

The Provenance of Elegance in Computation – Essays Dedicated to Val Tannen

Tannen's Festschrift, May 24–25, 2024,
University of Pennsylvania, Philadelphia, PA, USA

Edited by

Antoine Amarilli

Alin Deutsch



Editors

Antoine Amarilli 

Télécom Paris, France
a3nm@a3nm.net

Alin Deutsch

University of California, San Diego, USA
abdeutsch@ucsd.edu

ACM Classification 2012

Information systems

ISBN 978-3-95977-320-1

Published online and open access by

Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH, Dagstuhl Publishing, Saarbrücken/Wadern, Germany. Online available at <https://www.dagstuhl.de/dagpub/978-3-95977-320-1>.

Publication date

June, 2024

Bibliographic information published by the Deutsche Nationalbibliothek

The Deutsche Nationalbibliothek lists this publication in the Deutsche Nationalbibliografie; detailed bibliographic data are available in the Internet at <https://portal.dnb.de>.

License

This work is licensed under a Creative Commons Attribution 4.0 International license (CC-BY 4.0):
<https://creativecommons.org/licenses/by/4.0/legalcode>.



In brief, this license authorizes each and everybody to share (to copy, distribute and transmit) the work under the following conditions, without impairing or restricting the authors' moral rights:

- Attribution: The work must be attributed to its authors.

The copyright is retained by the corresponding authors.

Digital Object Identifier: 10.4230/OASlcs.Tannen.2024.0

ISBN 978-3-95977-320-1

ISSN 1868-8969

<https://www.dagstuhl.de/oasics>

OASlcs – OpenAccess Series in Informatics

OASlcs is a series of high-quality conference proceedings across all fields in informatics. OASlcs volumes are published according to the principle of Open Access, i.e., they are available online and free of charge.

Editorial Board

- Daniel Cremers (TU München, Germany)
- Barbara Hammer (Universität Bielefeld, Germany)
- Marc Langheinrich (Università della Svizzera Italiana – Lugano, Switzerland)
- Dorothea Wagner (*Editor-in-Chief*, Karlsruher Institut für Technologie, Germany)

ISSN 1868-8969

<https://www.dagstuhl.de/oasics>

To Val Tannen

■ Contents

Preface	
<i>Antoine Amarilli and Alin Deutsch</i>	0:ix–0:x
List of Authors	
.....	0:xi
Regular Papers	
Explaining Enterprise Knowledge Graphs with Large Language Models and Ontological Reasoning	
<i>Teodoro Baldazzi, Luigi Bellomarini, Stefano Ceri, Andrea Colombo, Andrea Gentili, Emanuel Sallinger, and Paolo Atzeni</i>	1:1–1:20
Traversal-Invariant Characterizations of Logarithmic Space	
<i>Siddharth Bhaskar, Steven Lindell, and Scott Weinstein</i>	2:1–2:17
Semiring Provenance in the Infinite	
<i>Sophie Brinke, Erich Grädel, Louro Mrkonjić, and Matthias Naaf</i>	3:1–3:26
Annotation and More Annotation: Some Problems Posed by (and to) Val Tannen	
<i>Peter Buneman and Stijn Vansummeren</i>	4:1–4:8
Chasing Parallelism in Aggregating Graph Queries	
<i>Alin Deutsch</i>	5:1–5:14
Fishing Fort: A System for Graph Analytics with ML Prediction and Logic Deduction	
<i>Wenfei Fan and Shuhao Liu</i>	6:1–6:18
A Note on Logical PERs and Reducibility. Logical Relations Strike Again!	
<i>Jean Gallier</i>	7:1–7:12
AutoML for Explainable Anomaly Detection (XAD)	
<i>Nikolaos Myrtakis, Ioannis Tsamardinos, and Vassilis Christophides</i>	8:1–8:23
On the Impact of Provenance Semiring Theory on the Design of a Provenance-Aware Database System	
<i>Pierre Senellart</i>	9:1–9:10
Different Differences in Semirings	
<i>Dan Suciu</i>	10:1–10:20
An Intensional Expressiveness Gap of Comprehension Syntax	
<i>Limsoon Wong</i>	11:1–11:13

■ Preface

This Festschrift volume accompanies a colloquium held at the University of Pennsylvania on May 24–25, 2024 in celebration of the distinguished career of Val Tannen. Attendants gathered from all over the world to express their admiration for Val as a researcher as well as their love for him as a person. We are lucky to have him as a role model: a teacher, a mentor, a collaborator, a colleague, a friend (several categories often apply).

The articles presented here are scientific offerings from some of us to Val. They pertain to some of the many areas of his research interests. Our one regret is that, to preserve the element of surprise, we could not collaborate with Val on these papers, nor could we at least ask for his feedback. This is a pity, as Val is famously unfailingly insightful, always honest, and extremely generous with his time.

Val has contributed seminaly to the principles of both programming languages and databases and also to the cross-pollination and unification of the two areas. He also contributed to bioinformatics and to systematic and evolutionary biology. For lack of space, we cannot do justice here to his manifold contributions, and can only include a few highlights.

One of Val’s major contributions is the use of *structural recursion*, together with other ideas from functional programming and type theory, to inform the design of query languages for post-relational data. Besides providing the theoretical underpinning for query optimization over nested-relational, complex-valued and object-oriented data, his work yielded – through the use of comprehensions – a standard technique for embedding relational databases in programming languages. Modern database systems support user-defined aggregates using a template that is an instance of Val’s techniques.

Val was instrumental in unifying a series of classic database optimization techniques that had been previously developed independently and implemented in different phases of the optimizer, with only limited interaction. Examples include rewriting using materialized views, join minimization, semantic optimization, index- and hash-based query evaluation, all of which were unified by reduction to query minimization under constraints. This enabled a novel, chase-based optimization approach in which these techniques, as well as other techniques that had not been explicitly articulated, are implicitly considered simultaneously. This allows them to feed off each other synergistically to yield plans that standard phase-based optimization will necessarily miss even if given unbounded computational resources. Val’s work on chase-based optimization brought a purely abstract concept, the “chase”, introduced for theoretical studies of logical constraints, to the attention of developers of query optimizers.

A highly celebrated outcome of Val’s work is the invention of *provenance semirings*, which yield a generalization of many adjuncts to relational databases, such as probabilistic databases, C-tables, bag semantics, and even database security, enabling their unified treatment. In addition, provenance semirings provide a widely adopted general formalism for defining, capturing, storing, reasoning about, and optimizing data provenance. By now, the elegant concept of K-relations is widely known, well beyond the database research community, and has been applied in domains as diverse as operating systems, programming languages, and verification. Closer to home, Val’s work has inspired database researchers to use K-relations as a tool in analyzing the fine-grained complexity of query evaluation, or to extend relational query optimization techniques to tensor processing systems.

Val’s work features a common leitmotif: the surprising unification of seemingly disparate concepts and theories. Such unification is not achieved by devising complicated hybrid unions of these theories, but rather by distilling them down to their essence in sublimely



0:x Preface

elegant style, thus exposing their commonality. Unsurprisingly, some of his PhD students nicknamed Val the Great Unifier. We hereby dub him the Amazingly Insightful Great Unifier, a well-deserved title whose acronym is moreover a nod to the sharp wit that earned him a reputation as a delightful conversationalist. We are looking forward to many more conversations with him, both scientific and social in nature.


May 2024


Antoine Amarilli
Peter Buneman
Daniel Deutch
Alin Deutsch
Zack Ives
Dan Suciu


Also on behalf of the contributing authors listed below.

■ List of Authors

Paolo Atzeni  (1)
Università Roma Tre, Italy

Teodoro Baldazzi  (1)
Università Roma Tre, Italy

Luigi Bellomarini  (1)
Banca d'Italia, Roma, Italy


Siddharth Bhaskar  (2)
Department of Computer Science, James
Madison University, Harrisonburg, VA, USA

Sophie Brinke  (3)
RWTH Aachen University, Germany


Peter Buneman (4)
University of Edinburgh, UK

Stefano Ceri  (1)
Politecnico di Milano, Italy

Vassilis Christophides  (8)
ETIS Laboratory, CY Cergy Paris Université,
ENSEA, France


Andrea Colombo  (1)
Politecnico di Milano, Italy

Alin Deutsch (5)
University of California, San Diego, CA, USA


Wenfei Fan  (6)
Shenzhen Institute of Computing Sciences,
China; University of Edinburgh, UK;
Beihang University, Beijing, China

Jean Gallier (7)
University of Pennsylvania,
Philadelphia, PA, USA


Andrea Gentili  (1)
Banca d'Italia, Roma, Italy

Erich Grädel  (3)
RWTH Aachen University, Germany

Steven Lindell (2)
Department of Computer Science,
Haverford College, PA, USA

Shuhao Liu  (6)
Shenzhen Institute of Computing Sciences,
China


Lovro Mrkonjić  (3)
RWTH Aachen University, Germany


Nikolaos Myrtakis  (8)
Department of Computer Science, University of
Crete, Heraklion, Greece;
ETIS Laboratory, CY Cergy Paris Université,
ENSEA, France

Matthias Naaf  (3)
RWTH Aachen University, Germany


Emanuel Sallinger  (1)
TU Wien, Austria; University of Oxford, UK


Pierre Senellart  (9)
DI ENS, ENS, PSL University, CNRS, Paris,
France;
Inria, Paris, France;
Institut Universitaire de France, Paris, France;
CNRS@CREATE LTD, Singapore;
IPAL, CNRS, Singapore

Dan Suciu  (10)
University of Washington, Seattle, WA, USA

Ioannis Tsamardinos  (8)
Department of Computer Science,
University of Crete, Heraklion, Greece

Stijn Vansummeren  (4)
UHasselt, Data Science Institute, Belgium

Scott Weinstein  (2)
Department of Philosophy, University of
Pennsylvania, Philadelphia, PA, USA

Limsoon Wong  (11)
School of Computing, National University of
Singapore, Singapore

The Provenance of Elegance in Computation – Essays Dedicated to Val Tannen.
Editors: Antoine Amarilli and Alin Deutsch



OpenAccess Series in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

