

# Declarative Networking



# Synthesis Lectures on Data Management

## Editor

M. Tamer Özsu, *University of Waterloo*

Synthesis Lectures on Data Management is edited by Tamer Özsu of the University of Waterloo. The series will publish 50- to 125 page publications on topics pertaining to data management. The scope will largely follow the purview of premier information and computer science conferences, such as ACM SIGMOD, VLDB, ICDE, PODS, ICDT, and ACM KDD. Potential topics include, but not are limited to: query languages, database system architectures, transaction management, data warehousing, XML and databases, data stream systems, wide scale data distribution, multimedia data management, data mining, and related subjects.

## Declarative Networking

Boon Thau Loo and Wenchao Zhou

2012

## Full-Text (Substring) Indexes in External Memory

Marina Barsky, Ulrike Stege, and Alex Thomo

2011

## Spatial Data Management

Nikos Mamoulis

2011

## Database Repairing and Consistent Query Answering

Leopoldo Bertossi

2011

## Managing Event Information: Modeling, Retrieval, and Applications

Amarnath Gupta and Ramesh Jain

2011

## Fundamentals of Physical Design and Query Compilation

David Toman and Grant Weddell

2011

[Methods for Mining and Summarizing Text Conversations](#)  
Giuseppe Carenini, Gabriel Murray, and Raymond Ng  
2011

[Probabilistic Databases](#)  
Dan Suciu, Dan Olteanu, Christopher Ré, and Christoph Koch  
2011

[Peer-to-Peer Data Management](#)  
Karl Aberer  
2011

[Probabilistic Ranking Techniques in Relational Databases](#)  
Ihab F. Ilyas and Mohamed A. Soliman  
2011

[Uncertain Schema Matching](#)  
Avigdor Gal  
2011

[Fundamentals of Object Databases: Object-Oriented and Object-Relational Design](#)  
Suzanne W. Dietrich and Susan D. Urban  
2010

[Advanced Metasearch Engine Technology](#)  
Weiyi Meng and Clement T. Yu  
2010

[Web Page Recommendation Models: Theory and Algorithms](#)  
Sule Gündüz-Ögüdücü  
2010

[Multidimensional Databases and Data Warehousing](#)  
Christian S. Jensen, Torben Bach Pedersen, and Christian Thomsen  
2010

[Database Replication](#)  
Bettina Kemme, Ricardo Jimenez Peris, and Marta Patino-Martinez  
2010

[Relational and XML Data Exchange](#)  
Marcelo Arenas, Pablo Barcelo, Leonid Libkin, and Filip Murlak  
2010

### User-Centered Data Management

Tiziana Catarci, Alan Dix, Stephen Kimani, and Giuseppe Santucci

2010

### Data Stream Management

Lukasz Golab and M. Tamer Özsu

2010

### Access Control in Data Management Systems

Elena Ferrari

2010

### An Introduction to Duplicate Detection

Felix Naumann and Melanie Herschel

2010

### Privacy-Preserving Data Publishing: An Overview

Raymond Chi-Wing Wong and Ada Wai-Chee Fu

2010

### Keyword Search in Databases

Jeffrey Xu Yu, Lu Qin, and Lijun Chang

2009

© Springer Nature Switzerland AG 2022

Reprint of original edition © Morgan & Claypool 2012

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—electronic, mechanical, photocopy, recording, or any other except for brief quotations in printed reviews, without the prior permission of the publisher.

Declarative Networking

Boon Thau Loo and Wenchao Zhou

ISBN: 978-3-031-00758-3      paperback

ISBN: 978-3-031-01886-2      ebook

DOI 10.1007/978-3-031-01886-2

A Publication in the Springer series

*SYNTHESIS LECTURES ON DATA MANAGEMENT*

Lecture #23

Series Editor: M. Tamer Özsu, *University of Waterloo*

Series ISSN

Synthesis Lectures on Data Management

Print 2153-5418    Electronic 2153-5426

# Declarative Networking

Boon Thau Loo and Wenchao Zhou  
University of Pennsylvania

*SYNTHESIS LECTURES ON DATA MANAGEMENT #23*

## ABSTRACT

*Declarative Networking* is a programming methodology that enables developers to concisely specify network protocols and services, which are directly compiled to a dataflow framework that executes the specifications. Declarative networking proposes the use of a declarative query language for specifying and implementing network protocols, and employs a dataflow framework at runtime for communication and maintenance of network state. The primary goal of declarative networking is to greatly simplify the process of specifying, implementing, deploying and evolving a network design. In addition, declarative networking serves as an important step towards an extensible, evolvable network architecture that can support *flexible, secure* and *efficient* deployment of new network protocols.

This book provides an introduction to basic issues in declarative networking, including language design, optimization and dataflow execution. The methodology behind declarative programming of networks is presented, including roots in Datalog, extensions for networked environments, and the semantics of long-running queries over network state. The book focuses on a representative declarative networking language called *Network Datalog* (NDlog), which is based on extensions to the Datalog recursive query language. An overview of declarative network protocols written in NDlog is provided, and its usage is illustrated using examples from routing protocols and overlay networks.

This book also describes the implementation of a declarative networking engine and NDlog execution strategies that provide eventual consistency semantics with significant flexibility in execution. Two representative declarative networking systems (P2 and its successor RapidNet) are presented. Finally, the book highlights recent advances in declarative networking, and new declarative approaches to related problems.

## KEYWORDS

declarative networking, datalog, recursive query processing

*To Foo Lee Yoong*

(Boon Thau Loo's mother)

*and*

*To Yiqing Ren*

(Wenchao Zhou's wife)

February 2012.



# Contents

	<b>Acknowledgments</b> .....	<b>xv</b>
	<b>Figure Credits</b> .....	<b>xvii</b>
<b>1</b>	<b>Introduction</b> .....	<b>1</b>
1.1	Overview of Declarative Networks .....	2
1.2	The Case for Declarative Networking .....	3
1.2.1	Ease of Programming .....	3
1.2.2	Optimizability .....	4
1.2.3	Balance of Extensibility and Safety .....	4
1.3	Organization .....	4
<b>2</b>	<b>Declarative Networking Language</b> .....	<b>7</b>
2.1	Introduction to Datalog .....	7
2.2	Network Datalog by Example .....	8
2.2.1	Overview of NDlog .....	9
2.2.2	From Query Specifications to Protocol Execution .....	10
2.2.3	Other Requirements of NDlog .....	13
2.3	Distributed Computation .....	13
2.4	Soft-state Data and Rules .....	14
2.4.1	Hard-state vs. Soft-state Data .....	15
2.4.2	Hard-state and Soft-state Rules .....	16
2.5	Incremental Maintenance of Network State .....	17
2.6	Summary of Network Datalog .....	18
2.7	Summary .....	18
<b>3</b>	<b>Declarative Networking Overview</b> .....	<b>21</b>
3.1	Architecture .....	21
3.2	DN Dataflow Engine .....	22
3.2.1	Dataflow elements .....	24
3.3	Network State Storage and Management .....	25
3.3.1	RapidNet Declarative Networking Engine .....	25

3.4	Summary .....	27
<b>4</b>	<b>Distributed Recursive Query Processing .....</b>	<b>29</b>
4.1	Centralized Plan Generation .....	29
4.1.1	Semi-naïve Evaluation .....	29
4.1.2	Dataflow Generation .....	31
4.2	Distributed Plan Generation .....	32
4.2.1	Localization Rewrite .....	32
4.2.2	Distributed Dataflow Generation .....	33
4.3	Relaxing Semi-naïve Evaluation .....	34
4.3.1	Pipelined Semi-naïve Evaluation .....	34
4.4	Processing in a Dynamic Network .....	36
4.4.1	Dataflow Generation for Incremental View Maintenance .....	37
4.4.2	Centralized Execution Semantics .....	39
4.4.3	Distributed Execution Semantics .....	40
4.5	Processing Soft-state Rules .....	41
4.5.1	Event Soft-state Rules .....	43
4.6	Summary .....	44
<b>5</b>	<b>Declarative Routing .....</b>	<b>45</b>
5.1	Motivation .....	45
5.2	Execution Model .....	46
5.3	Routing Protocols By Examples .....	48
5.3.1	Best-Path Routing .....	48
5.3.2	Distance-Vector Routing .....	49
5.3.3	Policy-Based Routing .....	50
5.3.4	Dynamic Source Routing .....	50
5.3.5	Link State .....	51
5.3.6	Multicast .....	51
5.4	Security Issues .....	52
5.5	Route Maintenance .....	53
5.6	Evaluation .....	55
5.6.1	Scalability of Path-Vector Protocol .....	55
5.6.2	Incremental Evaluation in Dynamic Networks .....	56
5.7	Summary .....	58

<b>6</b>	<b>Declarative Overlays</b> .....	<b>59</b>
6.1	Execution Model .....	59
6.2	Narada Mesh .....	60
6.2.1	Membership List Maintenance .....	62
6.2.2	Neighbor Selection .....	63
6.3	Chord Distributed Hash Table .....	64
6.3.1	Chord Network State .....	65
6.3.2	Joining the Chord Network .....	67
6.3.3	Chord Ring Maintenance .....	68
6.3.4	Finger Maintenance and Routing .....	69
6.3.5	Failure Detection .....	70
6.3.6	Summary of Chord .....	71
6.4	Evaluation .....	72
6.4.1	Narada Mesh Formation .....	72
6.4.2	Chord DHT .....	73
6.5	Summary .....	76
<b>7</b>	<b>Optimization of NDlog Programs</b> .....	<b>77</b>
7.1	Traditional Datalog Optimizations .....	77
7.1.1	Aggregate Selections .....	77
7.1.2	Magic Sets and Predicate Reordering .....	78
7.2	Multi-Query Optimizations .....	80
7.3	Hybrid Rewrites .....	81
7.4	Evaluation of Optimizations .....	82
7.4.1	Aggregate Selections .....	82
7.4.2	Magic Sets and Predicate Reordering .....	84
7.4.3	Opportunistic Message Sharing .....	85
7.4.4	Summary of Optimizations .....	85
7.5	Summary .....	85
<b>8</b>	<b>Recent Advances in Declarative Networking</b> .....	<b>87</b>
8.1	Language Extensions .....	87
8.2	Generating Safe Routing Implementations .....	88
8.2.1	Formally Safe Routing Toolkit .....	88
8.2.2	Declarative Network Verification .....	89
8.3	Securing Distributed Systems .....	89

8.3.1	Secure Network Datalog	90
8.3.2	Reconfigurable Security	92
8.3.3	Application-aware Anonymity	92
8.4	Debugging Distributed Systems	93
8.4.1	Network Provenance Model	93
8.4.2	Distributed Maintenance and Querying	94
8.4.3	Security and Temporal Extensions	95
8.5	Optimizing Distributed Systems	96
8.5.1	Use Cases: PUMA and COPE	96
8.5.2	Colog Language and Compilation	97
8.6	Summary	98
<b>9</b>	<b>Conclusion</b>	<b>99</b>
	<b>Bibliography</b>	<b>101</b>
	<b>Authors' Biographies</b>	<b>111</b>

# Acknowledgments

The technical contents from Chapter 1-7 are derived from Boon Thau Loo's doctoral dissertation [[Loo, 2006](#)] (advisors: Joseph M. Hellerstein and Ion Stoica, with significant input from Raghu Ramakrishnan), and collectively contributed by members of the P2 declarative networking team [[P2](#)] at the University of California Berkeley, and Intel Labs. Declarative networking research at the University of California Berkeley was funded by NSF grants (IIS-0205647, IIS-0209108, and ANI-0225660) and a gift from Microsoft.

At the University of Pennsylvania, our work on declarative networking has been funded by NSF (CNS-0721845, CNS-0831376, IIS-0812270, CCF-0820208, CNS-0845552, CNS-1040672, CNS-1065130, and CNS-1117052), AFOSR MURI grant FA9550-08-1-0352, DARPA SAFER award N66001-11-C-4020, and DARPA Air Force Research Laboratory (AFRL) Contract FA8750-07-C-0169. We thank our collaborators listed on the NetDB@Penn site [[NetDB@Penn](#)] for their contributions to the various research efforts described in this book, in particular, Chapter 8.

Finally, we thank Tamer Özsu for his useful comments and input into this book. We also thank Harjot Gill, Alex Gurney, Anduo Wang, and Zhuoyao Zhang for their careful proofreading of this book.

Boon Thau Loo and Wenchao Zhou  
February 2012

# Figure Credits

- Figure 3.4 *is from* Shivkumar C. Muthukumar, Xiaozhou Li, Changbin Liu, Joseph B. Kopena, Mihai Oprea, Ricardo Correa, Boon Thau Loo, and Prithwish Basu. RapidMesh: declarative toolkit for rapid experimentation of wireless mesh networks. In *Proc. ACM Int. Workshop on Wireless Network Testbeds, Experimental Evaluation and Characterization*, pages 1–10, 2009a. Copyright © 2009 ACM. Used with permission. And Shivkumar C. Muthukumar, Xiaozhou Li, Changbin Liu, Joseph B. Kopena, Mihai Oprea, and Boon Thau Loo. Declarative toolkit for rapid network protocol simulation and experimentation. In *Proc. ACM Int. Conf. on Data Communication*, pages 1–2, 2009b. Copyright © 2009 ACM. Used with permission.
- Figure 8.4 *is from* Wenchao Zhou, Qiong Fei, Shengzhi Sun, Tao Tao, Andreas Haeberlen, Zachary Ives, Boon Thau Loo, and Micah Sherr. NetTrails: A Declarative Platform for Provenance Maintenance and Querying in Distributed Systems. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 1323–1326, 2011c. Copyright ©2011 ACM. Used with permission.