time to start reading a book, not to mention finish it. This is another advantage of the cookbook format. You can browse through the contents and look for the solution to a particular problem. As the recipes are relatively short and organized in chapters, that task won't take much of your time. All the information will be in one place. In addition to that, you'll see which recipes are related to that one, so that you can learn even more.

The next time you encounter a problem, you might start on a completely different part of the book. Little by little, you'll read it all and hopefully become confident not only in using these recipes, but also those not present in this book. In other words, the moment you start creating your own solutions by combining principles laid out in this book, my goal of teaching through examples is accomplished. Feel free to drop me a note about it.

What this book covers

This book is a cookbook style of book, a book full of solutions with tips and techniques for writing advanced MDX calculations and queries. The recipes are organized in chapters, each covering one specific area to help you navigate the contents, find a recipe to the problem that you currently have and focus entirely on a particular topic by reading related or even all recipes in that chapter.

Here's a brief overview of the book's contents.

Chapter 1, Elementary MDX Techniques, illustrates several ways of how to perform common MDX-related tasks such as iterating on sets or applying Boolean logic on them, how to check for errors, and how to format, optimize, and debug your MDX calculations.

Chapter 2, Working with Time, covers the often-required time calculations such as the year-to-date, year-over-year, and the moving averages calculations, shows how to get the last date with data or calculate the difference between two events, and also presents several solutions to the problem of using today's date in your calculations.

Chapter 3, Concise Reporting, is all about how to do more with less - how to make your queries return only what matters by identifying important members and then either isolating them from the rest of the set, combining them with other members or other hierarchies, highlighting them, extracting information about them, or creating custom groups.

Chapter 4, Navigation, explores the benefits of having multilevel hierarchies and correct attribute relationships with recipes that guide you how to identify the current context, find related members, and implement typical calculations like relative percentages, averages, or ranks.

Chapter 5, Business Analytics, presents some of the typical business requests and a way to solve them in your cubes using MDX calculations and queries.

Chapter 6, When MDX is Not Enough, suggests rethinking where to solve the problem, using MDX or by modifying the cube by implementing relatively simple improvements such as adding a new attribute, new measure, new utility dimension, or using the new granularity.

Chapter 7, Context-aware Calculations, deals with calculations that can work on any cube, dimension, hierarchy, level, or measure and ends with a recipe how to implement them using a utility dimension.

Chapter 8, Advanced MDX Topics, is a collection of recipes that cover parent-child hierarchies, random values, and complex sorts and iterations.

Chapter 9, On the Edge, looks at expanding your horizons with recipes covering data management views (DMVs), stored procedures, drillthrough, and other exotic areas.

Conclusion, is a wrap-up chapter with suggestions for additional contents on the web you might find interesting after reading this book.

Appendix, contains the glossary of terms covered in this book or relevant to MDX in general.

What you need for this book

To follow the examples in this book you should have the following installed on your computer:

- ▶ Microsoft Windows XP or later for the workstation or Microsoft Windows Server 2003 or later for the server.
- ▶ Microsoft SQL Server 2008 R2 (full installation or at least the database engine, Analysis Services and the client components).
- ▶ Adventure Works sample database, built specifically for the Microsoft SQL Server 2008 R2.

Having an older version of Microsoft SQL Server (2005 or 2008) is also fine, most of the recipes will work. However, some of them will need adjustments because the Date dimension in respective older versions of the Adventure Works database has different set of years. To solve that problem simply shift the date-specific parts of the queries few years back in time, for example, turn the year 2006 into the year 2002 and Q3 of the year 2007 to Q3 of 2003.

Also, use the Developer, Enterprise, or the Trial Edition of Microsoft SQL Server together with the enterprise version of the Adventure Works cube project. Using the Standard Edition is also fine, but a few recipes might not work. If they are important to you, consider installing the above mentioned Trial Edition because it's free for 180 days.

Here's the link for the Trial Edition of Microsoft SQL Server 2008 R2:

<http://tinyurl.com/SQLServer2008R2Trial>

Here's the link for the Adventure Works sample database built specifically for Microsoft SQL Server 2008 R2:

<http://tinyurl.com/AdventureWorks2008R2>