SPECIAL ISSUE ON WORST-CASE VERSUS AVERAGE-CASE COMPLEXITY

Editors' Foreword

Oded Goldreich and Salil Vadhan

Average-case complexity, which examines the tractability of computational problems on 'random instances,' is a major topic in complexity theory with at least two distinct motivations. On one hand, it may provide a more realistic model than worst-case complexity for the problem instances actually encountered in practice. On the other hand, it provides us with methods to generate hard instances, allowing us to harness intractability for useful ends such as cryptography and derandomization. These two motivations are actually supported by a variety of different notions of average-case complexity (surveyed in [17, 13, 6]) and relating these notions is an important direction for research in the area.

An even more ambitious goal is to understand the relationship between average-case complexity and worst-case complexity, e.g. whether NP $\neq$ P implies that NP has problems that are hard on average. In recent years, there has been substantial progress on this front. This special issue aims to present a small sample of papers that are representative of the different types of results that have been obtained:

Positive Results for High Complexity Classes. There are many results showing equivalences between worst-case complexity and average-case complexity for high complexity classes such #P and EXP [21, 5, 4, 11, 10, 25, 26]. The paper of Trevisan and Vadhan in this issue is one of the most recent in this line, showing that if EXP has a problem that cannot be efficiently solved in the worst case by (uniform) probabilistic algorithms, then it has problem that cannot be efficiently solved on random instances noticeably better than guessing the answer at random. Previous results of this type either referred to hardness against nonuniform algorithms (i.e. Boolean circuits) [18, 25] or lost a substantial amount in the running time of those algorithms [19].

(Partial) Negative Results for NP. It would be very appealing to establish a worst-case/average-case connection like the ones above for NP, as this would be a first step towards basing cryptography on NP-hardness. Unfortunately, there have been a number of negative results on this question [12, 26, 7, 27, 3], which rule out a wide variety of natural approaches to establishing such a connection (under widely believed complexity assumptions). We regret that none of these works are represented in this special issue.

Positive Results for Specific Problems in NP. In a breakthrough, Ajtai [1] gave the first worst-case/average-case connection for a seemingly hard problem in NP, an approximate version of the SHORTEST VECTOR PROBLEM in high-dimensional lattices. Moreover, he showed that if this problem is hard in the worst case, then one can construct secure one-way functions and thus perform many cryptographic tasks. Unfortunately, Ajtai's result does not achieve the goal of basing cryptography on NP-hardness, because the problem used does not seem to be NP-hard [14, 8]. Nevertheless, it provides a new approach to building cryptographic systems, based on worst-case complexity assumptions. Subsequent works addressed a variety of issues regarding Ajtai's work, such as widening the applicability (e.g. to public-key cryptography), weakening the complexity assumption, and improving the efficiency [15, 2, 9, 22, 24, 23]. The paper of Micciancio in this issue falls into the latter category, obtaining a one-way function that can be computed in *nearly linear time* (in n) assuming worst-case hardness of computational problems involving cyclic lattices (of dimension n).

(Partial) Positive Results for All of NP. As mentioned above, it remains a major open problem to show that the worst-case hardness of NP implies the average-case hardness of NP, at least with respect to the usual notions of average-case hardness, where the problem should be hard with respect to a single, efficiently samplable distribution on instances. The paper of Gutfreund, Ta-Shma, and Shaltiel in this issue considers a different notion of average-case complexity, whereby a problem is considered 'easy' if there is an efficient algorithm that works well on *every* efficiently samplable distribution. Remarkably, they show that if NP is worst-case hard, then it is not average-case easy in this sense. That is, if $\mathrm{NP} \not\subseteq \mathrm{BPP}$, then for every efficient algorithm, there exists an efficiently samplable distribution of SAT instances on which the algorithm errs with constant probability. Needless to say, obtaining average-case hardness in the usual sense remains an intriguing open problem.

There are of course many other aspects of research in average-case complexity not addressed in this special issue. One is the study of reductions and completeness in average-case complexity, as initiated by Levin [20]. Another is the study of hardness amplification, where one seeks to convert problems that are mildly hard on average into ones that are very hard on average, in the style of Yao's XOR Lemma [28]. We refer the reader to the surveys [16, 13, 6].

Acknowledgements

We are grateful to the authors for submitting their papers to this special issue and for their patience with the process. We also thank Joachim von zur Gathen for giving us the opportunity to assemble this special issue.

Salil Vadhan's work on this issue was supported by NSF grants CCF-0133096 and CNS-0430336.

References

[1] Miklós Ajtai. Generating hard instances of lattice problems (extended abstract). In *Proceedings of the twenty-eighth annual ACM symposium on the theory of computing – STOC'96*, pages 99–108, Philadelphia, Pennsylvania, 22 May 1996. ACM.

[2] Miklós Ajtai and Cynthia Dwork. A public-key cryptosystem with worst-case/average-case equivalence. In *Proceedings of the Twenty-Ninth Annual ACM Symposium on Theory of Computing*, pages 284–293, El Paso, Texas, 4–6 May 1997.

[3] Adi Akavia, Oded Goldreich, Shafi Goldwasser, and Dana Moshkovitz. On basing one-way functions on NP-hardness. In *STOC'06: Proceedings of the 38th Annual ACM Symposium on Theory of Computing*, pages 701–710, New York, 2006. ACM.

[4] László Babai, Lance Fortnow, and Carsten Lund. Nondeterministic exponential time has two-prover interactive protocols. *Computational Complexity*, 1(1):3–40, 1991.

[5] Donald Beaver and Joan Feigenbaum. Hiding instances in multioracle queries. In *7th Annual Symposium on Theoretical Aspects of Computer Science*, volume 415 of *Lecture Notes in Computer Science*, pages 37–48. Springer, 1990.

[6] Andrej Bogdanov and Luca Trevisan. *Average-Case Complexity*, volume 2 (1) of *Foundations and Trends in Theoretical Computer Science*. now publishers, December 2006.

[7] Andrej Bogdanov and Luca Trevisan. On worst-case to average-case reductions for NP problems. *SIAM Journal on Computing*, 36(4):1119–1159 (electronic), 2006.

[8] Jin-Yi Cai and Ajay Nerurkar. A note on the non-NP-hardness of approximate lattice problems under general Cook reductions. *Information Processing Letters*, 76(1–2):61–66, November 2000.

[9] Jin-Yi Cai and Ajay P. Nerurkar. An improved worst-case to average-case connection for lattice problems (extended abstract). In *38th annual symposium on foundations of computer science – FOCS'97*, pages 468–477, Miami Beach, Florida, 20 October 1997. IEEE.

[10] Jin-Yi Cai, A. Pavan, and D. Sivakumar. On the hardness of the permanent. In *16th International Symposium on Theoretical Aspects of Computer Science*, Lecture Notes in Computer Science, Volume 1563, pages 90–99. Springer-Verlag, 1999.

[11] Uriel Feige and Carsten Lund. On the hardness of computing the permanent of random matrices. *Computational Complexity*, 6(2):101–132, 1996.

[12] Joan Feigenbaum and Lance Fortnow. Random-self-reducibility of complete sets. *SIAM J. on Computing*, 22(5):994–1005, 1993.

[13] Oded Goldreich. Notes on levin's theory of average-case complexity. Technical Report TR97-058, Electronic Colloquium on Computational Complexity, December 1997.

[14] Oded Goldreich and Shafi Goldwasser. On the limits of nonapproximability of lattice problems. *Journal of Computer and System Sciences*, 60(3):540–563, 2000.

[15] Oded Goldreich, Shafi Goldwasser, and Shai Halevi. Collision-free hashing from lattice problems. Technical Report TR96-056, Electronic Colloquium on Computational Complexity (ECCC), 1996.

[16] Oded Goldreich, Noam Nisan, and Avi Wigderson. On Yao's XOR lemma. Technical Report TR95–050, Electronic Colloquium on Computational Complexity, 1995. `http://www.eccc.uni-trier.de/eccc`.

[17] Russell Impagliazzo. A personal view of average-case complexity. In *Proceedings of the Tenth Annual Structure in Complexity Theory Conference*, pages 134–147. IEEE, 1995.

[18] Russell Impagliazzo and Avi Wigderson. $P = BPP$ if E requires exponential circuits: Derandomizing the XOR lemma. In *Proceedings of the Twenty-Ninth Annual ACM Symposium on Theory of Computing*, pages 220–229, 1997.

[19] Russell Impagliazzo and Avi Wigderson. Randomness vs time: derandomization under a uniform assumption. *J. Comput. System Sci.*, 63(4):672–688, 2001. Special issue on FOCS 98.

[20] Leonid A. Levin. Average case complete problems. *SIAM J. on Computing*, 15(1):285–286, 1986.

[21] Richard Lipton. New directions in testing. In *Proceedings of DIMACS Workshop on Distributed Computing and Cryptography*, volume 2, pages 191–202. ACM/AMS, 1991.

[22] Daniele Micciancio. Almost perfect lattices, the covering radius problem, and applications to Ajtai's connection factor. *SIAM Journal on Computing*, 34(1):118–169 (electronic), 2004.

[23] Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on Gaussian measure. *SIAM Journal on Computing*, 37(1):267–302, 2007.

[24] Oded Regev. New lattice-based cryptographic constructions. *Journal of the ACM*, 51(6):899–942 (electronic), 2004.

[25] Madhu Sudan, Luca Trevisan, and Salil Vadhan. Pseudorandom generators without the XOR lemma. *Journal of Computer and System Sciences*, 62:236–266, 2001.

[26] Emanuele Viola. The complexity of constructing pseudorandom generators from hard functions. *Comput. Complexity*, 13(3-4):147–188, 2004.

[27] Emanuele Viola. On constructing parallel pseudorandom generators from one-way functions. In *Proceedings of the Twentieth Annual Conference on Computational Complexity*, pages 183– 197. IEEE, 2005.

[28] Andrew C. Yao. Theory and applications of trapdoor functions (extended abstract). In *23rd Annual Symposium on Foundations of Computer Science*, pages 80–91. IEEE, 1982.

Oded Goldreich
Department of Computer Science
Weizmann Institute of Science
Rehovot, ISRAEL
oded.goldreich@weizmann.ac.il

Salil Vadhan
School of Engineering
and Applied Sciences
Harvard University
Cambridge, MA 02138, USA
salil@eecs.harvard.edu