

# Contents

Data Management Technology Kairometer: The Historical Context xxvii

Foreword xxix

Preface xxxi

Introduction 1

*Michael L. Brodie*

A Brief History of Databases 1

Preparing to Read the Stories and What You Might Find There 6

A Travel Guide to Software Systems Lessons in Nine Parts 7

## **PART I 2014 ACM A.M. TURING AWARD PAPER AND LECTURE 13**

The Land Sharks Are on the Squawk Box 15

*Michael Stonebraker*

Off to a Good Start 16

First Speedbumps 22

Another High 26

The High Does Not Last 28

The Future Looks Up (Again) 30

The Good Times Do Not Last Long 30

The Stories End 31

Why a Bicycle Story? 32

The Present Day 35

References 36

## **PART II MIKE STONEBRAKER'S CAREER 39**

### **Chapter 1 Make it Happen: The Life of Michael Stonebraker 41**

*Samuel Madden*

Synopsis 41

Early Years and Education 42

Academic Career and the Birth of Ingres 43

The Post-Ingres Years 45

Industry, MIT, and the New Millennium 46

Stonebraker's Legacy 47

Companies 48

Awards and Honors 49

Service 49

Advocacy 50

Personal Life 50

Acknowledgments 50

Mike Stonebraker's Student Genealogy Chart 53

The Career of Mike Stonebraker: The Chart 55

## **PART III MIKE STONEBRAKER SPEAKS OUT: AN INTERVIEW WITH MARIANNE WINSLETT 57**

### **Chapter 2 Mike Stonebraker Speaks Out: An Interview 59**

*Marianne Winslett*

## **PART IV THE BIG PICTURE 85**

### **Chapter 3 Leadership and Advocacy 87**

*Philip A. Bernstein*

Systems 87

Mechanisms 90

Advocacy 91

### **Chapter 4 Perspectives: The 2014 ACM Turing Award 93**

*James Hamilton*

**Chapter 5 Birth of an Industry; Path to the Turing Award 97***Jerry Held*

- Birth of an Industry (1970s) 97
- Ingres—Timing 98
- Ingres—Team 99
- Ingres—Competition 100
- Ingres—Platform 101
- Adolescence with Competition (1980s and 1990s) 101
- Competing with Oracle 102
- Competing with Oracle (Again) 102
- Maturity with Variety (2000s and 2010s) 103
- Vertica 104
- VoltDB 104
- Tamr 105
- The Bottom Line 105

**Chapter 6 A Perspective of Mike from a 50-Year Vantage Point 107***David J. DeWitt*

- Fall 1970—University of Michigan 107
- Fall 1976—Wisconsin 108
- Fall 1983—Berkeley 111
- 1988–1995—No Object Oriented DBMS Detour for Mike 111
- 2000—Project Sequoia 112
- 2003—CIDR Conference Launch 113
- 2005—Sabbatical at MIT 113
- 2008—We Blog about “MapReduce” 114
- 2014—Finally, a Turing Award 114
- 2016—I Land at MIT 115
- 2017 115

**PART V STARTUPS 117****Chapter 7 How to Start a Company in Five (Not So) Easy Steps 119***Michael Stonebraker*

- Introduction 119
- Step 1: Have a Good Idea 119

Step 2: Assemble a Team and Build a Prototype	120
Step 3: Find a Lighthouse Customer	122
Step 4: Recruit Adult Supervision	122
Step 5: Prepare a Pitch Deck and Solicit the VCs	123
Comments	125
Summary	128

**Chapter 8 How to Create and Run a Stonebraker Startup— The *Real* Story 129**

*Andy Palmer*

An Extraordinary Achievement. An Extraordinary Contribution.	130
A Problem of Mutual Interest A Happy Discovery	132
The Power of Partnership	133
Fierce Pragmatism, Unwavering Clarity, Boundless Energy	135
A Final Observation: Startups are Fundamentally about People	138

**Chapter 9 Getting Grownups in the Room: A VC Perspective 139**

*Jo Tango*

My First Meeting	139
Context	139
StreamBase	140
A Playbook Is Set	142
Mike's Values	143
A Coda	143
A Great Day	144

**PART VI DATABASE SYSTEMS RESEARCH 145**

**Chapter 10 Where Good Ideas Come From and How to Exploit Them 147**

*Michael Stonebraker*

Introduction	147
The Birth of Ingres	147
Abstract Data Types (ADTs)	148
Postgres	149
Distributed Ingres, Ingres*, Cohera, and Morpheus	150
Parallel Databases	151
Data Warehouses	151

	H-Store/VoltDB	151
	Data Tamer	152
	How to Exploit Ideas	153
	Closing Observations	153
<b>Chapter 11</b>	<b>Where We Have Failed</b>	<b>155</b>
	<i>Michael Stonebraker</i>	
	The Three Failures	155
	Consequences of Our Three Failures	160
	Summary	164
<b>Chapter 12</b>	<b>Stonebraker and Open Source</b>	<b>165</b>
	<i>Mike Olson</i>	
	The Origins of the BSD License	165
	BSD and Ingres	166
	The Impact of Ingres	167
	Post-Ingres	168
	The Impact of Open Source on Research	169
<b>Chapter 13</b>	<b>The Relational Database Management Systems Genealogy</b>	<b>173</b>
	<i>Felix Naumann</i>	
<b>PART VII</b>	<b>CONTRIBUTIONS BY SYSTEM</b>	<b>181</b>
<b>Chapter 14</b>	<b>Research Contributions of Mike Stonebraker: An Overview</b>	<b>183</b>
	<i>Samuel Madden</i>	
	Technical Rules of Engagement with Mike	183
	Mike's Technical Contributions	185
<b>PART VII.A</b>	<b>RESEARCH CONTRIBUTIONS BY SYSTEM</b>	<b>191</b>
<b>Chapter 15</b>	<b>The Later Ingres Years</b>	<b>193</b>
	<i>Michael J. Carey</i>	
	How I Ended Up at the Ingres Party	193
	Ingres: Realizing (and Sharing!) a Relational DBMS	194

Distributed Ingres: One Was Good, So More Must be Better 198  
Ingres: Moving Beyond Business Data 200

**Chapter 16 Looking Back at Postgres 205**

*Joseph M. Hellerstein*

Context 205  
Postgres: An Overview 206  
Log-centric Storage and Recovery 213  
Software Impact 218  
Lessons 223  
Acknowledgments 224

**Chapter 17 Databases Meet the Stream Processing Era 225**

*Magdalena Balazinska, Stan Zdonik*

Origins of the Aurora and Borealis Projects 225  
The Aurora and Borealis Stream-Processing Systems 227  
Concurrent Stream-Processing Efforts 231  
Founding StreamBase Systems 232  
Stream Processing Today 233  
Acknowledgments 234

**Chapter 18 C-Store: Through the Eyes of a Ph.D. Student 235**

*Daniel J. Abadi*

How I Became a Computer Scientist 235  
The Idea, Evolution, and Impact of C-Store 238  
Building C-Store with Mike 240  
Founding Vertica Systems 242

**Chapter 19 In-Memory, Horizontal, and Transactional:  
The H-Store OLTP DBMS Project 245**

*Andy Pavlo*

System Architecture Overview 246  
First Prototype (2006) 247  
Second Prototype (2007–2008) 247  
VoltDB (2009–Present) 250  
H-Store/VoltDB Split (2010–2016) 251  
Conclusion 251

**Chapter 20 Scaling Mountains: SciDB and Scientific Data Management 253***Paul Brown*

- Selecting Your Mountain 254
- Planning the Climb 256
- Expedition Logistics 259
- Base Camp 260
- Plans, Mountains, and Altitude Sickness 263
- On Peaks 267
- Acknowledgments 268

**Chapter 21 Data Unification at Scale: Data Tamer 269***Ihab Ilyas*

- How I Got Involved 269
- Data Tamer: The Idea and Prototype 270
- The Company: Tamr Inc. 273
- Mike's Influence: Three Lessons Learned. 276

**Chapter 22 The BigDAWG Polystore System 279***Tim Mattson, Jennie Rogers, Aaron J. Elmore*

- Big Data ISTC 279
- The Origins of BigDAWG 280
- One Size Does Not Fit All and the Quest for Polystore Systems 282
- Putting it All Together 284
- Query Modeling and Optimization 285
- Data Movement 286
- BigDAWG Releases and Demos 287
- Closing Thoughts 288

**Chapter 23 Data Civilizer: End-to-End Support for Data Discovery, Integration, and Cleaning 291***Mourad Ouzzani, Nan Tang, Raul Castro Fernandez*

- We Need to Civilize the Data 292
- The Day-to-Day Life of an Analyst 292
- Designing an End-to-End System 294
- Data Civilizer Challenges 295
- Concluding Remarks 300

**PART VII.B CONTRIBUTIONS FROM BUILDING SYSTEMS 301**

**Chapter 24 The Commercial Ingres Codeline 303**

*Paul Butterworth, Fred Carter*

Research to Commercial 304

Conclusions 309

Open Source Ingres 309

**Chapter 25 The Postgres and Illustra Codelines 311**

*Wei Hong*

Postgres: The Academic Prototype 311

Illustra: “Doing It for Dollars” 313

PostgreSQL and Beyond 317

Open Source PostgreSQL 318

Final Thoughts 319

**Chapter 26 The Aurora/Borealis/ StreamBase Codelines: A Tale of Three Systems 321**

*Nesime Tatbul*

Aurora/Borealis: The Dawn of Stream Processing Systems 322

From 100K+ Lines of University Code to a Commercial Product 326

Encounters with StreamBase Customers 327

“Over My Dead Body” Issues in StreamBase 328

An April Fool’s Day Joke, or the Next Big Idea? 330

Concluding Remarks 331

Acknowledgments 332

**Chapter 27 The Vertica Codeline 333**

*Shilpa Lawande*

Building a Database System from Scratch 333

Code Meets Customers 334

Don’t Reinvent the Wheel (Make It Better) 335

Architectural Decisions: Where Research Meets Real Life 336

Customers: The Most Important Members of the Dev Team 339

Conclusion 340

Acknowledgments 340



**Chapter 28 The VoltDB Codeline 341***John Hugg*

- Compaction 342
- Latency 344
- Disk Persistence 346
- Latency Redux 347
- Conclusion 348

**Chapter 29 The SciDB Codeline: Crossing the Chasm 349***Kriti Sen Sharma, Alex Poliakov, Jason Kinchen*

- Playing Well with Others 349
- You Can't Have Everything (at Once) 351
- In Hard Numbers We Trust 352
- Language Matters 353
- Security is an Ongoing Process 354
- Preparing for the (Genomic) Data Deluge 354
- Crossing the Chasm: From Early Adopters to Early Majority 355

**Chapter 30 The Tamr Codeline 357***Nikolaus Bates-Haus*

- Neither Fish nor Fowl 358
- Taming the Beast of Algorithmic Complexity 359
- Putting Users Front and Center 361
- Scaling with Respect to Variety 362
- Conclusion 365

**Chapter 31 The BigDAWG Codeline 367***Vijay Gadepally*

- Introduction 367
- BigDAWG Origins 370
- First Public BigDAWG Demonstration 371
- Refining BigDAWG 373
- BigDAWG Official Release 375
- BigDAWG Future 376

**PART VIII PERSPECTIVES 377**

**Chapter 32 IBM Relational Database Code Bases 379**

*James Hamilton*

Why Four Code Bases? 379

The Portable Code Base Emerges 381

Looking Forward 384

**Chapter 33 Aurum: A Story about Research Taste 387**

*Raul Castro Fernandez*

**Chapter 34 Nice: Or What It Was Like to Be Mike's Student 393**

*Marti Hearst*

**Chapter 35 Michael Stonebraker: Competitor, Collaborator, Friend 397**

*Don Haderle*

**Chapter 36 The Changing of the Database Guard 403**

*Michael L. Brodie*

Dinner with the Database Cognoscenti 403

The *Great* Relational-CODASYL Debate 404

Mike: More Memorable than the Debate, and Even the Cheese 406

A Decade Later: Friend or Foe? 407

**PART IX SEMINAL WORKS OF MICHAEL STONEBRAKER AND HIS COLLABORATORS 409**

OLTP Through the Looking Glass, and What We Found There 411

*Stavros Harizopoulos, Daniel J. Abadi,*

*Samuel Madden, Michael Stonebraker*

Abstract 411

1 Introduction 412

2 Trends in OLTP 416

3 Shore 418

4	Performance Study	424
5	Implications for Future OLTP Engines	433
6	Related Work	436
7	Conclusions	436
8	Acknowledgments	437
9	Repeatability Assessment	437
	References	437

**“One Size Fits All”: An Idea Whose Time Has Come and Gone 441**

*Michael Stonebraker*

*Uğur Çetintemel*

	Abstract	441
1	Introduction	441
2	Data Warehousing	443
3	Stream Processing	445
4	Performance Discussion	448
5	One Size Fits All?	455
6	A Comment on Factoring	458
7	Concluding Remarks	460
	References	460

**The End of an Architectural Era (It’s Time for a Complete Rewrite) 463**

*Michael Stonebraker, Samuel Madden,*

*Daniel J. Abadi, Stavros Harizopoulos,*

*Nabil Hachem,*

*Pat Helland*

	Abstract	463
1	Introduction	464
2	OLTP Design Considerations	466
3	Transaction, Processing and Environment Assumptions	470
4	H-Store Sketch	473
5	A Performance Comparison	479
6	Some Comments about a “One Size Does Not Fit All” World	483
7	Summary and Future Work	486
	References	488

C-Store: A Column-Oriented DBMS 491

*Mike Stonebraker, Daniel J. Abadi,  
Adam Batkin, Xuedong Chen,  
Mitch Cherniack, Miguel Ferreira,  
Edmond Lau, Amerson Lin,  
Sam Madden, Elizabeth O'Neil,  
Pat O'Neil, Alex Rasin,  
Nga Tran, Stan Zdonik*

	<b>Abstract</b>	<b>491</b>
1	<b>Introduction</b>	<b>492</b>
2	<b>Data Model</b>	<b>496</b>
3	<b>RS</b>	<b>500</b>
4	<b>WS</b>	<b>502</b>
5	<b>Storage Management</b>	<b>502</b>
6	<b>Updates and Transactions</b>	<b>503</b>
7	<b>Tuple Mover</b>	<b>508</b>
8	<b>C-Store Query Execution</b>	<b>509</b>
9	<b>Performance Comparison</b>	<b>511</b>
10	<b>Related Work</b>	<b>515</b>
11	<b>Conclusions</b>	<b>516</b>
	<b>Acknowledgements and References</b>	<b>517</b>

The Implementation of POSTGRES 519

*Michael Stonebraker, Lawrence A. Rowe, Michael Hirohama*

I	<b>Introduction</b>	<b>520</b>
II	<b>The POSTGRES Data Model and Query Language</b>	<b>521</b>
III	<b>The Rules System</b>	<b>538</b>
IV	<b>Storage System</b>	<b>547</b>
V	<b>The POSTGRES Implementation</b>	<b>550</b>
VI	<b>Status and Performance</b>	<b>554</b>
VII	<b>Conclusions</b>	<b>555</b>
	<b>References</b>	<b>557</b>

The Design and Implementation of INGRES 561

*Michael Stonebraker,  
Eugene Wong,*

*Peter Kreps,  
Gerald Held*

1	Introduction	562
2	The INGRES Process Structure	571
3	Data Structures and Access Methods	575
4	The Structure of Process 2	585
5	Process 3	591
6	Utilities in Process 4	599
7	Conclusion and Future Extensions	602
	Acknowledgment	603
	References	603

The Collected Works of Michael Stonebraker 607

References 635

Index 645

Biographies 671