

2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS 2017)

**Berkeley, California, USA
15-17 October 2017**

Pages 1-527



**IEEE Catalog Number: CFP17053-POD
ISBN: 978-1-5386-3465-3**

**Copyright © 2017 by the Institute of Electrical and Electronics Engineers, Inc.
All Rights Reserved**

Copyright and Reprint Permissions: Abstracting is permitted with credit to the source. Libraries are permitted to photocopy beyond the limit of U.S. copyright law for private use of patrons those articles in this volume that carry a code at the bottom of the first page, provided the per-copy fee indicated in the code is paid through Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923.

For other copying, reprint or republication permission, write to IEEE Copyrights Manager, IEEE Service Center, 445 Hoes Lane, Piscataway, NJ 08854. All rights reserved.

****** This is a print representation of what appears in the IEEE Digital Library. Some format issues inherent in the e-media version may also appear in this print version.***

IEEE Catalog Number:	CFP17053-POD
ISBN (Print-On-Demand):	978-1-5386-3465-3
ISBN (Online):	978-1-5386-3464-6
ISSN:	1523-8288

Additional Copies of This Publication Are Available From:

Curran Associates, Inc
57 Morehouse Lane
Red Hook, NY 12571 USA
Phone: (845) 758-0400
Fax: (845) 758-2633
E-mail: curran@proceedings.com
Web: www.proceedings.com

CURRAN ASSOCIATES INC.
proceedings
.com

58th Annual IEEE Symposium on Foundations of Computer Science

FOCS 2017

Table of Contents

Preface.....	xiii
Organizing Committee and Sponsors.....	xiv
Program Committee.....	xv
External Reviewers	xvi
Awards.....	xx

Session 1A

A Nearly Optimal Lower Bound on the Approximate Degree of AC^0	1
<i>Mark Bun and Justin Thaler</i>	
On the Quantitative Hardness of CVP	13
<i>Huck Bennett, Alexander Golovnev, and Noah Stephens-Davidowitz</i>	
Distributed PCP Theorems for Hardness of Approximation in P	25
<i>Amir Abboud, Aviad Rubinfeld, and Ryan Williams</i>	
Short Presburger Arithmetic Is Hard	37
<i>Danny Nguyen and Igor Pak</i>	

Session 1B

On the Local Structure of Stable Clustering Instances	49
<i>Vincent Cohen-Addad and Chris Schwiegelshohn</i>	
Better Guarantees for k-Means and Euclidean k-Median by Primal-Dual Algorithms	61
<i>Sara Ahmadian, Ashkan Norouzi-Fard, Ola Svensson, and Justin Ward</i>	
Statistical Query Lower Bounds for Robust Estimation of High-Dimensional Gaussians and Gaussian Mixtures	73
<i>Ilias Diakonikolas, Daniel M. Kane, and Alistair Stewart</i>	
On Learning Mixtures of Well-Separated Gaussians	85
<i>Oded Regev and Aravindan Vijayaraghavan</i>	

Session 2A

On Small-Depth Frege Proofs for Tseitin for Grids	97
<i>Johan Håstad</i>	
Random $\Theta(\log n)$ -CNFs Are Hard for Cutting Planes.....	109
<i>Noah Fleming, Denis Pankratov, Toniann Pitassi, and Robert Robere</i>	
Random Formulas, Monotone Circuits, and Interpolation	121
<i>Pavel Hrubeš and Pavel Pudlák</i>	
Query-to-Communication Lifting for BPP	132
<i>Mika Göös, Toniann Pitassi, and Thomas Watson</i>	
A Rounds vs. Communication Tradeoff for Multi-Party Set Disjointness	144
<i>Mark Braverman and Rotem Oshman</i>	

Session 2B

A Time Hierarchy Theorem for the LOCAL Model	156
<i>Yi-Jun Chang and Seth Pettie</i>	
Distributed Exact Weighted All-Pairs Shortest Paths in $\tilde{O}(n^{\{5/4\}})$ Rounds	168
<i>Chien-Chung Huang, Danupon Nanongkai, and Thatchaphol Saranurak</i>	
Deterministic Distributed Edge-Coloring via Hypergraph Maximal Matching	180
<i>Manuela Fischer, Mohsen Ghaffari, and Fabian Kuhn</i>	
Fine-Grained Complexity of Analyzing Compressed Data: Quantifying Improvements over Decompress-and-Solve	192
<i>Amir Abboud, Arturs Backurs, Karl Bringmann, and Marvin Künnemann</i>	

Session 3A

Local List Recovery of High-Rate Tensor Codes & Applications	204
<i>Brett Hemenway, Noga Ron-Zewi, and Mary Wootters</i>	
Optimal Repair of Reed-Solomon Codes: Achieving the Cut-Set Bound	216
<i>Itzhak Tamo, Min Ye, and Alexander Barg</i>	
Average-Case Reconstruction for the Deletion Channel: Subpolynomially Many Traces Suffice	228
<i>Yuval Peres and Alex Zhai</i>	
Optimal Interactive Coding for Insertions, Deletions, and Substitutions	240
<i>Alexander A. Sherstov and Pei Wu</i>	
The Independence Number of the Birkhoff Polytope Graph, and Applications to Maximally Recoverable Codes	252
<i>Daniel Kane, Shachar Lovett, and Sankeerth Rao</i>	

Session 3B

Approximating Geometric Knapsack via L-Packings	260
<i>Waldo Gálvez, Fabrizio Grandoni, Sandy Heydrich, Salvatore Ingala, Arindam Khan, and Andreas Wiese</i>	
Removing Depth-Order Cycles among Triangles: An Efficient Algorithm Generating Triangular Fragments	272
<i>Mark de Berg</i>	
Scheduling to Minimize Total Weighted Completion Time via Time-Indexed Linear Programming Relaxations	283
<i>Shi Li</i>	
Fast & Space-Efficient Approximations of Language Edit Distance and RNA Folding: An Amnesic Dynamic Programming Approach	295
<i>Barna Saha</i>	
A Dichotomy for Regular Expression Membership Testing	307
<i>Karl Bringmann, Allan Grønlund, and Kasper Green Larsen</i>	

Session 4

A Dichotomy Theorem for Nonuniform CSPs	319
<i>Andrei A. Bulatov</i>	
A Proof of CSP Dichotomy Conjecture	331
<i>Dmitriy Zhuk</i>	

Session 5A

Learning Graphical Models Using Multiplicative Weights	343
<i>Adam Klivans and Raghu Meka</i>	
Active Classification with Comparison Queries	355
<i>Daniel M. Kane, Shachar Lovett, Shay Moran, and Jiapeng Zhang</i>	
Capacity of Neural Networks for Lifelong Learning of Composable Tasks	367
<i>Leslie G. Valiant</i>	
Efficient Bayesian Estimation from Few Samples: Community Detection and Related Problems	379
<i>Samuel B. Hopkins and David Steurer</i>	
Robust Polynomial Regression up to the Information Theoretic Limit	391
<i>Daniel Kane, Sushrut Karmalkar, and Eric Price</i>	

Session 5B

Quantum SDP-Solvers: Better Upper and Lower Bounds	403
<i>Joran Van Apeldoorn, András Gilyén, Sander Gribling, and Ronald de Wolf</i>	
Quantum Speed-Ups for Solving Semidefinite Programs	415
<i>Fernando G.S.L. Brandao and Krysta M. Svore</i>	
Local Hamiltonians Whose Ground States Are Hard to Approximate	427
<i>Lior Eldar and Aram W. Harrow</i>	
On Preparing Ground States of Gapped Hamiltonians: An Efficient Quantum Lovász Local Lemma	439
<i>András Pál Gilyén and Or Sattath</i>	
Variable-Version Lovász Local Lemma: Beyond Shearer's Bound	451
<i>Kun He, Liang Li, Xingwu Liu, Yuyi Wang, and Mingji Xia</i>	
Linear Algebraic Analogues of the Graph Isomorphism Problem and the Erdős-Rényi Model	463
<i>Yinan Li and Youming Qiao</i>	

Session 6A

Optimal Lower Bounds for Universal Relation, and for Samplers and Finding Duplicates in Streams	475
<i>Michael Kapralov, Jelani Nelson, Jakub Pachocki, Zhengyu Wang, David P. Woodruff, and Mobin Yahyazadeh</i>	
First Efficient Convergence for Streaming k-PCA: A Global, Gap-Free, and Near-Optimal Rate	487
<i>Zeyuan Allen-Zhu and Yuanzhi Li</i>	
Weighted k-Server Bounds via Combinatorial Dichotomies	493
<i>Nikhil Bansal, Marek Eliáš, and Grigorios Koumoutsos</i>	
An Input Sensitive Online Algorithm for the Metric Bipartite Matching Problem	505
<i>Krati Nayyar and Sharath Raghvendra</i>	

Session 6B

Learning Multi-Item Auctions with (or without) Samples	516
<i>Yang Cai and Constantinos Daskalakis</i>	
Oracle-Efficient Online Learning and Auction Design	528
<i>Miroslav Dudík, Nika Haghtalab, Haipeng Luo, Robert E. Schapire, Vasilis Syrgkanis, and Jennifer Wortman Vaughan</i>	
Prophet Inequalities Made Easy: Stochastic Optimization by Pricing Non-Stochastic Inputs	540
<i>Paul Dütting, Michal Feldman, Thomas Kesselheim, and Brendan Lucier</i>	

Tight Lower Bounds for Differentially Private Selection	552
<i>Thomas Steinke and Jonathan Ullman</i>	

Session 7A

How to Achieve Non-Malleability in One or Two Rounds	564
<i>Dakshita Khurana and Amit Sahai</i>	
Two-Round and Non-Interactive Concurrent Non-Malleable Commitments from Time-Lock Puzzles	576
<i>Huijia Lin, Rafael Pass, and Pratik Soni</i>	
Garbled Protocols and Two-Round MPC from Bilinear Maps	588
<i>Sanjam Garg and Akshayaram Srinivasan</i>	
Obfuscating Compute-and-Compare Programs under LWE	600
<i>Daniel Wichs and Giorgos Zirdelis</i>	
Lockable Obfuscation	612
<i>Rishab Goyal, Venkata Koppula, and Brent Waters</i>	
White-Box vs. Black-Box Complexity of Search Problems: Ramsey and Graph Property Testing	622
<i>Ilan Komargodski, Moni Naor, and Eylon Yogev</i>	

Session 7B

Optimality of the Johnson-Lindenstrauss Lemma	633
<i>Kasper Green Larsen and Jelani Nelson</i>	
Optimal Compression of Approximate Inner Products and Dimension Reduction	639
<i>Noga Alon and Bo'az Klartag</i>	
Sample Efficient Estimation and Recovery in Sparse FFT via Isolation on Average	651
<i>Michael Kapralov</i>	
Fast Similarity Sketching	663
<i>Søren Dahlgaard, Mathias Bæk Tejs Knudsen, and Mikkel Thorup</i>	
Sublinear Time Low-Rank Approximation of Positive Semidefinite Matrices	672
<i>Cameron Musco and David P. Woodruff</i>	

Session 8

Hardness Results for Structured Linear Systems	684
<i>Rasmus Kyng and Peng Zhang</i>	
The Matching Problem in General Graphs Is in Quasi-NC	696
<i>Ola Svensson and Jakub Tarnawski</i>	

Session 9A

On the Power of Statistical Zero Knowledge	708
<i>Adam Boulund, Lijie Chen, Dhiraj Holden, Justin Thaler, and Prashant Nalini Vasudevan</i>	
The Power of Sum-of-Squares for Detecting Hidden Structures	720
<i>Samuel B. Hopkins, Pravesh K. Kothari, Aaron Potechin, Prasad Raghavendra, Tselil Schramm, and David Steurer</i>	
A Time-Space Lower Bound for a Large Class of Learning Problems	732
<i>Ran Raz</i>	
From Gap-ETH to FPT-Inapproximability: Clique, Dominating Set, and More	743
<i>Parinya Chalermsook, Marek Cygan, Guy Kortsarz, Bundit Laekhanukit, Pasin Manurangsi, Danupon Nanongkai, and Luca Trevisan</i>	

Session 9B

Faster (and Still Pretty Simple) Unbiased Estimators for Network (Un)reliability	755
<i>David R. Karger</i>	
Minor-Free Graphs Have Light Spanners	767
<i>Glencora Borradaile, Hung Le, and Christian Wulff-Nilsen</i>	
Polylogarithmic Approximation for Minimum Planarization (Almost)	779
<i>Ken-Ichi Kawarabayashi and Anastasios Sidiropoulos</i>	
Approximating the Held-Karp Bound for Metric TSP in Nearly-Linear Time	789
<i>Chandra Chekuri and Kent Quanrud</i>	

Session 10A

Derandomization Beyond Connectivity: Undirected Laplacian Systems in Nearly Logarithmic Space	801
<i>Jack Murtagh, Omer Reingold, Aaron Sidford, and Salil Vadhan</i>	
Deterministic Search for CNF Satisfying Assignments in Almost Polynomial Time	813
<i>Rocco A. Servedio and Li-Yang Tan</i>	
Fooling Intersections of Low-Weight Halfspaces	824
<i>Rocco A. Servedio and Li-Yang Tan</i>	
Exponentially-Hard Gap-CSP and Local PRG via Local Hardcore Functions	836
<i>Benny Applebaum</i>	

Session 10B

Testing Hereditary Properties of Ordered Graphs and Matrices	848
<i>Noga Alon, Omri Ben-Eliezer, and Eldar Fischer</i>	
A Characterization of Testable Hypergraph Properties	859
<i>Felix Joos, Jaehoon Kim, Daniela Kühn, and Deryk Osthus</i>	
Boolean Unateness Testing with $\tilde{O}(n^{3/4})$ Adaptive Queries	868
<i>Xi Chen, Erik Waingarten, and Jinyu Xie</i>	
Generalized Uniformity Testing	880
<i>Tugkan Batu and Clément L. Canonne</i>	

Session 11A

Much Faster Algorithms for Matrix Scaling	890
<i>Zeyuan Allen-Zhu, Yuanzhi Li, Rafael Oliveira, and Avi Wigderson</i>	
Matrix Scaling and Balancing via Box Constrained Newton's Method and Interior Point Methods	902
<i>Michael B. Cohen, Aleksander Madry, Dimitris Tsipras, and Adrian Vladu</i>	
Simply Exponential Approximation of the Permanent of Positive Semidefinite Matrices	914
<i>Nima Anari, Leonid Gurvits, Shayan Oveis Gharan, and Amin Saberi</i>	
Determinant-Preserving Sparsification of SDDM Matrices with Applications to Counting and Sampling Spanning Trees	926
<i>David Durfee, John Peebles, Richard Peng, and Anup B. Rao</i>	

Session 11B

Optimal Las Vegas Locality Sensitive Data Structures	938
<i>Thomas Dybdahl Ahle</i>	
Dynamic Minimum Spanning Forest with Subpolynomial Worst-Case Update Time	950
<i>Danupon Nanongkai, Thatchaphol Saranurak, and Christian Wulff-Nilsen</i>	
Fast and Compact Exact Distance Oracle for Planar Graphs	962
<i>Vincent Cohen-Addad, Søren Dahlgaard, and Christian Wulff-Nilsen</i>	

Session 12A

High Dimensional Expanders Imply Agreement Expanders	974
<i>Irit Dinur and Tali Kaufman</i>	
The Ising Partition Function: Zeros and Deterministic Approximation	986
<i>Jingcheng Liu, Alistair Sinclair, and Piyush Srivastava</i>	

Eldan's Stochastic Localization and the KLS Hyperplane Conjecture: An Improved Lower Bound for Expansion	998
<i>Yin Tat Lee and Santosh Srinivas Vempala</i>	

Session 12B

Weak Decoupling, Polynomial Folds and Approximate Optimization over the Sphere	1008
<i>Vijay Bhattiprolu, Mrinalkanti Ghosh, Venkatesan Guruswami, Euiwoong Lee, and Madhur Tulsiani</i>	
Subdeterminant Maximization via Nonconvex Relaxations and Anti-Concentration	1020
<i>Javad B. Ebrahimi, Damian Straszak, and Nisheeth K. Vishnoi</i>	
Hashing-Based-Estimators for Kernel Density in High Dimensions	1032
<i>Moses Charikar and Paris Siminelakis</i>	

Author Index