THEORY OF COMPUTING, Volume 8 (2012), pp. 55–68
www.theoryofcomputing.org

SPECIAL ISSUE IN HONOR OF RAJEEV MOTWANI

Rajeev Motwani (1962-2009)

Prabhakar Raghavan*

Received: September 13, 2010; published: March 31, 2012.

Abstract: Rajeev Motwani was a pre-eminent theoretical computer scientist of his generation, a technology thought leader, an insightful venture capitalist, and a mentor to some of the most influential entrepreneurs in Silicon Valley in the first decade of the 21st century. This article presents an overview of Rajeev's research, and provides a window to his early life and the various influences that shaped his research and professional career—it is a small celebration of his wonderful life and many achievements.

ACM Classification: K.2/People, F.2.2

AMS Classification: 01A70, 68-00

Key words and phrases: biography

1 Early Life and Education

Rajeev Motwani was born on March 24, 1962 in the Indian city of Jammu to Lieutenant Colonel Hotchand Motwani, an officer in the Indian Army, and Namita Motwani (née Sushila). His family included brothers Sanjeev and Suneev. Given his father's army career, Rajeev's family moved often and lived in various parts of India before settling down in New Delhi.

Those who knew Rajeev as a young boy remember him as an avid reader. When Rajeev was seven, his father was stationed in the scenic town of Devlali near Mumbai, India. His family would walk a kilometer to the local library to get books, and the seven-year-old Rajeev would be seen reading the books

*Based on contributions by Sanjeev Arora, Gautam Bhargava, Piotr Indyk, Asha Jadeja, Sanjeev Motwani, Prabhakar Raghavan, and G. D. Ramkumar.

DOI: 10.4086/toc.2012.v008a003

they had borrowed from the library as they walked home! His brothers recall that not a single day would pass when he did not read a whole book. Rajeev would read books of all types, including novels, comics, autobiographies and scientific books.

Rajeev's father had a huge influence on him. He indulged Rajeev's interest in books and encouraged him to read and pursue his interests to the fullest. In an interview in 2002, Rajeev recalled his debt to his parents for their belief in education and for sending him to great schools with inspiring teachers. Rajeev's love of learning and the breadth of his interests foreshadowed the depth and range of his future research.

A turning point in Rajeev's intellectual development came when his family moved to New Delhi in 1974. Rajeev's father wanted to send his children to the prestigious St Columba's High School in New Delhi. St Columba's High School administered a difficult entrance exam to admit students, and Rajeev studied hard the night before taking the exam. Not only did Rajeev pass the exam, he did so well that the principal admitted all three brothers into the school! This was perhaps the first time that Rajeev demonstrated his prowess in mathematics.

Encouraged by his parents, Rajeev soon developed a keen interest in learning about famous mathematicians. One of his childhood heroes was the German mathematician Carl Friedrich Gauss. Rajeev also loved music, particularly rock music.

When Rajeev finished high school in 1978, he applied to the Indian Institutes of Technology (IITs). He did extremely well in the highly competitive entrance examination, earning third place in the northern zone of India. Although Rajeev wanted to pursue his passion in mathematics, he was encouraged by his family and friends to study computer science, as it was considered a valuable skill to acquire at the time. He reluctantly agreed, not realizing until much later the close connections shared by the two subjects.

Even though IIT Kanpur had an outstanding computer science faculty in the late 1970s, formal computer science education at the undergraduate level was still in its infancy in India. Rajeev was a member of the very first cohort of undergraduate computer science students at IIT Kanpur. Rajeev often recalled that IIT Kanpur had attracted an amazing group of people, and there could not have been a better environment for studying computer science in India. As a student, Rajeev was inspired by Professor Kesav Nori, who taught Rajeev's first class on programming. Prof. Nori recalls that Rajeev had a knack for creating the most elegant and brief answers to the hardest of programming problems. Rajeev's undergraduate thesis (joint with Chilukuri K. Mohan and Amitabh Shah; advised by Professor Somenath Biswas) was quite theoretical: "Specification and Verification of Computer Communication Protocols." This simultaneous interest in theory and programming percolated through Rajeev's career.

Rajeev's friends recall that in college he was not only brilliant, but also incredibly fun-loving and always ready for a party. There would always be a stack of science-fiction books waiting for eager consumption in his dorm room, and he would spend endless hours solving difficult crossword puzzles, playing bridge or volleyball, and hanging out with friends. He never lost this undergraduate spirit. Through his time at Berkeley and as a faculty member at Stanford, Rajeev approached research and entrepreneurship with joy.

After graduating from IIT Kanpur in 1982, Rajeev joined the Ph. D. program in Computer Science at UC Berkeley, where Professor Richard Karp was his advisor.

2 Professional Career

As he approached the completion of his Ph. D. studies, Rajeev contemplated options that included returning to India and staying in the US. His ultimate decision was influenced by a chance meeting with Donald Knuth. Knuth was visiting Berkeley, and mentioned to Prof. Karp that Stanford wanted to hire a young faculty member in the area of algorithms. Prof. Karp suggested Rajeev's name, and Knuth asked Rajeev to join Stanford for a year and see if things worked out well. Rajeev was on the visiting faculty at Stanford for a year, after which he was given a permanent offer. Stanford University President John Hennessey later recalled how Knuth had enthusiastically presented the case for appointing Rajeev. President Hennessey said they all underestimated what a superb scholar, teacher, advisor, colleague, and friend Rajeev would become.

Rajeev met his soon-to-be wife Asha Jadeja in May 1989, around the time he started teaching at Stanford. Asha graduated from the University of Southern California and relocated to the San Francisco Bay area. Rajeev and Asha were married on March 22, 1990 in Delhi. They liked living in Stanford, and decided to stay. Asha started graduate school at UC Berkeley in 1991, studying urban planning and urban transportation, and joined the department of political science at Stanford in 1994. Two daughters were born from their marriage, Naitri (19) and Anya (6).

Rajeev was fascinated by the transformation of academic ideas into commercial ventures. To quote his friend and coauthor, Madhu Sudan from MIT, "Having worked with Rajeev at an early stage of my career I knew how smart he was, but it really took Stanford to bring out his true talent, which was the ability to recognize the importance of ideas quicker than anyone else around."

Rajeev was active in the venture industry, and had a reputation of being extremely helpful to entrepreneurs. Rajeev's students and industry colleagues remember him for his uncanny ability to connect people: he was a catalyst in bringing teams together and getting companies started with an element of inspiration and strategic guidance. Through Deutsche Bank, to which he had been an advisor, he became one of the first investors in PayPal, and Rajeev and Asha started a venture fund called Dot Edu Ventures in 2000. He was also a special advisor to Sequoia Capital, and was active in entrepreneurial student groups at Stanford, including BASES and Stanford Student Enterprises.

Rajeev was an inspiring teacher whose courses "Automata and Complexity Theory," "Randomized Algorithms," and "Advanced Data Structures and Algorithms" were enjoyed by generations of Stanford's students. His teaching style was understated and precise, yet alive with the beauty of the material, and he had a distinctive blackboard technique and handwriting style that many students who went on to become academics tried to emulate. His students remember the clarity of his lectures and his ability to explain complex technical material with patience and intuition.

Rajeev had the reputation among students of being a great advisor. Many of his students credit Rajeev's patient coaching and ability to be what the student needed, interacting with each student differently and building on an individual student's strengths and interests, as the reason each of them was able to find their way and go on to a successful career. Rajeev became the Director of Graduate Studies in the Computer Science Department at Stanford, and was one of the first to greet incoming Ph. D. students with his welcoming smile. Rajeev remained a kind, perceptive and always accessible mentor to many students well after their graduation, and into their academic and industry careers.

Rajeev was awarded the Distinguished Alumnus Award by IIT Kanpur in 2006. Quoting from IIT

Kanpur's web site, those in the institute who knew Rajeev as an undergraduate student remember him not only for his academic brilliance but also for the good cheer he always exuded. He was a very friendly person, and remained so throughout his life. His office door was ever open-to students, to young entrepreneurs, to academicians from across the world, and to his friends from IIT Kanpur. Rajeev was also closely involved with the Department of Computer Science and Engineering at IIT Kanpur as a mentor in its research program.

In the early days of the World Wide Web, Rajeev teamed up with other professors and future Google co-founder Sergey Brin, a Stanford graduate student at the time, to form a group called Mining Data at Stanford (MIDAS) that explored data mining on the web. Larry Page, the other future Google co-founder, also a Stanford graduate student at the time, had come up with an interesting idea for doing random walks on the web. Sergey soon became involved in the project, and Rajeev was the perfect sounding board for Sergey and Larry. They created a search engine called Backrub, running on Stanford servers until the traffic became too high. The search engine was re-located off campus, and its name changed to Google. Even after Google moved off campus, Rajeev spent considerable time working with the company. Brin acknowledged in a BBC Radio interview in 2009 that Google might not have been possible without Rajeev.

Rajeev Motwani died tragically in a swimming accident at his home in Atherton, California on June 5, 2009. He is sorely missed by his family, friends and colleagues.

3 Research

Rajeev Motwani's great interest in fundamental problems of theoretical computer science *and* a keen interest and awareness of applications sets him apart as a rarity among theoretical computer scientists. He repeatedly exhibited a grasp of issues affecting many areas of computing, and actively sought out and influenced new applications of his theoretical ideas.

3.1 Algorithmic combinatorics and graph theory

Motwani's early work, done during his Ph. D. studies and shortly thereafter, was focused on algorithms with a heavy combinatorial flavor. One of his first papers [55], written in 1988, was on constructive results related to the graph minor theorem, a topic he was fond of throughout his career. He followed it up with a paper the next year on the average-case analysis of matching algorithms. This paper showed that (a) there is a large gap between the average-case and worst-case performances of matching algorithms and (b) this gap can be explained by the expansion properties of random graphs. This result foreshadowed a major theme in Motwani's future work: the use of combinatorial and algorithmic tools to address issues arising in the practice of computing. The paper became the central component of Motwani's Ph. D. thesis.

A few months later, in collaboration with Joseph Naor and Moni Naor, Motwani turned to another algorithmic problem of combinatorial flavor. The resulting paper [51] showed how to find various combinatorial structures (e. g., low-discrepancy colorings) efficiently using parallel algorithms. The existence of such structures is a classic application of the probabilistic method pioneered by Erdős: the proof of the existence of a structure with a certain property proceeds by showing that a *randomly chosen* structure satisfies the property with positive probability. Unfortunately, this approach does not provide a method for computing the object in a *certifiable* and/or *deterministic* manner. To achieve the latter, one needs to set the random bits deterministically (i. e., to derandomize the algorithm). A popular method for achieving this goal is the method of conditional probabilities, which fixes the bits one by one, while maintaining a high (conditional) probability of success. At the time, this procedure seemed to be inherently sequential in nature. What Motwani, Naor and Naor showed, however, is that in fact one can simulate this task in only a polylogarithmic number of steps. (The same result was independently obtained by Berger and Rompel [10].) The solution proceeds in two phases. First, one generates a pseudorandom sequence of bits from a seed of only polylogarithmic length. The second (and the key) step is to show that the relevant conditional probabilities can still be computed in this setting. This approach led to the first NC algorithms for a variety of problems, including low discrepancy colorings, Ramsey graphs and hypergraph covers.

Another of Motwani's papers in this area [24], co-authored with Tomás Feder, introduced a general speed-up technique for graph algorithms, based on decomposing the input graph into small cliques. Algorithmic graph theory remained one of Motwani's pet topics over the years, leading to algorithms for finding longest paths [41, 26, 25], shortest paths [3], graph partitioning, and other tasks. Many of those papers were written with Feder, including one of Motwani's last papers [25].

3.2 Probabilistically Checkable Proofs

The PCP theorem, discovered by Motwani in collaboration with Sanjeev Arora, Carsten Lund, Madhu Sudan, and Mario Szegedy [4], is widely considered to be one of the deepest results in theoretical computer science. In a nutshell, the PCP theorem implies that computing even approximate solutions is hard for many NP-complete problems (assuming the widely-believed conjecture that these problems are hard to solve exactly). The authors of the theorem were awarded the prestigious Gödel Prize for their accomplishment in 2001.

The PCP theorem is rooted in exciting complexity-theoretic breakthroughs published in 1990 that gave probabilistic definitions of traditional complexity classes including PSPACE and NEXPTIME [48, 59, 7]. Work of Babai et al. [6], Feige et al. [27], and Arora and Safra [5] then showed that these new ideas could be extended to the class NP, and the resulting probabilistic definition of NP could be used to show that computing approximate solutions to the maximum clique problem is NP-hard.

The paper of Motwani and his coauthors that proved the PCP theorem represented the culmination of this line of research by showing two results. First, the definition of NP in [5] was enhanced to an "optimal" one. That is, it was shown that NP is *exactly* the set of languages for which membership proofs can be checked in polynomial time by a probabilistic verifier that uses only a logarithmic number of random bits and examines a fixed constant number of bits of the membership proof. Second, this new definition of NP was shown to imply that computing approximate solutions to a vast class of problems-called MAX-SNP, and identified in earlier work of [58]—is also NP-hard. The paper is highly cited, and has formed the cornerstone of all subsequent investigations into probabilistically checkable proofs. The area has been very active in the past two decades, leading to results showing the NP-hardness of computing approximate solutions to a host of optimization problems. Motwani's subsequent work on this topic includes [45] and [44].

3.3 Approximation algorithms

The discovery of the PCP theorem and its applications to proving hardness of approximation results provided new inspiration for investigating the approximability of NP-hard optimization problems. Further stimulus was provided by the 1994 paper of Goemans and Williamson [32], which showed that the classic approximation algorithm for the maximum cut problem could be drastically improved by relaxing the problem as a *semidefinite program (SDP)* and rounding the relaxed solution using the *random hyperplane* method.

In the next big step in this direction, Karger, Motwani and Sudan [42] presented an improved approximation algorithm for the *graph-coloring* problem (for the well-studied case of 3-colorable graphs). As with Goemans and Williamson's maximum cut algorithm, this algorithm started by computing an SDP relaxation of the problem and then rounding it to an integral solution. However, random hyperplanes do not work for graph coloring-for they result in an approximation factor slightly worse than the best-known prior algorithm. Instead, Karger, Motwani and Sudan introduced a more elaborate approach, based on rounding to the nearest random vector. This led to a significantly better approximation factor that became "low" for "low-degree" graphs. In order to achieve an improvement for general graphs, one needed an additional idea to take care of the high-degree nodes: this was achieved by resorting to an algorithm of Wigderson [60]. The final outcome became one of the classic results in the area of approximation algorithms, influencing a great body of later work.

As was typical of Motwani's research, his work on graph coloring was not limited to theoretical developments. For example, Hasan and Motwani [36] showed how to use (a different variant of) graph coloring to facilitate parallel query optimization in databases. In general, approximation algorithms became a recurrent topic in Motwani's work, with many of his papers investigating approximation and on-line algorithms for combinatorial problems including scheduling [53, 18, 21, 19, 20] and routing [1, 16].

3.4 Randomized geometric algorithms

As discussed above, Motwani's interest in geometric computation dates to the beginning of his career. His interests in this area would later shift toward robotics, with his first paper on this topic [35] originating from the following question, posed by his colleague, Leo Guibas: Suppose a robot has a map of the building it is in, and can determine (from a range-finder scan) the distances to the walls immediately visible from its position. How quickly can the robot determine its location on the map? Motwani and his collaborators discovered a data structure for processing such *robot localization* queries, based on partitioning the interior of the map into regions, within each of which the visibility of the robot is roughly the same.

Further work by Motwani and his collaborators in this area followed [43, 37, 33]. In particular, they presented a formalization and analysis of the highly successful *probabilistic roadmap algorithm* for robot path planning [43, 37]. Motwani also applied the tools of robotics and geometric pattern matching to drug design [28, 29, 40, 30].

3.5 Data mining and information retrieval

In the late 1990s, motivated by the emergence of massive datasets and the World Wide Web, Motwani became interested in designing efficient algorithms for data mining. His first work in this direction aimed at finding general and efficient frameworks for mining so-called *association* rules, popularized by Rakesh Agarwal and his colleagues in the then nascent field of data mining. Together with his colleagues in the Stanford database group, Motwani wrote several influential papers [12, 11, 23] proposing new frameworks and algorithms for finding association rules and their variants.

Arguably the most influential paper by Motwani, co-authored with Sergey Brin, Larry Page and Terry Winograd in 1998 [57], was never published in a conference or journal. This paper cast the PageRank concept of Brin and Page into the setting of Markov chain analysis, and formed the early basis of the computation of the search ranking function used by the Google search engine. At around this time, Motwani (being close to the workings of Google) wrote a number of well-cited papers spelling out the open algorithmic problems in web-scale search engines.

3.6 Similarity search

Motwani's research in massive data sets sparked his interest in the design of algorithms and data structures for nearest neighbor search. Initial work in the area was done by Indyk, Motwani, Raghavan and Vempala [39], whose paper proposed the notion of *locality-preserving hashing*: a higher-dimensional (but somewhat weaker) variant of the *non-expansive hashing* of Linial and Sasson [47]. However, it did not address the main challenge: how to design similarity search algorithms that scale *polynomially* with the dimension.

Approximation algorithms with the desired property were subsequently developed by Indyk and Motwani [38]. Their paper contained three main results: a reduction from approximate nearest neighbor to approximate near neighbor (a. k. a. point location in equal balls), and two solutions to the approximate near neighbor problem. One of the solutions was based on the concept of *locality sensitive hashing (LSH)*: a family of mappings with the property that close points have a higher probability of collisions than those that are far apart. In a follow-up paper, Gionis, Indyk, and Motwani [31] showed that the latter algorithm offered a significant speed-up over many other algorithms available in the literature.

Almost a decade later, in collaboration with Assaf Naor and Rina Panigrahy [50], Motwani returned to this question, proving a lower bound on the best running time exponent achievable using LSH method.

3.7 Streaming algorithms

Motwani's foray into databases led to a rich subsection of his technical output, including numerous papers on sampling large data sets and *streaming algorithms*. These algorithms are restricted to making only a small number of passes over the data, and using very little storage to perform computations. Motwani's first paper on this topic [15] focused on the following k-center problem: partition the input sequence of points into k clusters so that the largest diameter of the clusters is minimized. The paper presented a constant approximation algorithm for this problem using only one pass over the data and $O(k)$ storage. In fact, the algorithm provided a stronger guarantee: as the algorithm acquired more data, the old points were not re-grouped arbitrarily; instead, new clusters were only formed by merging old clusters.

A few years later, Guha, Mishra, Motwani and Raghavan [34] presented a streaming algorithm for the (much harder) k-median problem. Here, the goal is to minimize the sum of the distances from each point to its nearest cluster center. The key idea behind the algorithm was a composition theorem, which (roughly speaking) stated that if we split the data into blocks, cluster each block, and then cluster the clusters, we obtain a "good" clustering of the whole data set. This divide-and-conquer approach to data stream algorithms influenced many future developments, both in theory and in practice.

More streaming results by Motwani followed, including an algorithm for finding heavy hitters [49] and algorithms for the sliding window model [22, 8, 9]. Motwani was also very interested in sampling algorithms, especially in the context of databases [17, 14, 52, 13]. Ultimately, his interests in the area turned to a more applied direction: together with Jennifer Widom, he led the *Stanford Stream Data Manager Project* [56].

3.8 Anonymity and Privacy

The development of data mining tools raised concerns about preserving privacy of personal data. In some of his later work [2, 46], Motwani investigated frameworks and algorithms for data sanitization, query auditing and related problems.

3.9 Books

Soon after his arrival at Stanford, Motwani began teaching advanced graduate courses on randomized algorithms, advocating a school of thought popularized by his advisor Richard Karp. Around the same time, his friend and colleague Prabhakar Raghavan was teaching and developing course notes for a similar course at Yale University. In 1991 they began collaborating, turning these notes into a text book published by Cambridge University Press in 1995. The book benefited from the input of many colleagues who taught from or proofread early versions, particularly Allan Borodin at the University of Toronto and Donald Knuth at Stanford. To this day, few have recognized the somewhat esoteric allusion to randomized algorithms on the cover of the book!

In the late 1990s, John Hopcroft and Jeff Ullman were revising their classic text "Introduction to Automata Theory, Languages, and Computation." Motwani had taught extensively from the first edition of this book, and his mastery of the subject and detailed comments led the authors to invite Motwani to join them as a co-author of the second edition, published in 2001.

4 Back to the Beginning

In the spring of 1983, Rajeev was a junior graduate student at UC Berkeley enrolled in an advanced course on algorithms taught by Prof. Richard Karp. Prof. Karp presented the following open problem as a homework exercise: given a set S of n elements from an ordered set, there are two strategies for searching whether a query q is present in S. One could simply compare q with each of the n elements of S, incurring a cost of n comparisons. Or one could first sort S incurring an up-front cost of $O(n \log n)$ to build a binary search tree on S, following which we could process as many queries q as needed, paying only $\log n$ comparisons for each query. The former strategy is best suited to a small number of queries q, while the latter works better when the number of queries is known to be large enough to justify the search tree construction. The question posed by Prof. Karp was: what if the number r of queries is not known in advance?

Working with a fellow graduate student and Prof. Karp, Rajeev developed this question into a theory of *deferred data structures*: data structures that are built up incrementally, depending on the number of queries processed, in a way that no matter when the query stream were to end the total cost incurred up to that point would be optimal. In the simple search setting above, Rajeev and his collaborators showed that the optimal number of comparisons was $n \log r + r \log n$, and presented a deferred data structure that achieved this bound. The main idea was to gradually build up the search tree on S, going deeper as the number r became larger. They extended this idea, proving optimality for a number of other problems as well. This was Rajeev's first published paper [54], and he summarized it jokingly as his philosophy in life, with the words: *Never do work today that you can defer to tomorrow*.

References

[1] ALOK AGGARWAL, DON COPPERSMITH, SANJEEV KHANNA, RAJEEV MOTWANI, AND BARUCH SCHIEBER: The angular-metric traveling salesman problem. *SIAM J. Comput.*, 29(3):697–711, 1999. [doi:10.1137/S0097539796312721] 60

[2] GAGAN AGGARWAL, TOMÁS FEDER, KRISHNARAM KENTHAPADI, RAJEEV MOTWANI, RINA PANIGRAHY, DILYS THOMAS, AND AN ZHU: Anonymizing tables. In *Proc. 10th Internat. Conf. on Database Theory (ICDT'05)*, pp. 246–258, 2005. [doi:10.1007/978-3-540-30570-5_17] 62

[3] DONALD AINGWORTH, CHANDRA CHEKURI, AND RAJEEV MOTWANI: Fast estimation of diameter and shortest paths (without matrix multiplication). In *Proc. 7th Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA'96)*, pp. 547–553, 1996. 59

[4] SANJEEV ARORA, CARSTEN LUND, RAJEEV MOTWANI, MADHU SUDAN, AND MARIO SZEGEDY: Proof verification and the hardness of approximation problems. *J. ACM*, 45(3):501–555, 1998. Preliminary version in FOCS'92. [doi:10.1145/278298.278306] 59

[5] SANJEEV ARORA AND SHMUEL SAFRA: Probabilistic checking of proofs: A new characterization of NP. *J. ACM*, 45(1):70–122, 1998. Preliminary version in FOCS'92. [doi:10.1145/273865.273901] 59

[6] LÁSZLÓ BABAI, LANCE FORTNOW, LEONID A. LEVIN, AND MARIO SZEGEDY: Checking computations in polylogarithmic time. In *Proc. 23rd STOC*, pp. 21–31, 1991. [doi:10.1145/103418.103428] 59

[7] LÁSZLÓ BABAI, LANCE FORTNOW, AND CARSTEN LUND: Non-Deterministic Exponential Time has Two-Prover Interactive Protocols. *Comput. Complexity*, 1(1):3–40, 1991. [doi:10.1007/BF01200056] 59

[8] BRIAN BABCOCK, MAYUR DATAR, AND RAJEEV MOTWANI: Sampling from a moving window over streaming data. In *Proc. 13th Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA'02)*, pp. 633–634, 2002. 62

[9] BRIAN BABCOCK, MAYUR DATAR, RAJEEV MOTWANI, AND LIADAN O'CALLAGHAN: Maintaining variance and k-medians over data stream windows. In *Proc. 22nd Symp. on Principles of Database Systems (PODS'03)*, pp. 234–243, 2003. [doi:10.1145/773153.773176] 62

[10] BONNIE BERGER AND JOHN ROMPEL: Simulating $(\log^c n)$-wise independence in NC. *J. ACM*, 38(4):1026–1046, 1991. Preliminary version in FOCS'89. [doi:10.1145/115234.115347] 59

[11] SERGEY BRIN, RAJEEV MOTWANI, AND CRAIG SILVERSTEIN: Beyond market baskets: Generalizing association rules to correlations. In *1997 ACM SIGMOD Internat. Conf. on Managment of Data*, pp. 265–276, 1997. [doi:10.1145/253262.253327] 61

[12] SERGEY BRIN, RAJEEV MOTWANI, JEFFREY D. ULLMAN, AND SHALOM TSUR: Dynamic itemset counting and implication rules for market basket data. In *1997 ACM SIGMOD Internat. Conf. on Managment of Data*, pp. 255–264, 1997. [doi:10.1145/253262.253325] 61

[13] ANDREI Z. BRODER, MARCUS FONTOURA, VANJA JOSIFOVSKI, RAVI KUMAR, RAJEEV MOTWANI, SHUBHA U. NABAR, RINA PANIGRAHY, ANDREW TOMKINS, AND YING XU: Estimating corpus size via queries. In *15th Internat. Conf. on Information and Knowledge Managment (CIKM'06)*, pp. 594–603, 2006. [doi:10.1145/1183614.1183699] 62

[14] MOSES CHARIKAR, SURAJIT CHAUDHURI, RAJEEV MOTWANI, AND VIVEK R. NARASAYYA: Towards estimation error guarantees for distinct values. In *Proc. 19th Symp. on Principles of Database Systems (PODS'00)*, pp. 268–279, 2000. [doi:10.1145/335168.335230] 62

[15] MOSES CHARIKAR, CHANDRA CHEKURI, TOMÁS FEDER, AND RAJEEV MOTWANI: Incremental clustering and dynamic information retrieval. *SIAM J. Comput.*, 33(6):1417–1440, 2004. Preliminary version in STOC'97. [doi:10.1137/S0097539702418498] 61

[16] MOSES CHARIKAR, RAJEEV MOTWANI, PRABHAKAR RAGHAVAN, AND CRAIG SILVERSTEIN: Constrained TSP and low-power computing. In *Proc. 5th Internat. Workshop on Algorithms and Data Structures (WADS'97)*, pp. 104–115, 1997. [doi:10.1007/3-540-63307-3_51] 60

[17] SURAJIT CHAUDHURI, RAJEEV MOTWANI, AND VIVEK R. NARASAYYA: On random sampling over joins. In *1999 ACM SIGMOD Internat. Conf. on Managment of Data*, pp. 263–274, 1999. [doi:10.1145/304181.304206] 62

[18] CHANDRA CHEKURI, WAQAR HASAN, AND RAJEEV MOTWANI: Scheduling problems in parallel query optimization. In *Proc. 14th Symp. on Principles of Database Systems (PODS'95)*, pp. 255–265, 1995. [doi:10.1145/212433.212471] 60

[19] CHANDRA CHEKURI, RICHARD JOHNSON, RAJEEV MOTWANI, B. NATARAJAN, B. RAMAKRISHNA RAU, AND MICHAEL S. SCHLANSKER: Profile-driven instruction level parallel scheduling

with application to super blocks. In *Proc. 29th Ann. IEEE/ACM Internat. Symp. on Microarchitecture (MICRO'96)*, pp. 58–67, 1996. [doi:10.1109/MICRO.1996.566450] 60

[20] CHANDRA CHEKURI AND RAJEEV MOTWANI: Minimizing weighted completion time on a single machine. In *Proc. 10th Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA'99)*, pp. 873–874, 1999. 60

[21] CHANDRA CHEKURI, RAJEEV MOTWANI, B. NATARAJAN, AND CLIFFORD STEIN: Approximation techniques for average completion time scheduling. *SIAM J. Comput.*, 31(1):146–166, 2001. Preliminary version in SODA'97. [doi:10.1137/S0097539797327180] 60

[22] MAYUR DATAR, ARISTIDES GIONIS, PIOTR INDYK, AND RAJEEV MOTWANI: Maintaining stream statistics over sliding windows. *SIAM J. Comput.*, 31(6):1794–1813, 2002. Preliminary version in SODA'02. [doi:10.1137/S0097539701398363] 62

[23] MIN FANG, NARAYANAN SHIVAKUMAR, HECTOR GARCIA-MOLINA, RAJEEV MOTWANI, AND JEFFREY D. ULLMAN: Computing iceberg queries efficiently. In *Proc. 24th Internat. Conf. on Very Large Data Bases (VLDB'98)*, pp. 299–310, 1998. [VLDB:645924.671338] 61

[24] TOMÁS FEDER AND RAJEEV MOTWANI: Clique partitions, graph compression and speeding-up algorithms. *J. Comput. System Sci.*, 51(2):261–271, 1995. Preliminary version in STOC'91. [doi:10.1006/jcss.1995.1065] 59

[25] TOMÁS FEDER AND RAJEEV MOTWANI: Finding large cycles in Hamiltonian graphs. *Discr. Appl. Math.*, 158(8):882–893, 2010. Preliminary version in SODA'05. [doi:10.1016/j.dam.2009.12.006] 59

[26] TOMÁS FEDER, RAJEEV MOTWANI, AND CARLOS S. SUBI: Finding long paths and cycles in sparse Hamiltonian graphs. In *Proc. 32nd STOC*, pp. 524–529, 2000. [doi:10.1145/335305.335368] 59

[27] URIEL FEIGE, SHAFI GOLDWASSER, LÁSZLÓ LOVÁSZ, SHMUEL SAFRA, AND MARIO SZEGEDY: Interactive proofs and the hardness of approximating cliques. *J. ACM*, 43(2):268–292, 1996. Preliminary version in FOCS'91. [doi:10.1145/226643.226652] 59

[28] PAUL W. FINN, DAN HALPERIN, LYDIA E. KAVRAKI, JEAN-CLAUDE LATOMBE, RAJEEV MOTWANI, CHRISTIAN R. SHELTON, AND SURESH VENKATASUBRAMANIAN: Geometric manipulation of flexible ligands. In *Proc. Workshop on Appl. Comp. Geometry (WACG'96)*, volume 1148 of *LNCS*, pp. 67–78. Springer, 1996. [doi:10.1007/BFb0014486] 60

[29] PAUL W. FINN, LYDIA E. KAVRAKI, JEAN-CLAUDE LATOMBE, RAJEEV MOTWANI, CHRISTIAN R. SHELTON, SURESH VENKATASUBRAMANIAN, AND A. YAO: Rapid: Randomized pharmacophore identification for drug design. *Comput. Geom.*, 10(4):263–272, 1998. Preliminary version in SCG'97. [doi:10.1016/S0925-7721(98)00008-X] 60

[30] MARTIN GAVRILOV, PIOTR INDYK, RAJEEV MOTWANI, AND SURESH VENKATASUBRAMANIAN: Geometric pattern matching: A performance study. In *Proc. 15th Ann. Symp. on Computational Geometry (SCG'99)*, pp. 79–85, 1999. [doi:10.1145/304893.304916] 60

[31] ARISTIDES GIONIS, PIOTR INDYK, AND RAJEEV MOTWANI: Similarity search in high dimensions via hashing. In *Proc. 25th Internat. Conf. on Very Large Data Bases (VLDB'99)*, pp. 518–529. Morgan Kaufmann Publishers Inc., 1999. [VLDB:645925.671516] 61

[32] MICHEL X. GOEMANS AND DAVID P. WILLIAMSON: .878-approximation algorithms for MAX CUT and MAX 2SAT. In *Proc. 26th STOC*, pp. 422–431, New York, 1994. ACM Press. [doi:10.1145/195058.195216] 60

[33] MICHAEL H. GOLDWASSER AND RAJEEV MOTWANI: Intractability of assembly sequencing: Unit disks in the plane. In *Proc. 5th Internat. Workshop on Algorithms and Data Structures (WADS'97)*, volume 1272 of *LNCS*, pp. 307–320. Springer, 1997. [doi:10.1007/3-540-63307-3_70] 60

[34] SUDIPTO GUHA, NINA MISHRA, RAJEEV MOTWANI, AND LIADAN O'CALLAGHAN: Clustering data streams. In *Proc. 41st FOCS*, pp. 359–366, 2000. [doi:10.1109/SFCS.2000.892124] 62

[35] LEONIDAS J. GUIBAS, RAJEEV MOTWANI, AND PRABHAKAR RAGHAVAN: The robot localization problem in two dimensions. In *Proc. 3rd Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA'92)*, pp. 259–268, 1992. 60

[36] WAQAR HASAN AND RAJEEV MOTWANI: Coloring away communication in parallel query optimization. In *Proc. 21st Internat. Conf. on Very Large Data Bases (VLDB'95)*, pp. 239–250, 1995. 60

[37] DAVID HSU, JEAN-CLAUDE LATOMBE, AND RAJEEV MOTWANI: Path planning in expansive configuration spaces. *Int. J. Comput. Geometry Appl.*, 9(4-5):495–512, 1999. Preliminary version in ICRA'97. 60

[38] PIOTR INDYK AND RAJEEV MOTWANI: Approximate nearest neighbors: towards removing the curse of dimensionality. In *Proc. 30th STOC*, pp. 604–613. ACM Press, 1998. [doi:10.1145/276698.276876] 61

[39] PIOTR INDYK, RAJEEV MOTWANI, PRABHAKAR RAGHAVAN, AND SANTOSH VEMPALA: Locality-preserving hashing in multidimensional spaces. In *Proc. 29th STOC*, pp. 618–625, 1997. [doi:10.1145/258533.258656] 61

[40] PIOTR INDYK, RAJEEV MOTWANI, AND SURESH VENKATASUBRAMANIAN: Geometric matching under noise: Combinatorial bounds and algorithms. In *Proc. 10th Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA'99)*, pp. 457–465, 1999. 60

[41] DAVID R. KARGER, RAJEEV MOTWANI, AND G. D. S. RAMKUMAR: On approximating the longest path in a graph. *Algorithmica*, 18(1):82–98, 1997. Preliminary version in WADS'93. [doi:10.1007/BF02523689] 59

[42] DAVID R. KARGER, RAJEEV MOTWANI, AND MADHU SUDAN: Approximate graph coloring by semidefinite programming. *J. ACM*, 45(2):246–265, 1998. Preliminary version in FOCS'94. [doi:10.1145/274787.274791] 60

[43] LYDIA E. KAVRAKI, JEAN-CLAUDE LATOMBE, RAJEEV MOTWANI, AND PRABHAKAR RAGHAVAN: Randomized query processing in robot path planning. *J. Comput. System Sci.*, 57(1):50–66, 1998. Preliminary version in STOC'95. [doi:10.1006/jcss.1998.1578] 60

[44] SANJEEV KHANNA AND RAJEEV MOTWANI: Towards a syntactic characterization of PTAS. In *Proc. 28th STOC*, pp. 329–337, 1996. [doi:10.1145/237814.237979] 59

[45] SANJEEV KHANNA, RAJEEV MOTWANI, MADHU SUDAN, AND UMESH V. VAZIRANI: On syntactic versus computational views of approximability. *SIAM J. Comput.*, 28(1):164–191, 1998. Preliminary version in FOCS'94. [doi:10.1137/S0097539795286612] 59

[46] ALEKSANDRA KOROLOVA, RAJEEV MOTWANI, SHUBHA U. NABAR, AND YING XU: Link privacy in social networks. In *Proc. 24th Internat. Conf. on Data Engineering (ICDE'08)*, pp. 1355–1357. IEEE, 2008. [doi:10.1109/ICDE.2008.4497554] 62

[47] NATHAN LINIAL AND ORI SASSON: Non-expansive hashing. *Combinatorica*, 18(1):121–132, 1998. Preliminary version in STOC'96. [doi:10.1007/PL00009805] 61

[48] CARSTEN LUND, LANCE FORTNOW, HOWARD KARLOFF, AND NOAM NISAN: Algebraic methods for interactive proof systems. *J. ACM*, 39(4):859–868, 1992. [doi:10.1145/146585.146605] 59

[49] GURMEET SINGH MANKU AND RAJEEV MOTWANI: Approximate frequency counts over data streams. In *Proc. 28th Internat. Conf. on Very Large Data Bases (VLDB'02)*, pp. 346–357, 2002. [doi:10.1016/B978-155860869-6/50038-X] 62

[50] RAJEEV MOTWANI, ASSAF NAOR, AND RINA PANIGRAHY: Lower bounds on locality sensitive hashing. *SIAM J. Discrete Math.*, 21(4):930–935, 2007. Preliminary version in SCG'06. [doi:10.1137/050646858] 61

[51] RAJEEV MOTWANI, JOSEPH NAOR, AND MONI NAOR: The probabilistic method yields deterministic parallel algorithms. *J. Comput. System Sci.*, 49(3):478–516, 1994. Preliminary version in FOCS'89. [doi:10.1016/S0022-0000(05)80069-8] 58

[52] RAJEEV MOTWANI, RINA PANIGRAHY, AND YING XU: Estimating sum by weighted sampling. In *Proc. 34th Internat. Colloq. on Automata, Languages and Programming (ICALP'07)*, pp. 53–64, 2007. [doi:10.1007/978-3-540-73420-8_7] 62

[53] RAJEEV MOTWANI, STEVEN PHILLIPS, AND ERIC TORNG: Non-clairvoyant scheduling. *Theoret. Comput. Sci.*, 130(1):17–47, 1994. Preliminary version in SODA'93. [doi:10.1016/0304-3975(94)90151-1] 60

[54] RAJEEV MOTWANI AND PRABHAKAR RAGHAVAN: Deferred data structuring: Query-driven preprocessing for geometric search problems. In *Proc. 2nd Ann. Symp. on Computational Geometry (SCG'86)*, pp. 303–312, 1986. [doi:10.1145/10515.10548] 63

[55] RAJEEV MOTWANI, ARVIND RAGHUNATHAN, AND HUZUR SARAN: Constructive results from graph minors: Linkless embeddings. In *Proc. 29th FOCS*, pp. 398–409, 1988. [doi:10.1109/SFCS.1988.21956] 58

[56] RAJEEV MOTWANI, JENNIFER WIDOM, ARVIND ARASU, BRIAN BABCOCK, SHIVNATH BABU, MAYUR DATAR, GURMEET SINGH MANKU, CHRIS OLSTON, JUSTIN ROSENSTEIN, AND ROHIT VARMA: Query processing, approximation, and resource management in a data stream management system. In *Proc. Conf. on Innovative Data Systems Res. (CIDR'03)*, 2003. 62

[57] LAWRENCE PAGE, SERGEY BRIN, RAJEEV MOTWANI, AND TERRY WINOGRAD: The pagerank citation ranking: Bringing order to the web. Technical report, Stanford InfoLab, 1999. 61

[58] CHRISTOS H. PAPADIMITRIOU AND MIHALIS YANNAKAKIS: Optimization, approximation, and complexity classes. *J. Comput. System Sci.*, 43(3):425–440, 1991. [doi:10.1016/0022-0000(91)90023-X] 59

[59] ADI SHAMIR: IP = PSPACE. *J. ACM*, 39(4):869–877, 1992. [doi:10.1145/146585.146609] 59

[60] AVI WIGDERSON: A new approximate graph coloring algorithm. In *Proc. 14th STOC*, pp. 325–329, 1982. [doi:10.1145/800070.802207] 60

AUTHOR

Prabhakar Raghavan
Google
pragh@acm.org
http://theory.stanford.edu/~pragh/

ABOUT THE AUTHOR

PRABHAKAR RAGHAVAN works at Google and is a Consulting Professor at the Department of Computer Science, Stanford University. He was a long-time friend of Rajeev Motwani.