Letter from the President

Dear EATCS members,

Time flies somehow and the 40th ICALP conference is now upon us. I hope that many of you will attend the conference, which will be held in Riga, Latvia, in the period 8-12 July 2013. The conference programme is now available at

http://www.icalp2013.lu.lv/program/

and is packed with exciting talks and events. Apart from the usual ICALP-related events, on Tuesday, 9 July, the conference will feature a ceremony for the award of honorary doctorates to Jozef Gruska and Juris Hartmanis. The General Assembly of the EATCS will be held on Wednesday, 10 July, and I hope that many of the ICALP attendees will take part in this important yearly meeting for our organization. I remind you that, during the General Assembly in Riga, Kazuo Iwama will present the plans for having ICALP 2015 in Kyoto in the period 6-10 July 2015.

I look forward to the conference and to celebrating the 40th ICALP with you.

Following on the ICALP theme, the preparations for ICALP 2014 in Copenhagen are going well and the first call for papers for that conference will be available at ICALP 2013. The PC chairs for the three tracks of ICALP 2014 (Elias Koutsoupias (University of Oxford) for Track A, Javier Esparza (Technische Universität München) for Track B and Pierre Fraigniaud (Université Paris Diderot) for Track C) have selected very strong program committees and the names of the invited speakers will be available soon.

One of the goals of the EATCS is to recognize outstanding work done in Theoretical Computer Science, broadly construed, by means of awards. I am therefore happy to inform you that the EATCS has joined forces with IPEC to establish the so-called Nerode Prize for outstanding papers in the area of multivariate algorithmics. The Nerode Prize 2013 Committee, consisting of Georg Gottlob (University of Oxford, UK), Rolf Niedermeier (TU Berlin, Germany; chair), and Peter Widmayer (ETH Zurich, Switzerland), has unanimously decided to award Chris Calabro (Google Inc., Mountain View, USA), Russell Impagliazzo (UC San Diego, USA), Valentine Kabanets (Simon Fraser University, Canada), Ramamohan Paturi (UC San Diego, USA), and Francis Zane (Alcatel Lucent, Murray Hill, USA) the 2013 EATCS-IPEC Nerode Prize. You can read the full laudatio for the award, which will be delivered at IPEC 2013 during ALGO 2013, at

http://www.eatcs.org/images/awards/nerode13-laudatio.pdf.

This is an election year for the Council of the EATCS. The electronic election of eight new members of the Council of the EATCS will take place in September 2013. In order to prepare the election, at ICALP 2013 we will have to nominate a list of candidates.

I remind you that nominations of candidates can be made by any EATCS member. In order to allow me to bring the nominations to the General Assembly in Riga and to organize the discussion, you are kindly requested to send your nominations by email to me and to the EATCS Secretary Office

secretary@eatcs.org by June 23. Nominations can also be made by those present at the General Assembly in Riga.

Last, but not least, I take this opportunity to inform you of the changes that we have made regarding the editorship and production of the Bulletin of the EATCS. After serving as editor in chief of the Bulletin for a long time, Maria Serna will step down from this role after this issue. On behalf of the EATCS, I thank Maria for all the work she has put into the editorship of the Bulletin. This is an important service for the community.

From the October 2013 issue, and for a period of four years, the editor in chief of the Bulletin will be Kazuo Iwama (Kyoto University, Japan). I wish Kazuo the best of luck with his new position. I trust that you will offer your assistance in Kazuo's efforts to continue improving the quality and the impact of the flagship publication of the EATCS, in cooperation with the Council of the EATCS and following on Maria's footsteps.

The leadership of the EATCS has decided to decouple the production of the BEATCS from its editorship. Starting from the October 2013 issue, the BEATCS will be produced, printed and shipped by our Secretary Office in Greece. We think that this change was overdue, and that it will help us improve both the production and the scientific quality of the BEATCS. From the October 2013 issue, the editor in chief of the BEATCS will focus solely on the scientific content of the Bulletin.

Let me close by reminding you that you are most welcome to send me any criticism,

idea, opinion or suggestion you might have regarding the EATCS at president@eatcs.org.

Luca Aceto, Reykjavik
June 2013

Letter from the Bulletin Editor

Dear Reader,

Welcome to the June 2013 issue of the Bulletin of the EATCS, which will be my last issue as the Editor-in-chief. Kazuo Iwama will take over from me from issue 111. I wish him the best of luck in this challenging task.

I have been in charge of the Bulletin for five full years, editing 15 issues. I believe that in this period the Bulletin did a step forward achieving open-access publication. I still can see many opportunities for using the newest technologies ahead for the Bulletin, which I am sure Kazuo will help EATCS to achieve.

I wish to thank the EATCS for giving me this great opportunity, and specifically Georgio Ausiello, who entrusted me with this task. Beyond any other considerations, this has been an enriching task for me.

Last but not least, the credit must go to my column Editors and the bulletin contributors, for the work we accomplished and for the support they gave me during these years. Thanks to you all!

I leave now you to explore this volume no 110 of the Bulletin of the EATCS, containing the usual mixture of news and original contributions.

Maria Serna, Barcelona
June 2013

The EATCS Award 2013

Laudatio for Martin Dyer

Martin Dyer has made enormous and multifaceted contributions to Theoretical Computer Science. His work contains deep insights which have opened new frontiers and changed the research landscape. Many of his contributions are landmark results of great originality and beauty. Some of these have already become prominent textbook material. The full impact of these results is becoming increasingly clear as time progresses. We are pleased to recognise these contributions with the 2013 EATCS Award.

Scientific contributions

Dyer's contributions span a wide range of topics within Theoretical Computer Science, including the following.

Pioneering the development of linear-time algorithms for linear programming in a fixed number of dimensions. In the early 1980s, Dyer, and independently, Meggido, discovered the first linear-time algorithms for low-dimensional linear programs. Dyer's first path-breaking result showed that these can be solved in time linear in the number of constraints when the number of variables is at most 3. These techniques were improved by Dyer and Meggido (and subsequently, by others) leading to very-efficient linear-time algorithms which have important applications in computational geometry.

Developing probabilistic analysis of algorithms. Dyer and Frieze made seminal contributions to this area, showing that many NP-hard problems arising in combinatorial optimisation can be solved in polynomial expected time when the instances are drawn from natural distributions.

Discovering the first polynomial-time algorithm for estimating the volume of a highdimensional convex body. In a truly magnificent paper *A random polynomial-time algorithm for approximating the volume of convex bodies*, Dyer,

Frieze and Kannan gave a polynomial-time randomised algorithm for approximating the volume of a convex body in high-dimensional Euclidean space. This is essentially the problem of numerical integration. Classical approaches to solving this problem involve dissecting space into small cubes and adding the contributions from these. The running time of such methods necessarily grows exponentially in the number of dimensions. Dyer, Frieze and Kannan's paper describes the first algorithm for volume estimation whose running time scales polynomially in the number of dimensions. This paper was awarded the 1991 Fulkerson prize for outstanding publications in discrete mathematics. The citation for Kannan's 2011 Knuth prize referred to the result as *one of the most remarkable algorithmic achievements ever*. This result is now the textbook example of a case in which randomisation provably speeds up computation time. The result is one of the earliest, most exciting, and most important applications of the Markov Chain Monte Carlo (MCMC) method in approximation algorithms, and it shaped the way future research in this area has developed.

Introducing the elegant and useful path-coupling technique for bounding the mixing time of Markov chains. The MCMC method is one of the most important techniques for developing approximation algorithms, especially in applications related to approximate counting. These approximation algorithms are based on the simulation of rapidly-mixing Markov chains, and a key scientific challenge is determining the mixing rate of the chains. One of the most fruitful and elegant methods that has been discovered is the *path coupling* technique, which was discovered by Dyer together with his PhD student Bubley. Path coupling is a very simple and intuitive method for proving that a Markov chain is rapidly mixing. Despite its simplicity, it is very effective and easy to use. It has therefore become part of the repertoire of every researcher working in the area of rapid mixing Markov chain algorithms.

Discovering fast algorithms for approximate counting. Dyer has contributed numerous approximate counting algorithms, giving the fastest known algorithms for many significant problems arising in combinatorial enumeration, statistical physics and statisticalhypothesis testing. Most of these are randomised algorithms, but a notable exceptionis his dynamic-programming approach to approximation, which he used to give the firstpolynomial-time approximation algorithm for counting knapsack solutions and for other counting problems.

Classifying the complexity of counting problems. In addition to providing some of the fastest approximate counting algorithms that we have, Dyer has been a leading figure in developing a complexity theory of counting, classifying a wide domain of problems from homomorphism problems to constraint satisfaction problems. For example, his work with Frieze and Jerrum on the hardness of

approximately counting independent sets in bounded-degree graphs pioneered the area of exploiting phase transitions to achieve complexity-theoretic hardness. His work with Greenhill on classifying the complexity of counting graph homomorphisms opened up a whole area of research. Finally, his alternative proof, with Richerby, of Bulatov's #CSP dichotomy theorem is so accessible and elegant that it has spawned substantial further progress in the area.

To acknowledge these extensive and widely-recognized contributions to theoretical computer science, we present the EATCS Award 2013 to Martin Dyer.

The EATCS Awards Committee 2013:

Leslie Ann Goldberg
Friedhelm Meyer auf der Heide (chair)
Vladimiro Sassone

The Gödel Price 2012

Laudatio for

Antoine Joux, Dan Boneh and Matthew K. Franklin

The 2013 Gödel Prize for outstanding papers in Theoretical Computer Science is awarded jointly to the following two papers:

Antoine Joux
A One Round Protocol for Tripartite Diffie-Hellman
Journal of Cryptology, 17(4): 263-276, 2004.
(Conference version: ANTS 2000)

Dan Boneh and Matthew K. Franklin
Identity-Based Encryption from the Weil Pairing
SIAM Journal on Computing, 32(3): 586-615, 2003.
(Conference version: CRYPTO 2001)

ACM's Special Interest Group on Algorithms and Computation Theory (SIGACT) together with the European Association for Theoretical Computer Science (EATCS) will recognize three researchers for their contributions to cryptographic concepts and schemes that provide greater efficiency, flexibility, and security. Their respective papers established the field of pairing-based cryptography by supplying a precise definition of the security of this approach, and providing compelling new applications for it. These applications include better methods for users to exchange the cryptographic keys that enable them to communicate privately and securely with each other. The papers' authors are Antoine Joux, and the team of Dan Boneh and Matthew K. Franklin. They will receive the 2013 Gödel Prize for outstanding papers in theoretical computer science at the ACM Symposium on the Theory of Computing (STOC) June 1-4, in Palo Alto, CA.

In his paper *A One Round Protocol for Tripartite Diffie-Hellman*, Joux's work generalized the two-party key agreement to the multi-party key agreement protocol of Diffie and Hellman, with a focus on the three-party case. His work uses an approach to public-key cryptography based on the algebraic structure of elliptic curves. Joux showed how to implement an elegant tripartite key agreement protocol using pairings on elliptic curves developed by Weil and Tate, and demonstrated that only one broadcast is required for each party.

Boneh and Franklin, in their paper *Identity-Based Encryption from the Weil Pairing*, used Weil pairings on elliptic curves to develop a fully functional identity-based encryption scheme (IBE). It relies on a type of public-key encryption in which the user's public key can be simply the user's identity or email address combined with a single master public key common to all users. This approach replaces the sender's need to obtain a user's public key by direct interaction with the user or via a published database of user public keys, which may be susceptible to corruption.

Antoine Joux is a part-time professor at the Université de Versailles Saint-Quentin-en-Yvelines and a part-time senior security engineer at CryptoExperts. A former member of the Computer Science Department at L'Ecole Normale Supérieure in Paris, he was deputy scientific director of the Central Directorate of Security of Information Systems in France.

Dan Boneh is a professor Computer Science and Electrical Engineering at Stanford University. An editor of the Journal of the ACM, he received a Ph.D. degree in Computer Science from Princeton University. He is a recipient of the Packard Award, the Alfred P. Sloan Award, the Terman Award, and the RSA Award.

A professor of Computer Science at the University of California, Davis, Matthew Franklin is a graduate of Columbia University with a Ph.D. degree in Computer Science. He received an M.A. degree in Mathematics from the University of California, Berkeley, and a B.A. degree in Mathematics from Pomona College. He is editor-in-chief of the Journal of Cryptology. He received a National Science Foundation Career Award, and was an AT&T Bell Labs Ph.D. Scholar.

The Presburger Award 2012

Laudatio for Erik Demaine

The Presburger Award Committee 2013, consisting of Peter Widmayer, Antonin Kucera, and Monika Henzinger (chair), has unanimously decided to propose

Erik Demaine (MIT, USA)

as recipient of the 2013 EATCS Presburger Award for young scientists.

Erik Demaine, born in 1981, has made outstanding contributions in several fields of algorithms, namely computational geometry, data structures, graph algorithms and recreational algorithms.

In data structures he has solved or made significant progress on classic problems such as the carpenter's rule problem, the hinged-dissection problem, the prefix-sum problem, and the dynamic optimality conjecture.

In computational geometry he used the powerful theory of graph minors to develop a suite of algorithms for approximately solving a general family of intractable problems. He also started the new field of computational origami, where his book is the leading authority in the field.

His work has shown promising applications to computer graphics, sensor networks, molecular biology, programmable matter, and manufacturing and engineering.

The committee recommends Erik Demaine as an exceptional young scientist that fully deserves the Presburger Award.

The committee would also like to mention that the quality of all nominations submitted this year was very high. The Presburger Award is attracting the best young scientists in the field of theoretical computer science worldwide.

Report from the Japanese Chapter

R. Uehara (JAIST)

EATCS-JP/LA Workshop on TCS and Presentation Awards

The eleventh *EATCS/LA Workshop on Theoretical Computer Science* was held at Research Institute of Mathematical Sciences, Kyoto University, January 28 to January 30, 2013. **Dr. Takayuki Kihara** (Japan Advanced Institute of Science and Technology) who presented the following paper, was selected at the eleventh EATCS/LA Presentation Award.

> *An application of Computability Theory to the decomposability problem on Borel functions* by Takayuki Kihara (JAIST).

The award will be given him at the Summer LA Symposium held in July 2013.

We also established another presentation award, named "EATCS/LA Student Presentation Award" to encourage students. This time two students, **Mr. Yuji Mihara** (Kyushu University) and **Mr. Hiroyuki Ota** (The University of Tokyo), were selected at the second EATCS/LA Student Presentation Award. They gave the following papers:

> *Randomized Approximate Counting of Forests and Connected Spanning Subgraphs*, by Yuji Mihara, Yukiko Yamauchi, Shuji Kijima, and Masafumi Yamashita (Kyushu University).
>
> *Relativizing Log-space Oracle Hierarchy*, by Hiroyuki Ota and Akitoshi Kawamura (The University of Tokyo).

The award has been given them at the last day, January 30, 2013.

Congratulations!

This workshop is jointly organized with *LA Symposium*, Japanese association of theoretical computer scientists. This symposium has been held since 1970. Its purpose is to give a place for discussing topics on all aspects of theoretical computer science. That is, this workshop is an unrefereed meeting. All submissions are accepted for the presentation. There should be no problem of presenting these papers in refereed conferences and/or journals. We hold it twice a year (January/February, and July/August). If you have a chance, I recommend you to attend it. You can find the program of the last workshop in Appendix of this report.

New TCS Project in Japan

This year, we are glad to announce that two relatively big projects on TCS have been awarded, their main bodies are from EATCS Japan Chapter. One is ELC

Project (Exploring the Limits of Computation, P.I. Osamu Watanabe, 2012–2016) and another is JST ERATO Kawarabayashi Large Graph Project (P.I. Kenichi Kawarabayashi, 2012–2016). Both are planning various research activities including organizing workshops, inviting researchers (short/long, young/senior), publishing extended surveys, hosting major conferences, etc. Please join us with these activities and visit us, Japan!

Appendix: Program of EATCS-JP/LA workshop on TCS (January 28 to January 30, 2013)

In the following program, each [Sx] means student talk, while [x] means ordinary talk (student talks are shorter). Each "**" indicates a student speaker, and "*" indicates just a speaker. Talks are given in the following order (some numbers are skipped due to request by presenters):

[1] Computability of conditional probability
**Kenshi Miyabe (Kyoto University)*

[2] Randomized Approximate Counting of Forests and Connected Spanning Subgraphs
***Yuji Mihara, Yukiko Yamauchi, Shuji Kijima, Masafumi Yamashita (Kyushu University)*

[3] Isomorphism for graphs of bounded width parameters for strong tree decompositions
*Yota Otachi, *Pascal Schweitzer (Japan Advanced Institute of Science and Technology)*

[4] Linear Time Ranking and Unranking of Derangements
**Kenji Mikawa (Center for Academic Information Service, Niigata University), Ken Tanaka (Faculty of Science, Kanagawa University)*

[7] Indexing maximum densities of substrings
**Yoshifumi Sakai (Tohoku University)*

[S1] Red-Black Grammar Compression Algorithm
***Makoto Nishida, Tomohiro I, Shunsuke Inenaga, HIdeo Bannai, Masayuki Takeda (Kyushu University)*

[S2] Lyndon factorization on compressed text
*Tomohiro I, **Yuto Nakashima, Shunsuke Inenaga, Hideo Bannai, Masayuki Takeda (Kyushu University)*

[S3] On computing reversed LZ77 factorization online
***Shiho Sugimoto, Tomohiro I, Shunsuke Inenaga, Hideo Bannai, Masayuki Takeda (Kyushu University)*

[S4] Development of Dynamic Hybrid CEGAR Verifier
***Yanase Ryo, Sakai Tatsunori, Sakai Makoto (Graduate School of Natural Science and Technology, Kanazawa University), Yamane Satoshi (Institute of Science and Engineering, Kanazawa University)*

[S5] Development of the Probabilistic Timed CEGAR Verification Machine with Java
***Shuhei KOIKE, Takashi HASEGAWA, Takaya SHIMIZU (Kanazawa University Graduate School of Natural Science and Technology), Satoshi YAMANE (Kanazawa University Faculty of Electrical and Computer Engineering, Institute of Science and Engineering)*

[9] a Generic Construction for Master Secret Key Leakage-Resilient Identity-Based Encryption
***Takayuki Otsubo, Manh Ha Nguyen (Tokyo Institute of Technology), Ryo Nishimaki (NTT Secure Platform Laboratories), Keisuke Tanaka (Tokyo Institute of Technology)*

[10] Two-Party Computation and Game Theory
***Haruna Higo, Keisuke Tanaka (Tokyo Institute of Technology), Kenji Yasunaga (Institute of Systems, Information Technologies and Nanotechnologies)*

[11] Method of Calculation of Importance Degree for Module
**Takaaki Goto, Tetsuro Nishino (The University of Electro-Communications), Kensei Tsuchida (Toyo University)*

[14] A note on the expansions of insertion systems
**Kaoru Fujioka (Kyushu University)*

[S6] An Analysis of the Width of ZDD Constructed by the Frontier Method
***Keiji Takano (Tokyo Institute of Technology)*

[S7] Layer Operations in Heterogeneous Multiply Layered Tabular Forms with a Hexadecimal Grid Graph Model
***Shinji Koka (Nihon University), Koichi Anada (Waseda University Senior High School), Takeo Yaku (Nihon University)*

[S8] Analysis of DNA restoration algorithms
***ban tomohiro (tokyo institute of technology)*

[S10] Transitivity of Laman Graphs in Distributed System
***Taufiqurrachman, Yukiko Yamauchi, Shuji Kijima, Masafumi Yamashita (Kyushu University)*

[S11] On graphs with a polynomial number of minimal separators
*Ryosuke Nagasawa, Tatsuya Kato, **Toru Kino, Koichi Yamazaki (Gunma Univ.)*

[S12] Deterministic Random Walks for Irrational Transition Probabilities
***Takeharu Shiraga, Yukiko Yamauchi, Shuji Kijima, Masafumi Yamashita (Kyushu University)*

[S13] Hitting time and cover time on dynamic graphs
***Koba Kosuke, Yukiko Yamauchi, Shuji Kijima, Masafumi Yamashita (Kyushu University)*

[S14] A randomized algorithm for comparison between streams
***Naoto Sonoda, Yukiko Yamauchi, Shuji Kijima, Masafumi Yamashita (Kyushu University)*

[S15] Uncapacitated Single Allocation Hub Location Problem
***Ryuta Ando (Chuo University)*

[S16] A Fourier-analytic approach to list-decoding for sparse random linear codes
***Ikko Yamane, Akinori Kawachi (Tokyo Institute of Technology)*

[15] Tight Analysis of Priority Queuing for Egress Traffic
*Koji Kobayashi (National Institute of Informatics), *Jun Kawahara (Japan Science and Technology Agency), Tomotaka Maeda (Kyoto University)*

[16] Combinatorial prediction using offline algorithm
***Takahiro Fujita, Kohei Hatano, Eiji Takimoto (Kyushu University)*

[17] Faster Exact Algorithms for Hybridization Number and rSPR Distance
**Zhi-Zhong Chen (Tokyo Denki University), Lusheng Wang (City University of Hong Kong)*

[S17] Touch typing trainer system

**Aruga Isao (Chuo Univercity)*

[S19] Physical zero-knowledge proof systems for instant insanity
***Keisuke Ueda (Osaka Prefecture University), Harumichi Nishimura (Nagoya University)*

[S18] The adjustment system of musical playback time adapted for user's preference
***Takuro Hidaka (Chuo University)*

[S20] Threshold Circuits for Computing the P_n^D Function
***Daiki Yashima (Graduate School of Information Sciences, Tohoku University), Kei Uchizawa (Graduate School of Science and Engineering, Yamagata University), Xiao Zhou (Graduate School of Information Sciences, Tohoku University)*

[S21] Energy-Efficient Threshold Circuits Detecting Global Pattern in 1-Dimensional Arrays
***Akira Suzuki (Graduate School of Information Sciences, Tohoku University), Kei Uchizawa (Graduate School of Science and Engineering, Yamagata University), Xiao Zhou (Graduate School of Information Sciences, Tohoku University)*

[18] An application of Computability Theory to the decomposability problem on Borel functions
**Takayuki Kihara (Japan Advanced Institute of Science and Technology)*

[19] On representations of analytic functions and polynomial-time computability of operators
**Akitoshi Kawamura (University of Tokyo), Martin Ziegler (Technische Universität Darmstadt), Norbert Müller (Universität Trier), Carsten Rösnick (Technische Universität Darmstadt)*

[20] Relativizing Log-space Oracle Hierarchy
***Hiroyuki Ota, Akitoshi Kawamura (The University of Tokyo)*

[21] On 2-neighborhood CA with discontinuous inverse transition relation
**Shuichi Inokuchi (Kyushu University), Toshikazu Ishida (Kyushu Sangyo University), Yasuo Kawahara (Kyushu University)*

[22] Gliders flying over the Penrose tiling
***Yasuyuki Tsukamoto, Yuhei Miyazaki, Hideki Tsuiki (Kyoto University)*

[23] Information Theory of Cellular Automata
**Hidenosuke Nishio (Kyoto University), Thomas Worsch (Karlsruhe University)*

[24] Reversible multi-head finite automata and space-bounded Turing machines
**Kenichi Morita (Hiroshima University)*

The Japanese Chapter

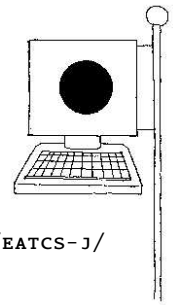

Chair:	Osamu Watanabe
V.Chair:	Kazuhisa Makino
Secretary:	Ryuhei Uehara
email:	eatcs-jp@is.titech.ac.jp
URL:	http://www.misojiro.t.u-tokyo.ac.jp/eatcs-j/

News from New Zealand

BY

C. S. Calude

Department of Computer Science, University of Auckland
Auckland, New Zealand
cristian@cs.auckland.ac.nz

1 Scientific and Community News

0. The latest CDMTCS research reports are (http://www.cs.auckland.ac.nz/staff-cgi-bin/mjd/secondcgi.pl):

432. Y. I. Manin. Zipf's Law and L. Levin's Probability Distributions. 02/2013

433. G. Altmann, I.-I. Popescu and D. Zotta. Stratification in Texts. 02/2013

434. C.S. Calude and K. Tadaki. Spectral Representation of Some Computably Enumerable Sets With an Application to Quantum Provability. 03/2013

435. K. Tadaki and N. Doi. Cryptography and Algorithmic Randomness. 04/2013

436. K. Tadaki. Phase Transition and Strong Predictability. 04/2013

437. A. Probert and M.J. Dinneen. Branchwidth, Branch Decompositions and b-parses. 04/2013

2 A Dialogue with Yuri I. Manin: My Life Is not a Conveyor Belt

Professor Yuri Manin, http://www.mpim-bonn.mpg.de/node/99*, is a member of three institutions based in three different countries: Max Planck Institut für Mathematik, Bonn, Germany, Steklov Mathematical Institute, Academy of Sciences, Moscow, Russia, and Northwestern University, Evanston, USA.*

Professor Manin was educated at Moscow University; his graduate studies have been supervised by I. R. Shafarevich. He obtained famous and important results in extremely diverse mathematical areas—algebraic geometry, number theory, mathematical logic, mathematical physics, informatics. "The field of quantum computing was first introduced by Yuri Manin in 1980 [2] and Richard Feynman in 1981 [3], [4]".[1]

Professor Manin is not "a monomaniac mathematician, but . . . a deep scholar with wide interest, for whom penetration into the mystery of knowledge is much more important than professional success"[2]*. He has also published extensively in literature, linguistics, mythology, semiotics, physics, computer science, philosophy of science and history of culture. His very long list of honours and awards includes invited lectures at universities and congresses*[3]*, prizes, membership in learned academies, honorary degrees and visiting appointments from prestigious organisations around the world. The books* Mathematics and Physics *and* Mathematics as Metaphor *provide a deep insight in his philosophy of science. He supervised 49 PhD students, some outstanding mathematicians themselves.*

CC: What was your motivation to move from an area to a completely different one, not once, but several times? Did these transitions affect your "productivity" in the short term?

YIM: I love mathematics, this great vast realm of human spirit, I am interested in its various aspects, I would like to understand as much of it as I can, and there is no other way than studying and working in various areas in turn. Productivity?? My life is not a conveyor belt, this word is not in my vocabulary.

[1] Wikipedia, "Quantum Computer" https://en.wikipedia.org/wiki/Quantum_computer.

[2] http://www-history.mcs.st-and.ac.uk/Biographies/Manin.html.

[3] Among them, a plenary and five invited lectures at International Congresses of Mathematics, 1966–2006.

CC: The inscription (promoted by academic bureaucrats) "publish or perish" is on all Graduate Schools walls. There is no better rebuttal than your statement: "Productivity? This word is not in my vocabulary".

YIM: "Publish or perish" is a joke, of course... And a mild joke with a grain of sadness is the best way to cope with existential anxiety.

CC: Mathematical logic is a not a core subject for the working mathematician. It is not even taught in many mathematics departments: its home is nowadays in philosophy departments and mainly in computer science departments. How did get interested in mathematical logic?

YIM: I started thinking about mathematical logic when I have already published several dozens research papers in the domains I was trained (algebraic geometry, number theory) and felt the need to expand my overall view of mathematics. Moreover, it was time when Matiyasevich made the last decisive step in the proof of the theorem that all enumerable sets are Diophantine. I could easily understand what are Diophantine sets, but enumerable ones required some study. Turning to books and articles on Logic, I met again what was already a familiar problem: I could not achieve understanding by just reading, other people's texts did not tell me what I felt I needed to know.

The remedy was also by this time well-rehearsed by me: I taught a course on mathematical logic, this time not even in Moscow University where I served as Professor (although my principal job was at the research Steklov Institute for Mathematics), but at the Moscow Institute of Electronic Engineering. My notes of the course were published in 1974, and they became the first draft of my book on Mathematical Logic published by Springer in 1977.

CC: You have developed your own formulation of Kolmogorov complexity in the idiosyncratic book on mathematical logic. In spite of being incomputable, you have successfully applied Kolmogorov complexity to mathematics, physics and linguistics. Your ideas, presented a few years ago at the CiE conference, the largest and arguably most important meeting dedicated to computability theory, have attracted a lot of interest.

YIM: But the mathematical theory of computations is interesting *exactly* because it delineates precise boundaries of the realm of computable, and most interesting things happen when we cross these boundaries! Kolmogorov complexity turned out the great bridge from the land of computable to the vaster realm of mathematics, unconstrained by computability, and, I hope, to physics as well.

I am very happy that during the last several years this vague feeling found justification in three papers of mine where Kolmogorov complexity plays the role of "energy" in three very different contexts: renormalisation in computation, asymptotic boundaries for error-correcting codes as phase transition curves (joint work with Matilde Marcolli), and quite recently, a mathematical explanation of Zipf's

law[4].

CC: Please explain one result in which Kolmogorov complexity works as "energy".

YIM: Consider a finite alphabet A consisting of q letters. An error-correcting block code is a subset $C \subset A^n$, $n \geq 1$ (I am speaking about unstructured codes, but the results I will explain hold, with appropriate modifications, also for linear codes etc.). Each such code C defines a point in the unit square of the plane ($R:=$ *transmission rate,* $\delta :=$ *relative minimal distance*). Now, what will you see if you look at the cloud of *all* code points (in a fixed alphabet)?

In 1981 I have proved that this cloud is everywhere dense below the graph of a certain continuous function, $R \leq \alpha_q(\delta)$ whereas above it each graph point is isolated. This curve $R = \alpha_q(\delta)$ ("silver lining" of the cloud) is called *the asymptotic bound.* In dozens of papers various upper/lower estimates of this function were obtained, but up to now, its exact values are unknown (outside the trivial range), and even whether this function is theoretically computable is open.

In our paper with Matilde we constructed a partition function on the set of all codes, in the sense of statistical physics, in which the energy of the code is exactly its (logarithmic) Kolmogorov complexity. It turned out that after a simple renaming of coordinates, the asymptotic bound becomes the phase transition curve in the (*temperature, density*) plane.

CC: What led you to think about quantum computing?

YIM: First, contemporary computers were electronic devices. They were supposed to produce *exact* results at the level of single bits, hence they had to be designed to suppress quantum effects inherent to any electronic device. Ongoing micro–minituarization was making this task more and more difficult, so it was natural to think about *using* quantum effects rather than suppressing them.

Second, as I have written then, *"the quantum configuration space is much more spacious that the relevant classical one: where in classical physics we have N discrete states, in the quantum theory allowing their superposition there is about c^N states. The union of two classical systems produces N_1N_2 states, whereas the quantum version has $c^{N_1N_2}$ ones."*

CC: The ancient Greek poet Archilocus observed that "the fox knows many things, but the hedgehog knows one big thing". For Freeman Dyson, foxes are mathematical birds (who "fly high in the air and survey broad vistas of mathematics out to the far horizon") and hedgehogs are mathematical frogs (who "live in the mud below and ... solve problems one at a time"). Mathematics needs both birds and frogs.

[4]Zipf's law states that in a natural language corpus, the frequency of a word is inversely proportional to its rank in the frequency table.

Dyson called you a bird. While "globally" this seems to be true, I think that "locally" you have alternated between a bird and a frog. Are you a Francis Bacon's mathematical bee (who "extracts material from the flowers of the gardens and meadows, and digests and transforms it by its own powers")?

YIM: The Swiss writer Max Frisch published in 1953 the ironic comedy "Don Juan oder die Liebe zur Geometrie"[5]. I find its title a wonderfully concise description of my mathematical personality.

Yes, I am a mathematical Don Juan. I still love all my loves, and when I meet my old flame, she can seduce me once again.

Yes, perhaps, all these loves are just incarnations of my deep love for "Geometry" (the latter including algebraic geometry, homotopy topology, quantum field theory . . . and what not).

CC: In a memorable interview published in 1998, you said that "the mathematics of the 20th century is best presented around programmes". Is this trend visible in the 21th century? Is the same true for computer science?

YIM: Probably, it is too early to speak about the trends of the 21th century: imagine an interview on the trends of the 20th century in 1913 . . .

But anyway, I see the full of energy development of programmes I most cherished in the 20th century mathematics: Grothendieck's vast enterprise expanding in many directions thanks to efforts of many strong minds; Langlands' programme; Turing and von Neumann's programmes (each new computer virus is a poisonous descendant of von Neumann's imagination).

CC: The "computer-assisted proofs, as well as computer-unassisted ones, can be good or bad. A good proof is a proof that makes you wiser" is nowadays less controversial than fifteen years ago when you made it. How can a computer-assisted proof make you wiser? What about a quantum computer proof?

YIM: A good proof starts with a project connecting your expected theorem with other results, opening vistas to interesting variations and generalisations. When you develop a detailed plan of it, it might happen that on the road you will have to make a complete list of some exceptional cases, or marginal situations, in which something is not quite as it is "generally", and that making such a list requires a search in a finite but vast set of a priori possibilities.

Then a good computer program that makes this work for you will not spoil the quality of your proof. If the computer is quantum one, then of course you must additionally convince yourself and others that the answer given with "probability close to one" is in fact a correct one.

CC: A paraphrase by C. Anderson (Wired Magazine) falsely attributed to Google's research director Peter Norvig, claims that "all models are wrong, and

[5]Don Juan, or the Love of Geometry.

increasingly you can succeed without them". An example is Google capability to match ads to content without any knowledge or assumptions about the ads or the content, and to translate languages without actually "knowing" them. Why? Google doesn't know why the webpage A is better than the webpage B. However, "if the statistics of incoming links say it is, that's good enough. No semantic or causal analysis is required". From here to the aggressive proclamation of the death of science was just one step: data deluge makes the scientific method obsolete.

YIM: I'll start with stressing that we are speaking about the science rather than market managing.

Now, what Chris Anderson calls "the new availability of huge amounts of data" by itself is not very new: after spreading of printing, astronomic observatories, scientific laboratories, and statistical studies, the amount of data available to any visitor of a big public library was always huge, and studies of correlations proliferated for at least the last two centuries.

Charles Darwin himself collected the database of his observations, and the result of his pondering over it was the theory of evolution.

Even if the sheer volume of data has by now grown by several orders of magnitude, this is not the gist of Anderson's rhetoric.

What Anderson actually wants to say is that human beings are now – happily! – free from thinking over these data. Allegedly, computers will take this burden upon themselves, and will provide us with correlations – replacing the old–fashioned "causations" (that I prefer to call scientific laws) – and expert guidance.

Leaving aside such questions as how "correlations" might possibly help us understand the structure of Universe or predict the Higgs boson, I would like to quote the precautionary tale from J. Groopman. The Body and the Human Progress, *New York Review of Books*, Oct. 27, 2011:

"[...] *in 2000 Peter C. Austin, a medical statistician at the University of Toronto, and his colleagues conducted a study of all 10,674,945 residents of Ontario aged between eighteen and one hundred. Residents were randomly assigned to different groups, in which they were classified according to their astrological signs. The research team then searched through more than two hundred of the most common diagnoses of hospitalization until they identified two where patients under one astrological sign had a significantly higher probability of hospitalization compared to those born under the remaining signs combined: Leos had a higher probability of gastrointestinal haemorrhage while Sagittarians had a higher probability of fracture of the upper arm compared to all other signs combined.*

It is thus relatively easy to generate statistically significant but spurious cor-

relations when examining a very large data set and a similarly large number of potential variables. Of course, there is no biological mechanism whereby Leos might be predisposed to intestinal bleeding or Sagittarians to bone fracture, but Austin notes, 'It is tempting to construct biologically plausible reasons for observed subgroup effects after having observed them.' Such an exercise is termed 'data mining', and Austin warns, 'Our study therefore serves as a cautionary note regarding the interpretation of findings generated by data mining' [. . .]"

Hence my answer to Anderson's question: "What can science learn from Google" is very straightforward: "Think! Otherwise no Google will help you."

CC: Ramsey theory has shown that complete disorder (true randomness) is an impossibility. Every large database (of numbers, points or objects) necessarily contains a highly regular pattern. Most patterns are not computable. Can content–based correlations be distinguished from Ramsey-type correlations?

YIM: I am not an expert. I did not wander far away from the notorious motto[6] about "lies, damned lies, and statistics".

CC: Richard Hamming famously said that "the purpose of computing is insight, not numbers". Do you agree? Do you think that mathematics will continue to be relevant to computer science?

YIM: Yes, and yes.

CC: Will mathematics die? But linguistics?

YIM: Archilocus' fable on the Fox and Hedgehog was re–introduced in our contemporary cultural household by Isaiah Berlin. Berlin had a keen ear for compressed wisdom, and called another of his book of essays after Immanuel Kant: "The Crooked Timber of Humanity". Berlin's message was that all global social projects were doomed: one cannot build a house from crooked timber.

However, we want to be optimists and to believe that human civilisation as we know it for the last two thousand years survives. Then mathematics will survive as well. It is incredibly resilient! My favourite example recently was Pappus' hexagon theorem (Alexandria, about 330 AD), a jump through millennia from Euclid to modernity.

CC: How interesting. Can you explain some details?

YIM: Trying to explain its statement without a picture, I would first suggest to imagine six points in plane, numbered cyclically ("vertices of a hexagon"). Two consecutive points define a line, passing through them, "one side" of this hexagon, there are all in all six sides. Two opposite points define another line, "a diagonal" of this hexagon, there are all in all three diagonals. For each diagonal, there are exactly two sides intersecting this diagonal *not* in vertices. I will say that this diagonal has *Pappus property* if this diagonal and the respective two sides

[6] Benjamin Disraeli.

have just one common point. *The Pappus Theorem* now says that if two of three diagonals have Pappus property, than the third one has it as well.

What immediately strikes anyone looking at the Pappus theorem, is its totally "non–Euclidean" character: neither its statement, nor its proof depends on angles and distances. In fact, it took more than a millennium to understand that Pappus theorem refers to the (real) projective plane, uses only the relation of incidence between lines and points, and, in a hidden form, basic properties of addition and multiplication of real numbers.

A couple centuries later, it became clear that Pappus plane's combinatorics is *completely equivalent* to the axiomatics of abstract fields and abstract projective geometry over them: essentially, his statement taken as an *axiom* is equivalent to the fact that the combinatorics of the incidence relation is an instance of (linear) projective geometry.

Then the whole non–linear algebraic geometry over algebraically closed commutative fields was rewritten in the incidence terms, vastly generalising Pappus, using the theory of models, a chapter in mathematical logic.

And during the last twenty years the abstract Pappus theorem/axiom was used in order to achieve an essential progress in the Alexander Grothendieck's *anabelian programme*.

CC: Many thanks.

The Distributed Computing Column

BY

Panagiota Fatourou

Department of Computer Science, University of Crete
P.O. Box 2208 GR-714 09 Heraklion, Crete, Greece
and
Institute of Computer Science (ICS)
Foundation for Research and Technology (FORTH)
N. Plastira 100. Vassilika Vouton
GR-700 13 Heraklion, Crete, Greece
faturu@csd.uoc.gr

Bulletin of the EATCS no 110, pp. 35–35, June 2013

Computing with Advice: when Knowledge Helps

Stefan Dobrev
Slovak Academy of Sciences, Bratislava, Slovakia.
Stefan.Dobrev@savba.sk

Rastislav Královič
Comenius University, Bratislava, Slovakia.
kralovic@dcs.fmph.uniba.sk

Richard Královič
ETH Zurich, Switzerland.
Google Zurich, Switzerland.
riso@google.com

Abstract

In several areas of computer science the possibility and efficiency of the solution is determined by information that is not accessible to the algorithm. Traditionally, a qualitative approach to the study of this information has been pursued, in which the impact of enhancing the algorithm with various specific types of information has been studied. Recently, a number of authors have proposed a quantitative approach, where the amount of the added information is studied in relation with the improvement of the quality or efficiency of the solution. We survey several recent examples of this approach from the area of distributed and online computing.

1 Introduction

From a high-level point of view, computation is usually thought of as information processing: The input instance contains some implicit, hidden information, and the role of the algorithm is to make this information explicit, which usually

means to produce some specified form of output. While from the information-theoretic point of view, all the relevant information is contained in the input, there may be several reasons why some of this information is not available to the algorithm. First, it may not be possible at all to obtain the required information in the given model of computation, or the algorithm may have limited resources that prevent it from extracting the required information. The questions how the bounded resources relate to the possibility to extract information are among the most notoriously difficult problems of computer science.

There are some cases, however, where information is inaccessible to the algorithm not due to some limited computational resources, but from the nature of the setting. A prominent example is the area of distributed computing where the global state of the system is usually not known to the computing entities, yet it often plays a crucial role in the efficiency (or even feasibility) of the solution. Indeed, many works have been studying the impact of knowledge of the network topology on the efficiency and feasibility of various distributed tasks. Other pieces of information that influence the distributed algorithm are the knowledge of (some) identifiers, and, possibly, the knowledge of failure patterns. In Sections 2 and 3, we survey the results that analyze the impact of topology knowledge from a quantitative point of view: how much information has to be supplied to the computing entities in order to be able to solve tasks efficiently. A similar situation to the distributed computing comes to play in another field of computer science: online algorithms. Here, the algorithm must make irreversible decisions based only on partial knowledge about the input, and coping with the fact that the yet unknown remainder of the input may be crucial for the solution. Again, there is an extensive research concerning the augmentation of the algorithm with some a-priori information about the input, an approach known as semi-online algorithms. Here, too, recently effort has been made to analyze the additional information in a quantitative way, and we survey some of the recent results in Section 4.

Results from both areas show complex behavior of the relationship between the amount of the additional information and the increase of the solution quality: in some cases, there are trade-off relations, where increasing the knowledge gives better solutions, on the other hand, examples of threshold values are known, where adding more bits of advice does not help.

2 Message based distributed computing

Let us consider distributed systems consisting of independent entities connected in a network that can communicate by some form of exchange of messages. There

are two basic views on such systems, which are in essence equivalent, but yield themselves to different types of questions – either the active components are the nodes of the network, and messages are pieces of data send among them, or the active components are the messages (agents) that traverse the network, and the nodes passively provide resources for computation and communication. Typical problems solved in the message-based systems include communication tasks such as broadcasting, wake-up, leader election, or computational problems where some graph-theoretic objects are to be constructed, like, e. g., various spanners, colorings, independent or dominating sets, etc. On the other hand, typical problems tackled in the agent-based view include many variants of graph exploration, map drawing, agent rendezvous, and similar.

The quantitative study of the topological information in message based systems was introduced in the work of Fraigniaud, Ilcinkas, and Pelc [49], where a distinction has been shown between two similar problems: broadcasting, and wakeup. Both problems are considered in an asynchronous setting where nodes have distinct identities, messages are delivered between neighboring pairs of nodes, and in both problems there is an initiator that starts with a message that must be delivered to all other nodes. The distinction is that in the broadcast problem, control messages may be spontaneously sent among vertices from the beginning of the algorithm, whereas in the wakeup problem, only vertices that have previously received a message may send a message (except for the initiator). While it has been known that without any information about the topology, the wakeup requires $\Omega(m)$ messages on a graph with n vertices, and m edges [5], in specific topologies (e. g., [28, 35]) wakeup can be done using $O(n)$ messages. $O(n)$ messages are also sufficient when the network is equipped with the sense of direction [42]. In [49] the authors model the topological information in the following way: a-priori, each node knows its identity, and the local labeling of incident edges. Before the algorithm starts, each node v is provided with a binary string $f(v)$. The function $f : V \mapsto \{0,1\}^*$ is called an oracle, and $\sum_{v \in V} |f(v)|$ is its size. The smallest number of messages, over all oracles of a given size, exchanged by the algorithm is considered as a complexity measure. It is shown that an oracle of size $\Theta(n \log n)$ is needed to perform wakeup with linearly many messages, while linear broadcast can be accomplished with an oracle of size $\Theta(n)$.

The same notion of oracle size has been addressed for a number of other problems in the synchronous setting. Fusco and Pelc [52] consider wakeup in a rooted tree in the one-port model, where each node may send in each step only one message, and the aim is to minimize the number of steps. To evaluate the algorithm they use the competitive analysis. The competitive ratio is the ratio of the broadcasting time of the algorithm with oracle of size q, to the optimal (offline) algorithm with full topological knowledge. The main result shows that with linear-sized oracle,

the broadcasting can be done optimally, and for $\sqrt{n} \leq q \leq n$ the competitive ratio is between $\Omega(n^{1-\varepsilon}/q)$ and $O(n \log^2 n/q)$ for arbitrary small ε. However, advice smaller than $\sqrt{n}$ does not help, since for $q < \sqrt{n}$, the competitive ratio is $\Theta(\sqrt{n})$, the same as without any advice.

A similar situation (in the $\mathcal{LOCAL}$ model from [75],i. e.,, in a synchronous message-passing system with nodes that have unique identifiers) where adding advice does not help has been observed in [47], where the time needed for proper 3-coloring of cycles and trees is investigated. Without any advice, cycles and oriented trees can be 3-colored in time $O(\log^* n)$, and oracle of size $\Omega(n/\log^{(k)} n)$ for any constant k is needed to beat the $O(\log^* n)$ bound, where $\log^{(k)} n$ is the k-th iteration of the logarithm. Moreover, for unoriented trees, the same oracle size is needed for 3-coloring in time $\Theta(\log^* n)$; almost as much information as specifying the color for each node.

In the $\mathcal{LOCAL}$ model, the construction of minimum spanning tree (MST) has been studied as well: in [51], authors consider a setting where each node has access to the weights of incident links, and the goal is to find a distributed representation of a minimum spanning tree. Instead of the overall length of advice strings, they studied the maximum, over all nodes. They show that with constant advice in each vertex, the MST can be constructed in logarithmic time, whereas without any advice, $\Omega(\sqrt{n}/\log n)$ rounds are needed [76].

Broadcasting in radio networks has been considered in [57], where a trade-off between the size of advice and broadcasting time has been devised.

Finally, let us note that the above mentioned work is tightly connected with the study of informative labeling schemes (see, e. g., [20, 22, 45, 62, 63, 64] and references therein): here, the aim is to label vertices of the graph in such a way that it is possible to extract, based solely on the labels of a subset of vertices $V' \subseteq V$ some parameter concerning V' (e. g., if V' is any two-element set, and the parameter is distance, the scheme is called distance labelling scheme).

3 Agent based distributed computing

Alternatively to the model of distributed systems with active nodes that communicate by exchanging messages, one can consider systems where the nodes are passive, and the computation is driven by active messages (agents). The most studied problems in this setting comprise various variants of graph exploration: the agents have to collaboratively explore the network, with possible goals including visiting all vertices, drawing a map, etc. It is always assumed that the incident links in each node are locally distinguishable; in some cases, the nodes

may have also unique identifiers.

The problems related to graph exploration are presumably the oldest graph-theoretic problems (e. g., [39]). The first studied variants concerned a single agent with full topology knowledge, and were focused on the existence of various types of walks. Supposedly the first algorithm for traversing unknown graphs is due to Shannon [80]. In the late 70ties, attention has turned to the problem of a finite automaton navigating in an unknown graph (e. g., [3, 17, 18]), which developed into a series of results concerning the size of memory, and the number of moves of the agent needed for successful exploration (e. g., [48, 72, 73]). At the same time, extensive research has been conducted on teams of cooperating agents (see, e. g., [24, 37, 46]). If not explicitly mentioned, we shall consider undirected graphs (i. e., the agent can always return along the link it arrived). Directed graphs have been treated, e. g., in [2, 8, 25, 41]. Apart from the various variants of graph exploration, problems like rendezvous (e. g., [7, 19, 27, 66]) or black hole search (e. g., [23, 29, 30]) have been investigated.

When considering the additional information, and how it affects the exploration, one should note that the local labelling of the incident links is a potential source of information. On the one hand, if the agent has no means to locally distinguish the incident links in a node, the exploration process can no longer be deterministic, and the adversary may force the agent to traverse a single edge back and forth. Assigning local labels to the incident links in every node is a natural way to circumvent this problem; another approach that is used in some cases is to assign labels to nodes, and to allow the agent to see the labels of neighboring nodes. When using the model with local link identifiers, it is assumed that the labeling is chosen by an adversary, and the agent(s) must be able to perform the task under any labelling. A series of papers [31, 53, 56, 65, 82] investigates how the properly chosen labeling may help the algorithm. It is proven that a memoryless agent (e. g., one using a right-hand-on-the-wall rule) can perform a fast periodic exploration of the network when the local port labels are set appropriately.

The oracle-based approach where additional advice strings can be placed in the nodes was applied in [21], where it is proven that 2 bits in every node are sufficient for a finite automaton to explore all graphs, a task that is not possible without any information.

In the problem of drawing a map of an unlabeled graph, investigated in [26], the symmetry of the graph plays a crucial role: there is a graph invariant called multiplicity (μ), which in some sense expresses the symmetry properties of the graphs, such that for graphs with $\mu = 1$, oracle of size $\varphi(n)$ is sufficient for any function $\varphi = \omega(1)$. On the other hand, for graphs with $\mu > 1$, oracle size $\Theta(m \log \mu)$ is needed, where m is the number of edges. Again, without any information, the task

is not solvable.

In [50], the following problem has been investigated: the agent, starting from a node v, has to traverse all edges of an unknown tree. Obviously, with full information about the tree, this can be done in an optimal number of moves $\texttt{Opt} = 2(n-1) - ecc(v)$ where $ecc(v)$ is the eccentricity of the starting node, i. e., the longest distance from v to another vertex w. For an algorithm $\texttt{A}$, the authors consider the competitive ratio, i. e., the ratio $\text{cost}(\texttt{A})/\texttt{Opt}$. It can be seen that without any further information, the best possible ratio attainable is 2. In order to avoid problems with information given in the port labeling, they use an oracle of the form $f : T \mapsto \{0,1\}^*$, where T is an unlabeled tree T (i. e., the advice given to the agent by the oracle is a bit string that is the same for all labellings of a given tree). The main result shows a tight bound of $\log\log D$ bits in order to get competitive ratio better than 2.

As we already mentioned, there are alternatives to the local port labeling. In the so-called fixed graph scenario introduced by Kalyanasundaram and Pruhs in [59], the nodes have identifiers, and when the agent arrives at a node $v \in G$, it learns all incident edges, their endpoints, and, if the graph is weighted, their weights. While learning the endpoints of the incident edges is stronger than the typical exploration scenario, it does have a justification (see [59] and [69]); it also corresponds to the previously studied neighbourhood sense of direction [43].

In [32], the following problem was addressed from the point of view of advice size: the agent starts at a node v of an undirected labeled graph with n nodes, where each edge has a non-negative cost. The agent has no knowledge about the graph, and has to visit every node of the graph and return to v. The agent can move only along the edges, each time paying the respective edge cost. Clearly, the optimum corresponds to the minimum travelling salesman route (TSP) on the metric closure of the graph (since it is allowed to visit a node more than once). A simple and fast heuristic for the traditional offline setting which has been extensively studied is the greedy algorithm Nearest Neighbor (NN): Once at a node u, go to the closest yet unexplored node, and repeat the process until all nodes have been explored. This algorithm can be also performed by an agent, achieving a competitive ratio of $\Theta(\log n)$ ([78]), which is tight even on planar unit-weight graphs ([55]). Despite many partial results ([4, 59, 70, 69]), the main question, whether there exists a constant-competitive algorithm is still open: the best known lower bound on the competitive ratio is $5/2 - \varepsilon$ ([32]), and the upper bound for general graphs is $O(\log n)$. Moreover, in [32] the authors were able to obtain constant competitive ratio with an advice of size $O(n)$ (for optimality, $\Omega(n \log n)$ bits are needed). Reducing the advice of a constant-competitive algorithm to $o(n)$ bits remains an open problem.

The same concept of advice has been also applied to the graph searching problem in [71].

4 Online Computing

In distributed computing, the information that is not known to the algorithm is the information about the topology of the network. The term online is used for problems where the input comes in parts, and the algorithm must produce output in an incremental fashion, too. Formally, the input $\mathbf{x}$ is a sequence of requests $\mathbf{x} = (x_1, \ldots, x_n)$. The output $\mathbf{y}$ is a sequence of answers $\mathbf{y} = (y_1, \ldots, y_n)$ computed by the algorithm in such a way that each y_i is a function of $x_1, \ldots, x_i$ (for randomized algorithms, it is also a function of the random bits used so far). The goal is to maximize or minimize a cost function defined over the whole output $\mathbf{y}$. For an exposition to online algorithms, we refer the reader to [15].

An archetypal online problem is paging, where the algorithm has to maintain a buffer of k items (pages). The input is a sequence of pages; when a page is requested that is in the buffer, the page is served, and no output is produced. However, if the page is not in memory, the algorithm must select a victim that is removed from the buffer, and is replaced by the requested page; this is called page fault, and the algorithm pays a penalty of 1 for each fault.

The notion of a competitive ratio was introduced by Sleator and Tarjan [81]: a minimization (analogous definition is for maximization) algorithm $\mathtt{A}$ is called c-competitive, if it always produces an output where $\text{cost}(\mathtt{A}) \leq c \cdot \mathtt{Opt} + \alpha$ for some constant α (sometimes, it is required that $\alpha = 0$; this requirement is termed strong competitiveness). Note that in online problems the main concern is not the computational complexity, but the inherent loss of performance due to the unknown future.

Online computation has received considerable attention over the past decades as a natural way of modeling real-time processing of data. A classical result from [81] states that no deterministic paging algorithm can be better than k-competitive. In general, since the algorithm does not know the future input, and because it is compared to the offline optimum in the worst case, many problems have no good competitive algorithms. In order to make the situation less unfair for the algorithm, randomization is often employed. Here, the algorithm has additional access to a random string. In order to be c-competitive, it is sufficient that $\mathrm{E}[\text{cost}(\mathtt{A})] \leq c \cdot \mathtt{Opt} + \alpha$ where the expectation is taken over all random strings. Another well known result [40] states that for paging, randomization helps, since the randomized paging is $\Theta(\log k)$-competitive. The help of randomization in the

case of paging comes, intuitively, from the fact that a particular mistake creates a single page fault: the algorithm pays for it, but it is easy to correct it later. In other cases, many decisions are critical: Consider, e. g., a problem when vertices of a graph are revealed one by one, and the algorithm has to select the largest possible independent set from the resulting graph, while seeing, at each time step, the subgraph induced by the arrived vertices. Clearly, there are situations, where incorrectly selecting a particular vertex prohibits the algorithm from selecting any other vertex for the rest of the input. Hence, it is not difficult to construct a graph and a vertex arrival sequence, where the randomized algorithm selects an expected constant number of vertices, while the graph admits a linearly-sized independent set.

In a way similar to the area of distributed computing, many results have been proven about enhancing the algorithm with a particular type of information about the input. The method of access graphs [16, 58] restricts the sequence of requests to be a walk in an a-priori known graph. A similar approach using entropy has been taken in [74]. Lookahead (e. g., [1]) reveals to the algorithm some limited number of future requests. In many scheduling and graph problems, specific forms of the input sequence have been considered (graphs with certain parameters, jobs arriving in certain order, etc.).

The first attempt to analyze the impact of added information quantitatively was due to Halldórsson et al [54]. The authors considered the problem of finding the maximum independent set online, and introduced a model where the algorithm can maintain a set of solutions. The final solution produced by the algorithm is the best one from the set at the time of the last input request. If the algorithm is allowed to maintain $r(n)$ solutions, this model can be interpreted as running the algorithm with $\log r(n)$ bits of advice describing the particular input. The results of the paper show that when $r(n)$ is constant, the competitive ratio is $\Omega(n)$, i. e., constant advice does not help. However, when $r(n)$ is polynomial, the competitive ratio is $\Theta(n/\log n)$.

In [34], the authors start a systematic quantitative treatment of the problem-specific information. In the proposed model, the algorithm received, in each step, a (possibly empty) advice string, and the advice complexity was defined as the sum of lengths of these strings. However, this model suffered from the fact that information was encoded also in the empty requests, as pointed out in [38]. Here, the authors used an advice string of fixed length attached to each request, and studied the advice complexity of k-server and metrical task system problems. Using this approach, however, it is not possible to analyze information that is sublinear in the number of requests. Therefore, in the model from [13], the whole advice is given to the algorithm at the beginning as a single binary string. In this way, the

model is equivalent to both the model from [54], and the model from agent based systems where the advice is given to the agent. Moreover, it forms an analogy to the model of randomized algorithms – instead of a string of random bits and the expected outcome, the string of best possible bits and the corresponding outcome is considered. Hence, the comparison between randomization and advice is attractive. In the remainder of this section we shall consider the model from [13].

Obviously, there are two trivial ways to obtain an optimal algorithm: either to encode the whole input in the advice (recall that in the treatment of online problems, computational resources are usually disregarded), or to encode the whole output; hence the advice is upper bounded by the minimum of Kolmogorov complexity of those two. In certain cases, however, significantly lower advice is sufficient. A number of problems have been considered in this model, including paging [13], k-server [12], knapsack [14], set cover [61], metrical task systems [38] (the results from the paper hold in both models), buffer management [36], job shop scheduling [13], independent sets in various classes of graphs (general graphs [54], interval graphs [13], bipartite graphs [33]), and various variants of online coloring (bipartite graphs [9], paths [44], 3-colorable graphs [79], $L(2,1)$ coloring [10]).

In general, there are three questions that are usually asked about a problem:

- What advice is needed to get optimal solution?
- What advice is needed to get the competitive ratio of the best possible randomized algorithm?
- What is the relationship between the size of the advice and the competitive ratio?

Usually, the first question is the easiest one to answer, and it turns out that for many problems, large information is needed to be optimal. However, even this large information is sometimes smaller that the trivial bound. For the paging problem, e. g., $O(n)$ bits of advice are sufficient to obtain an optimal algorithm (see [13]) in the following way: with every page in the buffer, the algorithm stores a flag indicating whether the page will be used by a reference optimal algorithm before replacement. When a page fault occurs, the algorithm is safe to remove any page that will not be needed by the optimal solution. The new page is inserted into the cache, and the new flag is read from the advice. On the other hand, to encode the input or the output, $\Omega(n \log k)$ bits would be needed.

The comparison of advice and randomization is an interesting point, since the two approaches use different properties of the solution space: to have a randomized algorithm with good expected performance, many good witnesses for each input

instance are needed. On the other hand, for good performance of advice algorithms, only one witness is sufficient for a given instance, but the space of possible witnesses must be small. In general, it holds (see [12] for minimization problems; analogous statement holds for maximization) that if there is a randomized algorithm with expected worst case competitive ratio $E(n)$, then for any constant $\varepsilon > 0$ there is an algorithm with advice of $O(\log n + \log\log|\mathcal{I}(n)|)$ bits with competitive ratio $(1+\varepsilon)E(n)$, where $\mathcal{I}(n)$ is the set of all input instances consisting of n requests. However, in many cases significantly smaller advice is sufficient to be on par with randomization. For paging, e. g., $\log k$ bits (i. e., independent of n) of advice is sufficient to get competitive ratio $O(\log k)$ ([13]) which equals to the randomized competitive ratio. An interesting point is that an $O(\log k)$-competitive randomized algorithm can be obtained with $O(\log k)$ random bits ([60]), so the random bits and advice bits exhibit the same power in this case.

For an illustration, consider the k-server problem, which is a generalization of paging. In a metric space (finite or infinite) there are k servers located in some points of the space. Each request is some point x_i in the space. To fulfill the request, the algorithm must make sure that there is some server located on x_i; if it is not the case, some server must be moved there, paying the cost of the travelled distance. In the deterministic case, the competitive ratio is known to be $\Theta(k)$, and a famous conjecture states that the randomized competitive ratio is $\Theta(\log k)$. The closest in proving the conjecture is the breakthrough result from [6] that for a finite metric space with β points, there is a randomized k-server algorithm with expected competitive ratio $O(\log^2 k \log^3 \beta \log\log\beta)$. From that follows that there is an algorithm with advice $O(\log n + \log\log\beta)$ having the same competitive ratio. However, the algorithm from [12] runs in exponential time: it first simulates the randomized algorithm on all inputs and all possible random strings, and produces a dictionary of polynomial size. The advice is then a pointer to the dictionary. Existence of a polynomial-time algorithm for k-server that achieves competitive ratio $O(\log k)$ using $O(\log n)$ bits is an open problem.

The relationship between the size of the advice and the competitive ratio is a complex one. In some cases, a trade-off relation exists, where increasing the advice yields a better competitive ratio, as is, e. g., the case of constant competitive ratio of paging [13]. On the other hand, there are thresholds, where increasing the advice does not help, e. g., for simple knapsack [14], no algorithm using less than $\log n$ bits can be better than $(2-\varepsilon)$ competitive, but with $(3(\varepsilon+1)/\varepsilon)\log n + o(\log n)$ bits, competitive ratio $(1+\varepsilon)$ can be achieved for any constant ε.

Notable is also the approach from [11], where an artificial problem of string guessing is analyzed, and a reduction is used to prove lower bounds on the advice complexity of online set cover.

5 Conclusion

Recently, there have been several attempts to analyze the impact of the hidden information in a quantitative way. Although they are applied in different areas, they share a common framework: The algorithm is enhanced by some information about the unknown part of the input, which may be of any type, but of bounded size. This approach may deepen the understanding of the structure of the respective problems. Finally, we note that the term advice complexity has traditionally been used as a synonym for relativized complexity (i. e., a sequential computation where the Turing machine gets an advice that depends on the length of the input), which may cause some confusion. Also, we note similar approaches in the treatment of the problem of factorization ([68, 77]) where the number of queries to a yes/no oracle needed to determine the factors of a number was studied.

References

[1] S. Albers. On the influence of lookahead in competitive paging algorithms. *Algorithmica*, 18(3):283–305, 1997.

[2] S. Albers and M. R. Henzinger. Exploring unknown environments. *SIAM J. Comput.*, 29(4):1164–1188, 2000.

[3] H. Antelmann, L. Budach, and H.-A. Rollik. On universal traps. *Elektronische Informationsverarbeitung und Kybernetik*, 15(3):123–131, 1979.

[4] Y. Asahiro, E. Miyano, S. Miyazaki, and T. Yoshimuta. Weighted nearest neighbor algorithms for the graph exploration problem on cycles. *Information Processing Letters*, 110(3):93 – 98, 2010.

[5] B. Awerbuch, O. Goldreich, D. Peleg, and R. Vainish. A trade-off between information and communication in broadcast protocols. *J. ACM*, 37(2):238–256, 1990.

[6] N. Bansal, N. Buchbinder, A. Madry, and J. Naor. A polylogarithmic-competitive algorithm for the k-server problem. In R. Ostrovsky, editor, *FOCS*, pages 267–276. IEEE, 2011.

[7] L. Barrière, P. Flocchini, P. Fraigniaud, and N. Santoro. Rendezvous and election of mobile agents: Impact of sense of direction. *Theory Comput. Syst.*, 40(2):143–162, 2007.

[8] M. A. Bender and D. K. Slonim. The power of team exploration: Two robots can learn unlabeled directed graphs. In *FOCS*, pages 75–85. IEEE Computer Society, 1994.

[9] M. P. Bianchi, H.-J. Böckenhauer, J. Hromkovič, and L. Keller. Online coloring of bipartite graphs with and without advice. In *Proc. of the 18th Annual International*

Conference on Computing and Combinatorics (COCOON 2012), volume 7434 of *Lecture Notes in Computer Science*, pages 519–530, 2012.

[10] M. P. Bianchi, H.-J. Böckenhauer, J. Hromkovič, S. Krug, and B. Steffen. On the advice complexity of the online $L(2,1)$-coloring problem on paths and cycles. In D. Du and G. Zhang, editors, *COCOON*, volume 7936 of *Lecture Notes in Computer Science*. Springer-Verlag, 2013. to appear.

[11] H.-J. Böckenhauer, J. Hromkovic, D. Komm, S. Krug, J. Smula, and A. Sprock. The string guessing problem as a method to prove lower bounds on the advice complexity. to appear, 2013.

[12] H.-J. Böckenhauer, D. Komm, R. Královič, and R. Královič. On the advice complexity of the k-server problem. In L. Aceto, M. Henzinger, and J. Sgall, editors, *Proc. of the 38th International Colloquium on Automata, Languages and Programming (ICALP 2011)*, volume 6755 of *Lecture Notes in Computer Science*, pages 207–218. Springer-Verlag, 2011.

[13] H.-J. Böckenhauer, D. Komm, R. Královič, R. Královič, and T. Mömke. On the advice complexity of online problems. In Y. Dong, D.-Z. Du, and O. H. Ibarra, editors, *Proc. of the 20th International Symposium on Algorithms and Computation (ISAAC 2009)*, volume 5878 of *Lecture Notes in Computer Science*, pages 331–340. Springer-Verlag, 2009.

[14] H.-J. Böckenhauer, D. Komm, R. Královič, and P. Rossmanith. On the advice complexity of the knapsack problem. In D. Fernández-Baca, editor, *Proc. of the 10th Latin American Symposium on Theoretical Informatics (LATIN 2012)*, volume 7256 of *Lecture Notes in Computer Science*, pages 61–72. Springer-Verlag, 2012.

[15] A. Borodin and R. El-Yaniv. *Online Computation and Competitive Analysis*. Cambridge University Press, 1998.

[16] A. Borodin, S. Irani, P. Raghavan, and B. Schieber. Competitive paging with locality of reference (preliminary version). In C. Koutsougeras and J. S. Vitter, editors, *STOC*, pages 249–259. ACM, 1991.

[17] L. Budach. On the solution of the labyrinth problem for finite automata. *Elektronische Informationsverarbeitung und Kybernetik*, 11(10-12):661–672, 1975.

[18] L. Budach. Environments, labyrinths and automata. In *FCT*, pages 54–64, 1977.

[19] J. Chalopin, S. Das, and P. Widmayer. Rendezvous of mobile agents in directed graphs. In Lynch and Shvartsman [67], pages 282–296.

[20] V. Chepoi, F. F. Dragan, B. Estellon, M. Habib, Y. Vaxès, and Y. Xiang. Additive spanners and distance and routing labeling schemes for hyperbolic graphs. *Algorithmica*, 62(3-4):713–732, 2012.

[21] R. Cohen, P. Fraigniaud, D. Ilcinkas, A. Korman, and D. Peleg. Label-guided graph exploration by a finite automaton. *ACM Transactions on Algorithms*, 4(4), 2008.

[22] R. Cohen, P. Fraigniaud, D. Ilcinkas, A. Korman, and D. Peleg. Labeling schemes for tree representation. *Algorithmica*, 53(1):1–15, 2009.

[23] J. Czyzowicz, S. Dobrev, R. Královic, S. Miklík, and D. Pardubská. Black hole search in directed graphs. In S. Kutten and J. Zerovnik, editors, *SIROCCO*, volume 5869 of *Lecture Notes in Computer Science*, pages 182–194. Springer, 2009.

[24] S. Das, P. Flocchini, S. Kutten, A. Nayak, and N. Santoro. Map construction of unknown graphs by multiple agents. *Theor. Comput. Sci.*, 385(1-3):34–48, 2007.

[25] X. Deng and C. H. Papadimitriou. Exploring an unknown graph. *Journal of Graph Theory*, 32(3):265–297, 1999.

[26] D. Dereniowski and A. Pelc. Drawing maps with advice. In Lynch and Shvartsman [67], pages 328–342.

[27] A. Dessmark, P. Fraigniaud, D. R. Kowalski, and A. Pelc. Deterministic rendezvous in graphs. *Algorithmica*, 46(1):69–96, 2006.

[28] K. Diks, S. Dobrev, E. Kranakis, A. Pelc, and P. Ruzicka. Broadcasting in unlabeled hypercubes with a linear number of messages. *Inf. Process. Lett.*, 66(4):181–186, 1998.

[29] S. Dobrev, P. Flocchini, R. Kralovic, P. Ruzicka, G. Prencipe, and N. Santoro. Black hole search in common interconnection networks. *Networks*, 47(2):61–71, 2006.

[30] S. Dobrev, P. Flocchini, G. Prencipe, and N. Santoro. Searching for a black hole in arbitrary networks: optimal mobile agents protocols. *Distributed Computing*, 19(1), 2006.

[31] S. Dobrev, J. Jansson, K. Sadakane, and W.-K. Sung. Finding short right-hand-on-the-wall walks in graphs. In A. Pelc and M. Raynal, editors, *SIROCCO*, volume 3499 of *Lecture Notes in Computer Science*, pages 127–139. Springer, 2005.

[32] S. Dobrev, R. Královic, and E. Markou. Online graph exploration with advice. In G. Even and M. M. Halldórsson, editors, *SIROCCO*, volume 7355 of *Lecture Notes in Computer Science*, pages 267–278. Springer, 2012.

[33] S. Dobrev, R. Královič, and R. Královič. Independent set with advice: The impact of graph knowledge. In T. Erlebach and G. Persiano, editors, *WAOA*, Lecture Notes in Computer Science, page to appear. Springer, 2012.

[34] S. Dobrev, R. Královič, and D. Pardubská. Measuring the problem-relevant information in input. *RAIRO Theoretical Informatics and Applications*, 43(3):585–613, 2009.

[35] S. Dobrev and P. Ruzicka. Broadcasting on anonymous unoriented tori. In J. Hromkovic and O. Sýkora, editors, *WG*, volume 1517 of *Lecture Notes in Computer Science*, pages 50–62. Springer, 1998.

[36] R. Dorrigiv, M. He, and N. Zeh. On the advice complexity of buffer management. In K.-M. Chao, T. sheng Hsu, and D.-T. Lee, editors, *ISAAC*, volume 7676 of *Lecture Notes in Computer Science*, pages 136–145. Springer, 2012.

[37] M. Dynia, J. Lopuszanski, and C. Schindelhauer. Why robots need maps. In G. Prencipe and S. Zaks, editors, *SIROCCO*, volume 4474 of *Lecture Notes in Computer Science*, pages 41–50. Springer, 2007.

[38] Y. Emek, P. Fraigniaud, A. Korman, and A. Rosén. Online computation with advice. *Theoretical Computer Science*, 412(24):2642–2656, 2011.

[39] L. Euler. Solutio problematis ad geometriam situs pertinentis. *Novi Commentarii Academiae Scientarium Imperialis Petropolitanque*, 7:9–28, 1758-59.

[40] A. Fiat, R. M. Karp, M. Luby, L. A. McGeoch, D. D. Sleator, and N. E. Young. Competitive paging algorithms. *J. Algorithms*, 12(4):685–699, 1991.

[41] R. Fleischer and G. Trippen. Exploring an unknown graph efficiently. In G. S. Brodal and S. Leonardi, editors, *ESA*, volume 3669 of *Lecture Notes in Computer Science*, pages 11–22. Springer, 2005.

[42] P. Flocchini, B. Mans, and N. Santoro. On the impact of sense of direction on message complexity. *Inf. Process. Lett.*, 63(1):23–31, 1997.

[43] P. Flocchini, B. Mans, and N. Santoro. Sense of direction in distributed computing. *Theor. Comput. Sci.*, 291(1):29–53, 2003.

[44] M. Forišek, L. Keller, and M. Steinová. Advice complexity of online coloring for paths. In *Proc. of the 6rd International Conference on Language and Automata Theory and Applications (LATA 2012)*, pages 228–239, 2012.

[45] P. Fraigniaud. Informative labeling schemes. In S. Abramsky, C. Gavoille, C. Kirchner, F. Meyer auf der Heide, and P. G. Spirakis, editors, *ICALP (2)*, volume 6199 of *Lecture Notes in Computer Science*, page 1. Springer, 2010.

[46] P. Fraigniaud, L. Gasieniec, D. R. Kowalski, and A. Pelc. Collective tree exploration. *Networks*, 48(3):166–177, 2006.

[47] P. Fraigniaud, C. Gavoille, D. Ilcinkas, and A. Pelc. Distributed computing with advice: information sensitivity of graph coloring. *Distributed Computing*, 21(6):395–403, 2009.

[48] P. Fraigniaud, D. Ilcinkas, G. Peer, A. Pelc, and D. Peleg. Graph exploration by a finite automaton. *Theor. Comput. Sci.*, 345(2-3):331–344, 2005.

[49] P. Fraigniaud, D. Ilcinkas, and A. Pelc. Oracle size: a new measure of difficulty for communication tasks. In E. Ruppert and D. Malkhi, editors, *PODC*, pages 179–187. ACM, 2006.

[50] P. Fraigniaud, D. Ilcinkas, and A. Pelc. Tree exploration with advice. *Inf. Comput.*, 206(11):1276–1287, 2008.

[51] P. Fraigniaud, A. Korman, and E. Lebhar. Local mst computation with short advice. *Theory Comput. Syst.*, 47(4):920–933, 2010.

[52] E. G. Fusco and A. Pelc. Trade-offs between the size of advice and broadcasting time in trees. In F. Meyer auf der Heide and N. Shavit, editors, *SPAA*, pages 77–84. ACM, 2008.

[53] L. Gasieniec, R. Klasing, R. A. Martin, A. Navarra, and X. Zhang. Fast periodic graph exploration with constant memory. *J. Comput. Syst. Sci.*, 74(5):808–822, 2008.

[54] M. M. Halldórsson, K. Iwama, S. Miyazaki, and S. Taketomi. Online independent sets. *Theor. Comput. Sci.*, 289(2):953–962, 2002.

[55] C. A. Hurkens and G. J. Woeginger. On the nearest neighbor rule for the traveling salesman problem. *Operations Research Letters*, 32(1):1 – 4, 2004.

[56] D. Ilcinkas. Setting port numbers for fast graph exploration. *Theor. Comput. Sci.*, 401(1-3):236–242, 2008.

[57] D. Ilcinkas, D. R. Kowalski, and A. Pelc. Fast radio broadcasting with advice. *Theor. Comput. Sci.*, 411(14-15):1544–1557, 2010.

[58] S. Irani, A. R. Karlin, and S. Phillips. Strongly competitive algorithms for paging with locality of reference. In G. N. Frederickson, editor, *SODA*, pages 228–236. ACM/SIAM, 1992.

[59] B. Kalyanasundaram and K. R. Pruhs. Constructing competitive tours from local information. *Theoretical Computer Science*, 130(1):125 – 138, 1994.

[60] D. Komm and R. Královič. Advice complexity and barely random algorithms. In I. Černá, T. Gyimóthy, J. Hromkovič, K. G. Jeffery, R. Královič, M. Vukolic, and S. Wolf, editors, *Proc. of the 37th International Conference on Current Trends in Theory and Practice of Computer Science (SOFSEM 2011)*, volume 6543 of *Lecture Notes in Computer Science*, pages 332–343. Springer-Verlag, 2011.

[61] D. Komm, R. Královič, and T. Mömke. On the advice complexity of the set cover problem. In E. A. Hirsch, J. Karhumäki, A. Lepistö, and M. Prilutskii, editors, *Proc. of the 7th Symposium on Computer Science in Russia (CSR 2012)*, volume 7353 of *Lecture Notes in Computer Science*, pages 241–252. Springer-Verlag, 2012.

[62] A. Korman. Labeling schemes for vertex connectivity. *ACM Transactions on Algorithms*, 6(2), 2010.

[63] A. Korman, S. Kutten, and D. Peleg. Proof labeling schemes. *Distributed Computing*, 22(4):215–233, 2010.

[64] A. Korman, D. Peleg, and Y. Rodeh. Constructing labeling schemes through universal matrices. *Algorithmica*, 57(4):641–652, 2010.

[65] A. Kosowski and A. Navarra. Graph decomposition for memoryless periodic exploration. *Algorithmica*, 63(1-2):26–38, 2012.

[66] D. R. Kowalski and A. Malinowski. How to meet in anonymous network. *Theor. Comput. Sci.*, 399(1-2):141–156, 2008.

[67] N. A. Lynch and A. A. Shvartsman, editors. *Distributed Computing, 24th International Symposium, DISC 2010, Cambridge, MA, USA, September 13-15, 2010. Proceedings*, volume 6343 of *Lecture Notes in Computer Science*. Springer, 2010.

[68] U. M. Maurer. On the oracle complexity of factoring integers. *Computational Complexity*, 5(3/4):237–247, 1995.

[69] N. Megow, K. Mehlhorn, and P. Schweitzer. Online graph exploration: New results on old and new algorithms. In L. Aceto, M. Henzinger, and J. Sgall, editors, *ICALP (2)*, volume 6756 of *LNCS*, pages 478–489. Springer, 2011.

[70] S. Miyazaki, N. Morimoto, and Y. Okabe. The online graph exploration problem on restricted graphs. *IEICE Transactions on Information and Systems*, 92(9):1620–1627, 2009.

[71] N. Nisse and D. Soguet. Graph searching with advice. *Theoretical Computer Science*, 410(14):1307 – 1318, 2009.

[72] P. Panaite and A. Pelc. Exploring unknown undirected graphs. *J. Algorithms*, 33(2):281–295, 1999.

[73] P. Panaite and A. Pelc. Impact of topographic information on graph exploration efficiency. *Networks*, 36(2):96–103, 2000.

[74] G. Pandurangan and E. Upfal. Can entropy characterize performance of online algorithms? In S. R. Kosaraju, editor, *SODA*, pages 727–734. ACM/SIAM, 2001.

[75] D. Peleg. *Distributed Computing: A Locality-Sensitive Approach*. Monographs on Discrete Mathematics and Applications. Society for Industrial and Applied Mathematics, 2000.

[76] D. Peleg and V. Rubinovich. A near-tight lower bound on the time complexity of distributed minimum-weight spanning tree construction. *SIAM J. Comput.*, 30(5):1427–1442, 2000.

[77] N. Robertson and P. D. Seymour. Graph minors. XIII. The disjoint paths problem. *Journal of Combinatorial Theory, Series B*, 63(1):65–110, 1995.

[78] D. J. Rosenkrantz, R. E. Stearns, and P. M. L. II. An analysis of several heuristics for the traveling salesman problem. *SIAM Journal on Computing*, 6(3):563–581, 1977.

[79] S. Seibert, A. Sprock, and W. Unger. Advice complexity of the online coloring problem. In *Proc. of the 8th International Conference on Algorithms and Complexity (CIAC 2013)*, volume to appear of *Lecture Notes in Computer Science*, page to appear. Springer-Verlag, 2013.

[80] C. E. Shannon. Presentation of a maze solving machine. In H. von Foerster, M. Mead, and H. L. Teuber, editors, *Cybernetics: Circular, Causal and Feedback Mechanisms in Biological and Social Systems, Transactions Eighth Conference, March 15–16, 1951, New York, NY*, pages 169–181, New York, NY, USA, 1951. Josiah Macy Jr. Foundation.

[81] D. D. Sleator and R. E. Tarjan. Amortized efficiency of list update and paging rules. *Communications of the ACM*, 28(2):202–208, 1985.

[82] M. Steinová. On the power of local orientations. In A. A. Shvartsman and P. Felber, editors, *SIROCCO*, volume 5058 of *Lecture Notes in Computer Science*, pages 156–169. Springer, 2008.

The Logic in Computer Science Column

by

Yuri Gurevich

Microsoft Research
One Microsoft Way, Redmond WA 98052, USA
gurevich@microsoft.com

From Reversible Logic Gates to Universal Quantum Bases

Alex Bocharov, Krysta M. Svore*

Abstract

With the anticipated end of Moore's law for integrated circuits [3, 4] fast approaching and continued advances in low-power electronics, interest in quantum computing has increased. This shifts the focus from deterministic logical circuits to potentially more powerful circuits based on controllable quantum systems. In this column, we present a mathematical tour of the quantum circuit model, beginning with reversible logic circuits and expanding to quantum circuits, gates, and measurement. We highlight quantum mechanical phenomena such as superposition, entanglement, and measurement, review the Gottesman-Knill theorem, which states that some subclasses of quantum operations can be simulated efficiently on a classical computer, and describe sets of quantum gates that are universal for quantum computation.

*Quantum Architectures and Computation Group, Microsoft Research, Redmond, WA 98052, USA. {alexeib,ksvore}@microsoft.com

Bulletin of the EATCS no 110, pp. 52–75, June 2013

1 Introduction

There is a growing list of problems for which a quantum algorithm delivers super-polynomial speedup over the corresponding classical algorithms. Most notably, *integer factorization* can be solved exponentially faster using a quantum algorithm than using the best-known classical algorithms [10]. Other problems include *Pell's equation* [11], computing the *unit group and class group* of a number field [12, 13], finding the *hidden shift* of a boolean function [14], solving *linear systems* of equations [15], *group order and membership* [16], *group isomorphism* [17], and *knot invariants* [18, 19]. These algorithms are based on the quantum circuit model of computation.

At the core of the quantum circuit model is *unitary evolution*, which by nature is physically reversible. According to Landauer's principle [2], in order for a computational process to be physically reversible it must be logically reversible. We begin in Section 2 with a short introductory discourse on reversible logic gates which are a special case of classical Boolean gates [5]. We then introduce in Section 3 controllable quantum states and observe that any reversible logic gate generates a unitary quantum gate, however the converse is not true. Quantum state spaces are incomparably richer than boolean logic, as is the hierarchy of controllable gates that we will present.

Since quantum gates are richer than boolean gates, it comes as a surprise that a subclass of quantum circuits can in fact be simulated efficiently on a classical computer [6]. This subclass is commonly called stabilizer circuits or *Clifford* circuits. In Section 4, we introduce the Gottesman-Knill theorem and review some of its key implications.

In Section 5, we introduce the notion of a *universal quantum basis*. After reviewing several such universal bases, we highlight a key result: the quantum analog of the classical *Toffoli* gate, with some help in the form of *measurement and classical feedback*, is universal when added to the group of Clifford circuits. Finally, we conclude in Section 6 with directions for future work.

2 Reversible Logic Gates

Consider an n-bit space $\{0,1\}^n$ and the complete set of Boolean functions of the form $f : \{0,1\}^n \to \{0,1\}^m$, where m is a positive integer. It was realized very early that any Boolean function can be represented as a nested composition of *logic gates*. In fact, in 1881, C.S. Pierce claimed that with just *one* gate, a NOR, or alternatively a NAND, any Boolean function can be realized. For NAND, this was first proven by H.M. Sheffer in 1913 (see [20]). Nested compositions of logic gates are called *logic circuits*.

In this column, we focus on *reversible* Boolean functions. A reversible function f is a computation that can be 'undone', that is, an arbitrary bit vector input x can be reconstructed from the corresponding output vector $f(x)$ and so no input information is erased during the computation of f. We define a *reversible Boolean function* with n arguments as a bijection $\{0,1\}^n \to \{0,1\}^n$. If we consider $\{0,1\}^n$ as a set with 2^n elements, we can also say that such a bijection is an arbitrary permutation of 2^n. Thus there is a one-to-one correspondence between reversible boolean functions and elements of the symmetric group S_{2^n}.

Interestingly, neither NAND$(a,b) = \neg(a \wedge b)$ nor NOR$(a,b) = \neg(a \vee b)$ is reversible. In fact, among those gates commonly appearing in disjunctive or conjunctive normal forms, only NOT: $\{0,1\} \to \{0,1\}$ is reversible. Naturally, for any n, the identity map I_n is also reversible. We denote the identity gate by I.

A binary reversible gate that plays a key role in reversible (and subsequently in quantum) logic is the controlled-NOT gate, written as CNOT, which maps $\{0,1\}^2 \to \{0,1\}^2$ and is defined as CNOT$(a,b) = (a, a \oplus b)$, where $\oplus$ is the exclusive OR (which can be alternatively viewed as addition modulo 2). Intuitively, the first input bit, or the *control* bit, controls the application of the NOT operation on the second bit, called the *target* bit. Evidently, when $a = 0$, b remains unchanged, and when $a = 1$ the second bit is flipped.

We can use simple gates such as NOT and CNOT during synthesis of $n \times n$ reversible Boolean functions $\{0,1\}^n \to \{0,1\}^n$.

To this end, NOT$[i], i = 1, \ldots, n$ denotes the NOT gate applied to the i-th input argument, i.e., NOT$[i]$ replaces the i-th bit of the bit vector $(x_1, \ldots, x_n)$ with NOT(x_i).

We define CNOT$[i, j], 1 \leq i, j \leq n, i \neq j$ as follows:

CNOT$[i, j]$ replaces the j-th bit of the bit vector $(x_1, \ldots, x_n)$ with $x_i \oplus x_j$.

In general, given a $k \times k$ Boolean function $f : \{0,1\}^k \to \{0,1\}^k$, some $n \geq k$, and a *multi-index* $\mathbf{i} = i_1, \ldots, i_k, 1 \leq i_l \leq n, l = 1, \ldots k$, we define the *extension* $f[\mathbf{i}] : \{0,1\}^n \to \{0,1\}^n$ to n-bit logic as follows: if $f(x_{i_1}, \ldots, x_{i_k}) = (y_{i_1}, \ldots, y_{i_k})$ then $f[\mathbf{i}](x_1, \ldots x_n) = replace((x_1, \ldots x_n), x_{i_l}, y_{i_l}, l = 1, \ldots, k)$.

Throughout, we use the $\circ$ symbol to denote the composition of reversible Boolean functions of the same arity, and replace it with a single space when this does not lead to ambiguities. We also use the term 'wire' to refer to logical input bits in the computation.

We now consider several examples of function composition.

Example 1.

1. *NOT*$[i]$ *NOT*$[i] = I$.

2. *CNOT*$[i, j]$ *CNOT*$[i, j] = I$.

Example 2.

1. *We introduce the two-bit SWAP gate that swaps the two logical input bits:*

$$SWAP = CNOT[1,2]\,CNOT[2,1]\,CNOT[1,2]$$

.

2. *Verify that for* $n \geq 2, 1 \leq i, j \leq n, i \neq j$,

$$SWAP[i,j] = CNOT[i,j]\,CNOT[j,i]\,CNOT[i,j].$$

3. *Verify that* $SWAP[i,j] = SWAP[j,i]$ *and* $SWAP[i,j]\,SWAP[i,j] = I$.

Conjugated composition of a reversible function with SWAP is equivalent to re-indexing the arguments of that function.

Example 3.

1. $NOT[j] = SWAP[i,j]\,NOT[i]\,SWAP[i,j]$.

2. *By direct verification on the* i, j, k, l *bits,*

$$CNOT[k,l] = SWAP[i,k]SWAP[j,l]\,CNOT[i,j]\,SWAP[j,l]\,SWAP[i,k].$$

Regarding individual bits as logical wires as in Example 3(2.), we can view the SWAP gate as a mechanism to move both the control and the target bits of a CNOT gate from wire to wire.

Given a set of elementary gates, we can define a n-bit logical *circuit* over that set of gates as a composition of a finite number of extensions of these gates to n-bit logic.

Example 4. *For* $n = 3$, $SWAP[1,2]\,SWAP[2,3]$ *is a circuit implementing a cyclic permutation of bits in a 3-bit vector:* $(x_1, x_2, x_3) \rightarrow (x_2, x_3, x_1)$.

How do we implement a Boolean function with a logical circuit? Example 4 describes a specific implementation where the resulting function is a simple composition of all gates in the circuit.

In general, however, we may need to implement a function with a given number of arguments with a circuit of *greater* arity. Let $f : \{0,1\}^k \rightarrow \{0,1\}^k$ be a Boolean function, and suppose $n > k$. Consider the extension $f[1,\ldots,k]$ of the function f to the n-bit space. Given an n-bit circuit c such that the composition of its gates is equal to $f[1,\ldots,k]$, we say that *circuit* c *implements* f *using* $n-k$ *ancillary bits*.

To get a taste of the constructive side of Boolean function synthesis, let us first characterize Boolean functions that are implementable using only the CNOT gate and no ancillary bits. To this end, let us view the elements of the $\{0,1\}^n$ as bit vectors and interpret the exclusive OR operation, $\oplus$, as bitwise addition of the vectors mod 2. We say that a reversible boolean function $f : \{0,1\}^n \to \{0,1\}^n$ performs a *linear transformation* $f(x \oplus y) = f(x) \oplus f(y)$ for any bit vectors of appropriate dimension.

Theorem 1. *A reversible Boolean function can be represented by a circuit constructed entirely of CNOTs if and only if the function performs a reversible linear transformation.*

Obviously, the identity map is a linear transformation, as is any extension CNOT[i, j] to any bit space. Therefore the 'only if' part of the above theorem is straightforward: composition of any number of linear transformations is a linear transformation. A proof of the less trivial 'if' part can be found, for example, in [21].

Example 5. *None of the extensions of the NOT gate are linear transformations. This can be seen from the fact that when an extension of the NOT gate is applied to the zero bit vector the result is non-zero. Therefore the NOT gate cannot be implemented by a CNOT circuit using any number of ancillary bits.*

Example 6. *Consider a controlled-CNOT gate, CCNOT, which is a ternary gate with the property* $CCNOT(x, y, z) = (x, y, (x \wedge y) \oplus z)$. *This gate is also called a* Toffoli *gate after its inventor Tommaso Toffoli [22]. Toffoli gate is not a linear transformation (e.g., compare* $CCNOT((1,0,0) \oplus (1,1,0))$ *and* $CCNOT(1,0,0) \oplus CCNOT(1,1,0)$*). Therefore it cannot be implemented as a CNOT circuit.*

Even more surprising is that the Toffoli gate cannot be implemented as a circuit combining NOTs and CNOTs. We omit the proof in this column. However, if we add a Toffoli gate to our small library of reversible gates and allow ancillary bits, we can represent all the reversible Boolean functions, as outlined in the following definition and theorem. We refer the reader to [21] for the corresponding proof.

Definition 1. *A composition of various extensions of the CNOT, NOT, and Toffoli gates is called a CNT-*circuit*.*

Theorem 2. *Any reversible Boolean function can be implemented by a CNT-circuit using at most one ancillary bit.*

Exercise 1. *Consider the 4-bit reversible boolean function f defined by the following rules:*

$$f(0,0,0,0) = (1,1,1,1),$$

$$f(0,1,0,1) = (0,0,0,0),$$
$$f(1,0,1,0) = (0,1,0,1),$$
$$f(1,1,1,1) = (1,0,1,0),$$

and f is identity on the remaining 12 bit vectors of $\{0,1\}^4$.

Show that the function f cannot be implemented with a CNT-circuit using no ancillary bits.

Hint As shown in [21], a reversible function can be implemented by a *CNT*-circuit without ancillary bits if and only if the function performs an even permutation of the bit space. Establish that the function f from the above exercise performs an odd permutation of the $\{0,1\}^4$.

3 Quantum States and Quantum Gates

Armed with an understanding of reversible logic gates and circuits, we can now introduce quantum gates and circuits. We begin by describing the principal information unit, the quantum bit, as a quantum system with two basis states. (On a physical level, such states may be represented, for example, by polarizations of a single photon or spin directions of a single electron.)

3.1 Quantum States

In a quantum computation, information is stored in a quantum bit, or *qubit*, which extends the concept of the classical bit. Whereas a classical bit has a state value $s \in \{0,1\}$, a state of a qubit $|\psi\rangle$ is actually a linear *superposition* of basis states:

$$|\psi\rangle = \alpha\,|0\rangle + \beta\,|1\rangle\,, \tag{1}$$

where the $\{0,1\}$ basis state vectors are represented in Dirac notation (ket vectors) as $|0\rangle = (1,0)^T$ and $|1\rangle = (0,1)^T$, respectively. The *amplitudes* α and β are complex numbers that satisfy the normalization condition: $|\alpha|^2 + |\beta|^2 = 1$. Upon *measurement* of the quantum state $|\psi\rangle$, either state $|0\rangle$ or $|1\rangle$ is observed with probability $|\alpha|^2$ or $|\beta|^2$, respectively.

Note that a n-qubit quantum state is a $2^n \times 1$-dimensional state vector, where each entry represents the amplitude of the corresponding basis state. Therefore, n qubits live in a 2^n-dimensional Hilbert space, and we can represent a superposition over 2^n states as:

$$|\psi\rangle = \sum_{i=0}^{2^n-1} \alpha_i\,|i\rangle\,, \tag{2}$$

where α_i are complex amplitudes that satisfy the condition $\sum_i |\alpha_i|^2 = 1$, and i is the binary representation of integer i. Note, for example, that the state $|0000\rangle$ is equivalent to writing the tensor product of the four states: $|0\rangle \otimes |0\rangle \otimes |0\rangle \otimes |0\rangle = |0\rangle^{\otimes 4} = (1,0,0,0,0,0,0,0)^T$. The ability to represent a superposition over exponentially many states with only a linear number of qubits is one of the essential ingredients of a quantum circuit — an innate massive parallelism.

Example 7.

1. *The two-qubit state* $(1/2)(|00\rangle - i\,|01\rangle + i\,|10\rangle + |11\rangle)$ *is a product of the single-qubit state* $(1/\sqrt{2})(|0\rangle + i\,|1\rangle)$ *on the first qubit and the* $(1/\sqrt{2})(|0\rangle - i\,|1\rangle)$ *state on the second qubit.*

2. *The two-qubit state* $(1/\sqrt{2})(|00\rangle + |11\rangle)$ *is not a product of two individual single-qubit states*

The state given in Example 7(2.) possesses a non-classical *entanglement* property. When multi-qubit state cannot be represented as a product of individual single-qubit states, we say that the state is *entangled*. Intuitively, this means that for at least two qubits in the system we cannot in principle identify, or separate out, their individual states.

It follows from the principles of quantum mechanics that two quantum states are indistinguishable if they differ only by a *phase factor* of the form $e^{i\theta}, \theta \in \mathbb{R}$. Thus, we can rewrite a qubit as $\cos(\theta)\,|0\rangle + e^{i\phi}\sin(\theta)\,|1\rangle$, where $0 \leq \phi < 2\pi, 0 \leq \theta \leq \pi/2$. We can interpret 2θ and ϕ as spherical angle coordinates and map the state onto a point on the unit sphere, allowing a geometrical interpretation of single-qubit states, as originally proposed by Felix Bloch [38].

Now consider the evolution of a quantum state. Such evolution would need to preserve the norm of the complex vectors representing the states and would also need to transform a superposition of states into a superposition of transformed states. A physically motivated operation would be, for example, a linear operator $\mathbb{C}^2 \to \mathbb{C}^2$ that preserves the inner product $\langle(\alpha,\beta),(\gamma,\delta)\rangle = \alpha\gamma^* + \beta\delta^*$, where $*$ is complex conjugate transpose. Such operetors are known as *unitary operators* (c.f., [43]).

A quantum computation proceeds through the *unitary* evolution of a quantum state; in turn, quantum operations are necessarily *reversible*. We refer to quantum unitary operations as quantum *gates*. In the multi-qubit case, an n-qubit quantum gate is a $2^n \times 2^n$ unitary matrix acting on an n-qubit quantum state.

We can more formally define a unitary operator U. For a given invertible linear transformation $U : \mathbb{C}^N \to \mathbb{C}^N$, we introduce

$$U^\dagger : \mathbb{C}^N \to \mathbb{C}^N$$

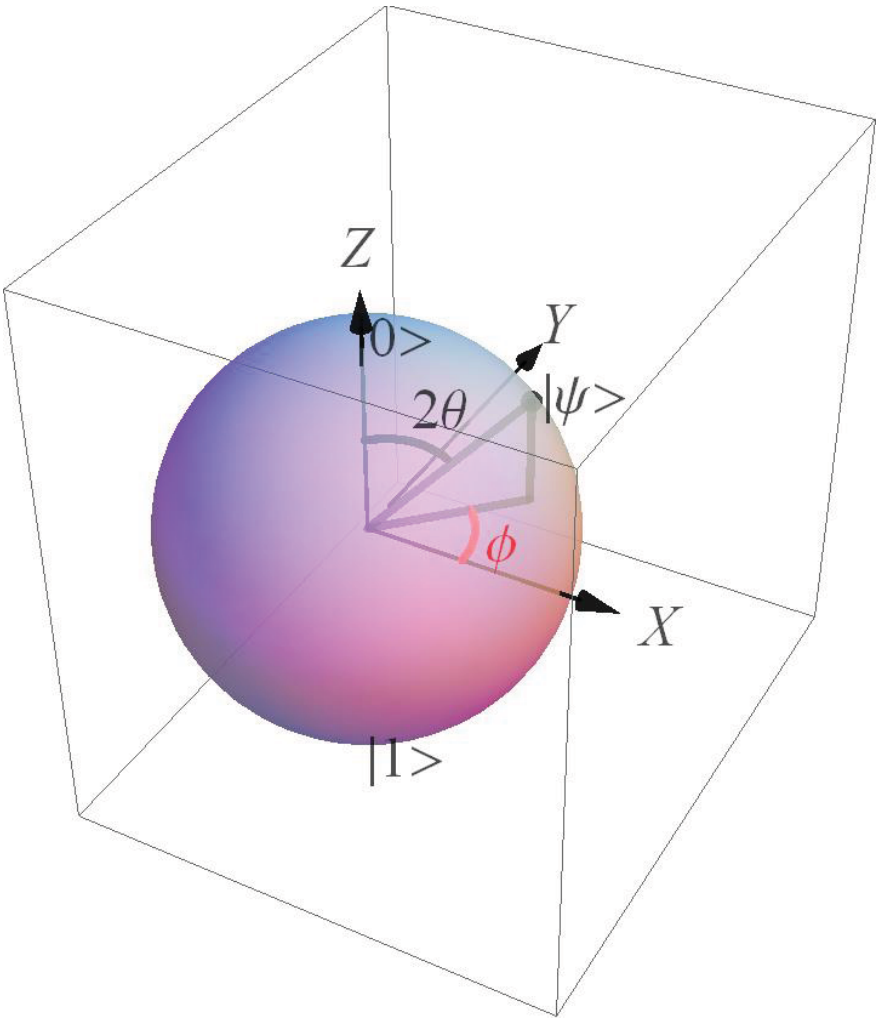

Figure 1: Bloch sphere representation of a single-qubit state. The two basis states, $|0\rangle$ and $|1\rangle$, sit at the two poles.

as the transformation defined by transposition and complex conjugation of the matrix of U.

Definition 2. *A linear operator $U : \mathbb{C}^N \to \mathbb{C}^N$ is* unitary *if $UU^\dagger = I$ (or in other words $U^\dagger$ is the inverse of U).*

Returning to the single-qubit Bloch sphere interpretation, because a unitary operator $A : \mathbb{C}^2 \to \mathbb{C}^2$ transforms valid single-qubit states into valid single-qubit states and because the states map onto points on the Bloch sphere (see Fig. 1), a single-qubit unitary can be interpreted as some transformation of the Bloch sphere. It is not difficult to see that this transformation is, in fact, an isometry. Conversely, each isometry from the special orthogonal group SO(3) corresponds to an equivalence class of single-qubit unitary operators. Because superposition states are defined up to am arbitrary global phase factor, two unitary operators that differ only by a multiplicative $e^{i\theta}$ are considered equivalent, in particular any

unitary is equivalent to one with determinant equal to 1.

3.2 Reversible 'Classical' Gates

Unitary operations on an n-qubit system are, by definition, reversible. Any reversible n-bit classical Boolean function f can be converted into the n-qubit unitary operator U_f by defining the action on the standard basis $|s\rangle$ as follows:

$$U_f(|s\rangle) = |f(s)\rangle\,.$$

We call a unitary operator U_f, where f is a reversible Boolean function, a *classical unitary gate*. From the definition, it follows that U_* is a functor preserving, for reversible Boolean functions f and g, the composition: $U_{f\circ g} = U_f \circ U_g$.

Boolean gates of the universal CNT basis generate the following unitary gates: $X = U_{\mathrm{NOT}}$ is a single-qubit gate with the matrix:

$$\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}.$$

U_{CNOT} is a two-qubit gate, denoted by $\Lambda(X)$, and called CNOT. Its matrix is

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}.$$

Similarly, $U_{\mathrm{CNOT}_{[2,1]}}$ is denoted by $\Lambda(X)[2,1]$ and its matrix is

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix}$$

Since we can compute $U_{f\circ g}$, we can obtain the U_{SWAP} gate:

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}.$$

U_{CCNOT} is a three-qubit, denoted by $\Lambda^2(X)$, and referred to as the Toffoli

gate. Its matrix is

$$\begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix}.$$

The Λ symbol surreptitiously introduced above is actually the functor of adding a control qubit to a unitary. For an n-qubit unitary G, $\Lambda(G)$ is the $n+1$-qubit unitary defined as follows in the standard basis:

$$\Lambda(G)\,|0\,s_1 \ldots s_n\rangle \equiv |0 s_1 \ldots s_n\rangle\,,$$

$$\Lambda(G)\,|1\,s_1 \ldots s_n\rangle \equiv |1\rangle\, G\,|s_1 \ldots s_n\rangle\,.$$

Unlike the classical case where the control bit is a logical switch on the application of the target gate, the meaning of the control qubit is more complex, since the qubit can be in superposition. Nevertheless, at the matrix level, $U_{C(f)} = \Lambda(U_f)$ for any reversible Boolean function f.

An important non-classical single-qubit gate is the *Hadamard* gate H, which maps a quantum state into a quantum superposition state, as follows:

$$H\,|0\rangle = (1/\sqrt{2})(|0\rangle + |1\rangle),$$

$$H\,|1\rangle = (1/\sqrt{2})(|0\rangle - |1\rangle),$$

where the unitary matrix is given by

$$H = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}. \tag{3}$$

There are $2^n!$ reversible Boolean functions and as many 'classical' gates on n qubits. Although the size of this set is double-exponential in n, the set turns out to be very sparse in the infinite continuous group of unitary operators.

3.3 Pauli Gates

The single-qubit *Pauli group* is generated by compositions of the following unitary gates, called *Pauli gates*:

$$I = U_I = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, X = U_{\mathrm{NOT}} = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, Y = \begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix}, Z = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}.$$

Of these generating gates, only X and I are classical. Note that $X^2 = Y^2 = Z^2 = I$.

An easy matrix algebra exercise shows that the Pauli gates generate a group of 16 elements with a 4-element center $z = \{\pm I, \pm i\, I\}$. Since $-1 = e^{i\pi}$ and $\pm\, i = e^{\pm i\,\pi/2}$ are phase factors, each element of z is equivalent to the identity as a single-qubit unitary gate and each element of the Pauli group is equivalent to one of the four gates in $\{I, X, Y, Z\}$

Recall that a single-qubit unitary can be interpreted as a rotation of the Bloch sphere and apply this interpretation to the Pauli gates. Then it is easy to verify that $\{X, Y, Z\}$ are rotations by angle π about axes $\{x, y, z\}$, respectively. Proceeding with the Bloch sphere exercise, we find an easy recipe for writing out the rotation corresponding to any single-qubit unitary A.

Exercise 2. *(1) Prove that given a single-qubit unitary A and $P \in \{X, Y, Z\}$, $A\,P\,A^\dagger = a_P\,X + b_P\,Y + c_P\,Z$ where a_P, b_P, c_P are uniquely defined by* real *coefficients.*

Hint. *Due to the unitary condition $A\,A^\dagger = I$, matrix A is defined up to a phase factor by its first row. More specifically, $A = e^{i\alpha} * B$, where $\alpha \in \mathbb{R}$ and $B = [[u, v], [-v^*, u^*]]$. It is easy to see that $B = Re(u)\,I + i(Im(v)\,X + Re(v)\,Y + Im(u)\,Z)$.*

(2) Prove that in the context of (1)

$$\begin{bmatrix} a_X & a_Y & a_Z \\ b_X & b_Y & b_Z \\ c_X & c_Y & c_Z \end{bmatrix}$$

is a special orthogonal matrix.

(3) Prove that the matrix in (2) defines the rotation of the Bloch sphere corresponding to A.

The multi-qubit Pauli group is, again, generated by the $\{I, X, Y, Z\}$ using tensor products along with the compositions. Let us start with a formal description of the tensor product of unitary operators. Let $n = k + m$ be an integer partition of the natural integer n. We note that the complex vector space $\mathbb{C}^{2^n}$ is represented as tensor product $\mathbb{C}^{2^k} \otimes \mathbb{C}^{2^m}$. One specific tensor representation is written in terms of standard bases in $\mathbb{C}^{2^k}$ and $\mathbb{C}^{2^m}$ by noting that if $\{a_1, \ldots, a_{2^k}\}$ is a basis in the former and $\{b_1, \ldots, b_{2^m}\}$ is a basis in the latter, then the formal pairs $(a_j, b_l), 1 \le j \le 2^k, 1 \le l \le 2^m$ form a basis in a complex vector space of dimension $(2^k)\times(2^m) = 2^n$. For the purposes of the tensor product representation we denote the new basis element (a_j, b_l) as $a_j \otimes b_l$ or simply $a_j b_l$.

Definition 3. *Given two linear operators $A : \mathbb{C}^M \to \mathbb{C}^M$ and $B : \mathbb{C}^N \to \mathbb{C}^N$, $M, N \in \mathbb{N}$, the* tensor product *of A and B is the linear operator $A \otimes B : \mathbb{C}^M \otimes \mathbb{C}^N \to \mathbb{C}^M \otimes \mathbb{C}^N$ uniquely defined by the property*

$$(A \otimes B)(a \otimes b) = (A\,a) \otimes (B\,b).$$

Iterating the construction and the definition, we can represent the n-qubit state space $\mathbb{C}^{2^n}$ as a tensor product of n copies of the one-qubit state space $\mathbb{C}^2$. In particular for any set of unitaries $A_j : \mathbb{C}^2 \to \mathbb{C}^2$, $j = 1, \ldots, n$, the tensor product $A_1 \otimes \cdots \otimes A_n$ is the unitary operator uniquely defined by the property

$$(A_1 \otimes \cdots \otimes A_n)(|\psi_1\rangle \ldots |\psi_n\rangle) = A_1(|\psi_1\rangle) \ldots A_n(|\psi_n\rangle).$$

If $A_1 = \ldots = A_n = A$, then the above tensor product is written as $A^{\otimes n}$.

It is easy to see that, in the context of the above definition, given pairs of operators $A, C : \mathbb{C}^M \to \mathbb{C}^M$ and $B, D : \mathbb{C}^N \to \mathbb{C}^N$,

$$(A \circ C) \otimes (B \circ D) = (A \otimes B) \circ (C \otimes D).$$

We are now ready to define the *n-qubit Pauli group* P_n as the group with respect to composition generated by $\{P_1 \otimes \ldots \otimes P_n | P_i \in \{I, X, Y, Z\}, i = 1, \ldots, n\}$. It is known (c.f., [24]) that this group has 2^{2n+2} elements. For any n , P_n has a 4-element center. Introducing $I_n = I^{\otimes n}$, we can describe the center as $\{\pm I_n, \pm i\, I_n\}$.

Pauli groups are important in quantum information theory. This is due to the fact that it forms the core of the so-called *Heisenberg* representation of quantum computing (see [6]), where quantum information is encoded in eigenstates of certain Pauli operators (we refer the reader to [6] for details.)

3.4 Clifford Group

Consider the group $U(2^n)$ of the n-qubit unitaries. The Pauli group $P_n \subset U(2^n)$ is a subgroup, and is tiny compared to the continuous infinite group $U(2^n)$. We want to look for meaningful ways to extend this group of operations into a larger and more powerful set.

One might logically ask what is a set of unitary operations that *preserves* the Pauli group? (This would be a set of operations that would preserve the Heisenberg computational model mentioned in the previous subsection.) In group theory language, the question would be: what is the *normalizer* of the Pauli group?

Before answering this question, we note that the center of the $U(2^n)$ consisting of the scalar operators of the form $e^{i\theta} I, \theta \in \mathbb{R}$ stabilizes all the elements of $U(2^n)$ and is trivially a part of the normalizer for P_n. This is not at all interesting. The question then must be: what other operators outside the center and the P_n are in the normalizer of P_n?

Here, the *Hadamard* gate H comes to prominence. Recall that H is a single-qubit gate defined as

$$H|0\rangle = (1/\sqrt{2})(|0\rangle + |1\rangle),$$
$$H|1\rangle = (1/\sqrt{2})(|0\rangle - |1\rangle).$$

It is easy to see that $H^2 = I; HXH = Z; HYH = (-1)Y; HZH = X$, therefore H is in the normalizer of the single-qubit Pauli group. Obviously, using H in a tensor product with other normalizer gates generates a normalizer element of the respective multi-qubit Pauli group. For example, $I \otimes H, H \otimes I$ and $H \otimes H$ are in the normalizer of the two-qubit Pauli group.

Another important gate that we must now introduce is the *phase gate* S defined as

$$S\,|0\rangle = |0\rangle\,; S\,|1\rangle = i\,|1\rangle\,.$$

Unlike other gates considered thus far, S is not involutive, rather we see immediately that $S^2 = Z$ so $S^4 = I$, making it an order-4 group element. The inverse is given by $S^\dagger = S^3$.

By direct computation,

$$SXS^\dagger = Y; SYS^\dagger = (-1)X; SZS^\dagger = Z.$$

Thus S is a normalizer element and so are its compositions and tensor products with other normalizer elements. For example, $S \otimes H, S \otimes S$ and $H \otimes S$ are all in the normalizer of the two-qubit Paulis.

We note that the already familiar CNOT gate $\Lambda(X)$ preserves the two-qubit Pauli group (as does any controlled-Pauli gate $\Lambda(P), P \in \{X, Y, Z, (-1)I, \pm i\,I\}$; this is easy to check by direct computation). We also note the following amusing two-qubit identity:

$$(H \otimes H) \circ \Lambda(X) \circ (H \otimes H) = \Lambda(X)[2, 1].$$

Theorem 3. *The normalizer of the n-qubit Pauli group in $U(2^n)$ is generated by the center $z(U(2^n))$ (the subgroup of scalar unitaries), tensor products of I, H, S operators, and various CNOT operators $\Lambda(X)[j, l], 1 \leq j < l \leq n$.*

A proof of an equivalent theorem can be found in Chapter 10 of [23].

Definition 4. *The* Clifford group *is the group of unitary operators, generated by*

(1) H and S in the single-qubit case

(2) Tensor products of I, H, S and all $\Lambda(X)[j, l], 1 \leq j < l \leq n$ in the n-qubit case, $n > 1$.

For any number of qubits, the Clifford group is the "non-trivial part" of the normalizer of the Pauli group.[1] Interestingly, although the term *Clifford group* is

[1]It is commonly claimed that the Clifford group *is* the normalizer of the Pauli group. Strictly speaking, this claim is incorrect. It is only meaningful 'modulo scalar operators'. More precisely, the central quotient of the Clifford group is the normalizer of the central quotient of the Pauli group in the central quotient $PU(2^n)$ of the unitary group.

universally accepted, its origin is not entirely clear (it is not directly related to the class of Clifford algebras). According to D. Gottesman, the first use of the term is attributed to Eric M. Rains.

If we view Clifford operators as "instructions" on a quantum computer, we get a rather large instruction set. As per [24], the n-qubit Clifford group has $2^{n^2+2n+3}\prod_{j=1}^{n}4^j-1$ distinct elements. For example, this amounts to 92,160 elements in the two-qubit case. Regrettably, this instruction set does not provide any speedups for quantum computers over classical computers. This remarkable result is the subject of the next section.

4 Gottesman-Knill Theorem

In order to harness the power of a quantum computer, we need to first step away from the unitary operator paradise and introduce some non-unitary operations. We start this section by discussing *quantum measurement* and *classical feedback* operations.

4.1 Measurement and Classical Feedback

It is one of the great mysteries of quantum mechanics that measurements are parameterized by Hermitian operators. An operator $M : \mathbb{C}^N \to \mathbb{C}^N$ is called *Hermitian* if $M = M^\dagger$. Obviously all the eigenvalues of such an operator are real. It follows that if M is both unitary and Hermitian, then it is also involutive, i.e., $M^2 = I$, with eignevalues ± 1. Note, for example, that a generator of the Clifford group is Hermitian iff it does not explicitly contain the S gate.

If $\{m_1, \ldots, m_l\}$ is the list of distinct eigenvalues of a Hermitian operator M, then it is a simple algebraic fact that

$$M = \sum_j m_j \, Pr_j,$$

where Pr_j is the projector onto the eigenspace of M corresponding to the eigenvalue m_j. Conceptually, according to the postulates of quantum mechanics, a *measurement* of the operator M on a quantum state $|\psi\rangle$ must produce one of the eigenvalues of M. Any of the eigenvalues may randomly result from the measurement.

We need a way to compute the probability p_j of observing a certain eigenvalue m_j in the measurement. To this end, consider the projection $Pr_j|\psi\rangle$ of a state vector $|\psi\rangle$ on the jth eigenspace. Note that the scalar product $\langle\psi|Pr_j|\psi\rangle$ is a measure of proximity of the state vector $|\psi\rangle$ to the eigenspace, very similar to

the cosine of the angle between a vector and a plane in Euclidean space (the larger the cosine, the smaller the angle).

It is the *measurement postulate* of quantum mechanics that defines p_j as $p_j = \langle\psi| Pr_j |\psi\rangle$. It is easy to see that $\sum_j p_j = \langle\psi|\psi\rangle = 1$.

The fundamental *state reduction* principle also states that *if the eigenvalue m_j is observed in measuring the operator M, the quantum state $|\psi\rangle$ is changed ('collapsed') into $Pr_j |\psi\rangle$ post-measurement*.

The probability of observing an eigenvalue of M can thus be also understood as the probability of the quantum state being forced into the corresponding eigenspace. It makes intuitive sense that such a probability is proportional to a proximity measure between the state and the eigenspace.

Example 8. *The simplest scenario is measuring the Pauli operator Z on the single-qubit state $\alpha|0\rangle + \beta|1\rangle$. The eigenvalues of Z are $+1$ and -1 and the probabilities of observing each one are $|\alpha|^2$ and $|\beta]^2$ respectively. Post-measurement the state collapses to either basis state $|0\rangle$ or to basis state $|1\rangle$. The standard notation for this measurement procedure is M_Z. In multi-qubit case we will use the notation $M_Z[i]$ for the Z-measurement applied to i^{th} qubit only.*

One key application of measurements in quantum computation is the feeding of the measurement results back into quantum circuits to be used as *classical control bits*. The non-unitary primitive that makes such feedback possible is called a *classically controlled gate*. Given an operator $G \in U(N)$, the classically controlled gate

$$BC(G) : (\{0,1\} \times \mathbb{C}^N) \rightarrow (\{0,1\} \times \mathbb{C}^N)$$

is defined by

$$BC(G)((0,v)) = (0,v);\ BC(G)((1,v)) = (1, G\,v).$$

4.1.1 An Important Toffoli-based Construction

We now consider an important measurement example that is significantly more sophisticated than Example 8, and introduces the concept of a classical feedback loop. We will implement a certain single-qubit rotation that will turn out to be important in the next section.

Consider a single-qubit state $|\psi\rangle$ and add two ancillary qubits prepared in state $|0\rangle$. To simplify notations, assign indices 1 and 2 to the ancillary qubits and assign index 3 to the qubit in the state $|\psi\rangle$. The resulting 3-qubit system is initially in the product state $|00\rangle|\psi\rangle$.

Consider operator $H \otimes H \otimes I$ that performs the Hadamard gate on qubits 1 and 2 and leaves qubit 3 unchanged; consider operator $I \otimes I \otimes S$ that leaves qubits 1 and 2 unchanged while performing the phase gate S on the third qubit. Build the

3-qubit circuit $U = (H \otimes H \otimes I)\, \Lambda^2(X)\, (I \otimes I \otimes S)\, \Lambda^2(X)\, (H \otimes H \otimes I)$, apply it to the 3-qubit system prepared in state $|00\rangle\,|\psi\rangle$, then apply M_Z measurement operator to qubits 1 and 2.

Note that the Toffoli gate $\Lambda^2(X)$ does not belong to the Clifford group and neither does the composite U. Given $|\psi\rangle = a\,|0\rangle + b\,|1\rangle$, $|a|^2 + |b|^2 = 1$, by direct computation we obtain

$$\begin{aligned} U\,|00\rangle\,|\psi\rangle = (1/4)\,((3+i)\,a\,|000\rangle + (1+3i)\,b\,|001\rangle + \\ (1-i)\,a\,|010\rangle - (1-i)\,b\,|011\rangle + (1-i)\,a\,|100\rangle - \\ (1-i)\,b\,|101\rangle + (i-1)\,a\,|110\rangle - (i-1)\,b\,|111\rangle). \end{aligned}$$

We introduce the Clifford gate $IIZ = I \otimes I \otimes Z$ that leaves qubits 1 and 2 unchanged and performs a Pauli-Z gate on the third qubit and introduce the classically controlled gate $BC(IIZ)(M_Z[1] \vee M_Z[2], *)$, where the notation reads: apply the IIZ gate *unless* $M_Z[1] = |0\rangle$ and $M_Z[2] = |0\rangle$. Now apply the composition $BC(IIZ)(M_Z[1] \vee M_Z[2], *) \circ U$ to the $|00\rangle\,|\psi\rangle$ state.

As per the state reduction principle, when $M_Z[1] = |0\rangle$ and $M_Z[2] = |0\rangle$, the $U\,|00\rangle\,|\psi\rangle$ is projected to the $+1$ eigenspace of $Z \otimes I \otimes I$, then to $+1$ eigenspace of $I \otimes Z \otimes I$. In other words, the state gets projected onto the two-dimensional space spanned by $|000\rangle$ and $|001\rangle$. From the above expression for $U\,|00\rangle\,|\psi\rangle$, we derive that the projected state vector is proportional to $(3+i)\,a\,|000\rangle + (1+3\,i)\,b\,|001\rangle$ and thus it is equivalent to $a\,|000\rangle + ((1+3\,i)/(3+i))\,b\,|001\rangle = a\,|000\rangle + (\frac{3+4\,i}{5})\,b\,|001\rangle$.

To summarize, the case when measurement outcomes are $|0\rangle$ is equivalent to applying the $V = [[1,0],[0,\frac{3+4\,i}{5}]]$ gate to the third qubit. Looking at the remainder of the expression for $U\,|00\rangle\,|\psi\rangle$, it is easy to see that all other outcomes are equivalent to applying the Pauli-Z gate to the third qubit, which is then canceled out by the $BC(IIZ)(M_Z[1] \vee M_Z[2], *)$ operator.

Finally, we estimate the probability of measurement outcomes being simultaneously $|0\rangle$. As per the measurement postulate above, that probability is $p_{00} = |(3+i)/4|^2\,|a|^2 + |(1+3\,i)/4|^2\,|b|^2 = 5/8\,(|a|^2 + |b|^2) = 5/8$. Note, for now, that using the above protocol, the probability of performing the gate V on the third qubit is higher than the probability of leaving the third qubit state unchanged.

4.2 The Theorem and Discussion

Daniel Gottesman [6] and, independently, Emmanuel Knill, conjectured (and later proved) that certain quantum circuits, when containing only a subset of quantum operations and measurements, could be efficiently computed on a classical computer. Informally, if the computer uses only, for example, the gates within the Clifford group and measurements in the computational basis, then it is no more

powerful (and in fact, more restricted) than a classical computer. One of the most common versions of this result is articulated in the following theorem.

Theorem 4. *The result of applying a sequence of Clifford gates followed by a Pauli measurement to the input state* $|\mathbf{0}\rangle = |0\rangle^{\otimes N}$ *can be simulated in polynomial time on a probabilistic classical computer.*

The practical corollary is that if we only use Clifford gates for both preparation of a quantum state and the evolution of the quantum state, then this computation will not have an exponential speed-up over the corresponding classical computation.

In [8], Maarten Van den Nest established a slightly more general result regarding the 'classicality' of certain circuits. Recall that the quantum Toffoli gate $\Lambda^2(X)$ does not belong to the Clifford group. (As an exercise, check $\Lambda^2(X)(X \otimes I \otimes I)\Lambda^2(X)$.) Consider, however, a type of circuit called *H-Toffoli* that consists of two decoupled parts: the first part is a multi-qubit Hadamard gate $H^{\otimes N}$ and the second part is an arbitrary N-qubit circuit composed of the NOT, CNOT and Toffoli gates. Because classical {NOT, CNOT, Toffoli} constitute a universal basis in the the group of reversible classical circuits, we note that the second part is a quantum wrapper around *arbitrary* reversible boolean function, i.e., it is of the form U_f, where f is the reversible boolean function computed by the classical circuit replicating the {NOT, CNOT, Toffoli} part of the quantum circuit.

Then any H-Toffoli circuit followed by a Pauli measurement has the same computational power as a probabilistic classical computation. Intuitively, the result may be not so unexpected given the circuit is $H^{\otimes N} U_f$, where f is classical. This is surprising though: if we allow the Hadamard gate and the quantum Toffoli gate to *interleave*, then we get a 'universal' circuit group that goes beyond classical computation and delivers the famous exponential speedups observed in some quantum algorithms. This phenomenon is discussed in the next section.

5 Universal Quantum Bases

Since any constructive set of operations is going to be finite or countably infinite, we need a different notion of universality and a different concept of circuit synthesis. Both are based on the notion of a *dense subgroup* of the unitary group $U(N)$.

Definition 5. *A subgroup* $G \subset U(N)$ *is everywhere* dense *if for every* $u \in U(N)$ *and for every* $\epsilon > 0$ *there exists a* $g_\epsilon \in G$ *such that the distance between* g_ϵ *and* u *is less than* ϵ*.*

There are several ways to define the distance on an operator group, for the purposes of this definition, they are equivalent. The distance most often used in quantum computing literature is the *trace distance* and is defined on $U(N)$ as

$$dist(U, V) = \sqrt{(N - |tr(U\,V^{\dagger})|)/N},$$

where tr stands for the operator trace. Since $tr(I_N) = N$, each operator is at zero distance with itself. It is not difficult to prove that the distance as defined is non-negative real and satisfies the triangle inequality.

Concept: A finite set of quantum gates forms a *pure universal quantum basis* in n-qubit space if they generate an everywhere dense subgroup of $U(2^n)$.

The best fault-tolerant implementations of operations in a pure universal quantum basis are not entirely unitary; they also require the use of non-unitary operations, including (but not limited to) *state preparation*, *measurement*, and *classical feedback*. These non-unitary operations may use varying numbers of *ancillary qubits*. (To get some taste of the amount of 'non-unitary help' required, an inquisitive reader may consult [33] or the appendix of [35].) As in the reversible logic world, allowing ancillary qubits may be more desirable than increasing the number of operations required for implementation, which relates to the following definition.

Definition 6. *We say that a k-qubit unitary $u \in U(2^k)$ is approximated to precision $\epsilon > 0$ by a circuit $c \in U(2^n)$, where $n \geq k$,* using $n - k$ ancillary qubits *if either $I_{n-k} \otimes u$ is at a distance less than ϵ from c or u is at a distance less than ϵ from a projection of c onto $U(2^k)$.*

In this definition, the term *projection* refers to a factorization map $U(2^n) \rightarrow U(2^k)$ related to some non-unitary operation(s).

We are finally ready to discuss universal quantum bases, which enable quantum computations that cannot be simulated classically. Exact and effective unitary reduction leads to the following result first published in [25]:

Theorem 5. *The circuit group generated by CNOT and all single-qubit unitary operators is purely universal in the multi-qubit space.*

This particular reduction puts an onus on implementing any single-qubit gate G that is universal in single-qubit space, possibly in combination with the single-qubit Clifford group. In fact, *any gate G has this property, unless the eigenvalues of G^2 are* ± 1.

In light of this it would seem that the gate $T = [[1, 0], [0, \sqrt{i}]]$ is the simplest and most logical choice, since the phase gate $S = T^2$ has one eigenvalue equal to i. The gate T, commonly known as the $\pi/8$-*gate*, was originally proposed in

[26] (albeit with a different rationale). While $\{H, S\}$ generate the finite single-qubit Clifford group, $\{H, T\}$ is a universal single-qubit basis and hence generates an infinite group, everywhere dense in $U(2)$, that, of course, contains $S = T^2$ and thus contains the entire Clifford group. Research based on this 'Clifford+T' basis has generated a steady stream of both theoretical and practical results over recent years (an incomplete selection includes [27, 28, 29, 30, 31]).

In fact, the use of this basis is so common that the research community has focused on developing fault-tolerant implementations of the T gate, while perhaps overlooking other more convenient universal bases. The best fault-tolerant implementations of T are based on the so-called *magic state distillation protocol* that consumes a number of ancillary qubits while also requiring non-unitary steps such as measurement and classical feedback (see, for example, [32, 33, 34]). Nevertheless, the T gate provides a convenient abstraction, where the non-unitary techniques needed for implementation are separate from the group-theoretical guarantees.

To give a taste of alternative universal bases, we will briefly sketch a more recent proposal, borrowing directly from reversible logic circuits. The alternative is based on the quantum Toffoli gate $\Lambda^2(X)$ and Subsection 4.1.1.

In 2002, Shi [9] offered an elegant proof that the Toffoli gate in combination with the Hadamard gate form a universal quantum basis when *one* ancillary qubit is allowed. However, the proof does not yield a constructive algorithm to perform the actual approximation of a unitary gate by a synthesized Toffoli/Hadamard circuit to a desired precision.

In contrast, the task of synthesizing Clifford+Toffoli circuits has recently become algorithmic. In [35], an algorithm for synthesizing efficient Clifford+V circuits, where $V = [[1, 0], [0, \frac{3+4i}{5}]]$ is the gate constructed in subsection 4.1.1, is presented. (Action of the T and V gates on the Bloch sphere are shown schematically in Fig. 2.) It shows that any single-qubit unitary can be effectively approximated to precision ϵ by a Clifford+V circuit containing no more than $4 \log_5(2/\epsilon)$ occurrences of the V gate.

By iterating over the circuit from 4.1.1, we can perform the V gate with probability 1. The actual number of iterations needed is a random variable, however its expected value is $5/8 * \sum_k k\,(3/8)^{k-1} = 8/5$. Thus a Toffoli-based circuit approximating a single-qubit target to precision ϵ will have on average $(64/5) \log_5(2/\epsilon)$ occurrences of the Toffoli gate.

As per Theorem 5, the ability to effectively approximate any single-qubit gate with a Toffoli-based circuit, combined with the two-qubit Clifford gate $CNOT$, implies the ability to effectively approximate any multi-qubit unitary by such a circuit. Note that we also have a specific upper bound on the number of occurrences of the Toffoli gate in the resulting approximation.

As defined, this solution currently consumes more resources than the most

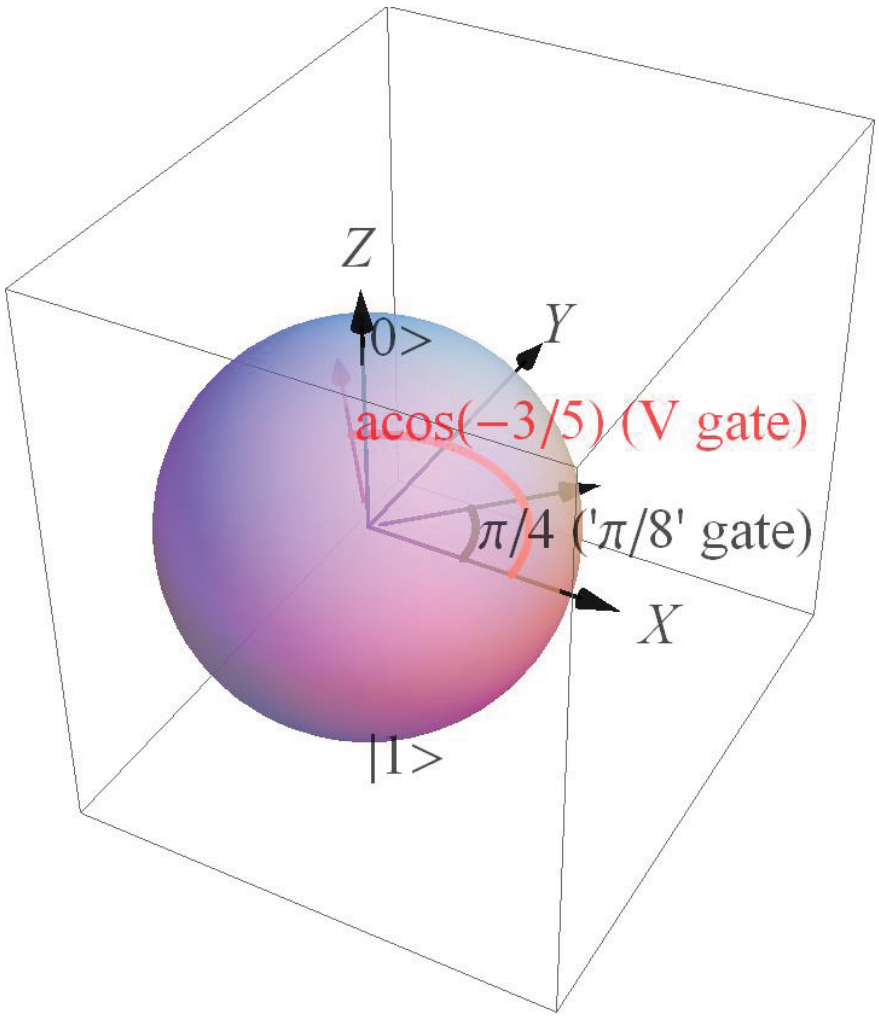

Figure 2: Action of the '$\pi/8$' and V gates on the Bloch sphere. The gates perform rotations about Z axis by the angles of $\pi/4$ and $cos^{-1}(-3/5)$ respectively (the latter angle is an irrational multiple of π).

recent solutions based on the Clifford+T basis (primarily because known fault-tolerant implementations of the Toffoli gate are even more expensive than those of the T gate [41, 42]). However we point out this alternative not just to prove the algorithmic feasibility of the universal Toffoli-based quantum circuits. There is evidence that in the multi-qubit space such circuits can be more aggressively optimized than those based on the Clifford+T approach.

6 Future Directions

At this stage of research, circuits based on universal quantum bases constitute the most popular framework for implementing quantum algorithms. The implementation of an algorithm begins with a definition of the required high-level unitary

and non-unitary operators, followed by a *quantum compilation* step, where the high-level operators are represented by circuits in a chosen basis.

Interestingly, the year 2012 could likely be referred to as the "Year of Quantum Circuit Decomposition". Until early 2012, the most popular and most efficient method for decomposing (or compiling) a high-level quantum circuit, in particular the single-qubit gates, into implementable and fault-tolerant quantum gates was the Dawson-Nielsen version of the Solovay-Kitaev theorem [36, 37]. Given a target unitary gate and a compilation precision ϵ, this method delivers circuits of depth $O(\log^{3.97}(1/\epsilon))$. A handful of theoretical results published over the last decade (c.f., [39]), however, suggested that much more efficient circuit depths $O(\log(1/\epsilon))$ could be achieved for some bases, but no constructive compilation algorithms to achieve these asymptotics were yet known.

Remarkably, in the course of 2012, *efficient* circuit compilation algorithms achieving circuit depths of $O(\log(1/\epsilon))$ have been discovered for two universal bases. Compared to the previous solution in [36, 37], the cost of a circuit implementing a typical single-qubit rotation in an algorithm such as Shor's factorization [10] has come down from millions of basis gates to mere dozens of gates. The latest compilation algorithms (see, for example, [35, 31, 29]) not only address the asymptotic circuit growth rate, but also come with specific upper bounds on the circuit depth.

We now look to an upcoming year to label as the "Year of Multi-qubit Decomposition". Depth upper bounds for multi-qubit circuits is the next research frontier for circuit compilation. Most of the algorithms referenced in this column exploit the following two facts:

1. any single-qubit unitary can be decomposed, effectively and exactly, into at most three axial rotations,

2. any controlled single-qubit unitary can be decomposed, effectively and exactly, into at most three uncontrolled single-qubit unitaries (interleaved with at most two CNOTs).

Sidestepping either of these two intermediate decomposition steps would slash the depth of a circuit implementing a general controlled unitary by a factor of 3 (bypassing both has the potential of reducing the constant coefficient in front of the $\log(1/\epsilon)$ by a factor of 9). In this respect, using multi-qubit primitive gates (such as Toffoli) hold much promise for the future of practical compilation of quantum algorithms.

7 Acknowledgments

The Authors wish to thank Andreas Blass, Yuri Gurevich, and Nathan Cody Jones for very useful discussions.

References

[1] A. Lubotsky, R. Phillips, P. Sarnak. Hecke operators and distributing points on S^2, I and II. *Comm. Pure and Appl. Math., 34, 149–186, and 40, 401–420*, 1986,1987.

[2] R. Landauer. Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development, 5, 183–191*, 1961.

[3] G. E. Moore. Cramming more components onto integrated circuits. *Electronics Magazine, p.4*, 1965.

[4] R. Kurzweil. The Singularity is Near. *Penguin Books*, 2005.

[5] C. H. Bennett. Logical reversibility of computation. *IBM Journal of Research and Development, vol. 17, no. 6, 525–532*, 1973.

[6] D. Gottesman. The Heisenberg Representation of Quantum Computers. `arXiv:quant-ph/9807006v1`, 1998. (http://arxiv.org/abs/quant-ph/9807006v1).

[7] S. Aaronson, D. Gottesman. Improved Simulation of Stabilizer Circuits . *Phys. Rev. A, 70, 052328* , 2004. `arXiv:quant-ph/0406196`, 2004.

[8] M. Van den Nest. Classical simulation of quantum computation, the Gottesman-Knill theorem, and slightly beyond. `arXiv:quant-ph/0811.0898`, 2008. (http://arxiv.org/abs/0811.0898).

[9] Y. Shi. Both Toffoli and Controlled-NOT need little help to do universal quantum computation. `arXiv:quant-ph/0205115`, 2002. (http://arxiv.org/abs/quant-ph/0205115).

[10] P. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing, 26(5):1484-1509*, 2005. (also, arXiv:quant-ph/9508027)

[11] S. Hallgren. Polynomial-time quantum algorithms for Pell's equation and the principal ideal problem. *Proceedings of the 34th ACM Symposium on Theory of Computing*, 2002.

[12] S. Hallgren. Fast quantum algorithms for computing the unit group and class group of a number field. *Proceedings of the 37th ACM Symposium on Theory of Computing*, 2005

[13] A. Schmidt, U. Vollmer. Polynomial time quantum algorithm for the computation of the unit group of a number field. *Proceedings of the 37th Symposium on the Theory of Computing, pg. 475-480*, 2005.

[14] W. van Dam, S. Hallgren, L. Ip. Quantum algorithms for some hidden shift problems. *SIAM Journal on Computing, 36(3):763-778*, 2006. (also, arXiv:quant-h/0211140)

[15] A. Harrow, A. Hassidim, S. Lloyd. Quantum algorithm for solving linear systems of equations. *Physical Review Letters 15(103):150502*, 2009. (also, arXiv:0811.3171)

[16] J. Watrous. Quantum algorithms for solvable groups. *Proceedings of the 33rd ACM Symposium on Theory of Computing, pages 60-67*, 2001. (also, arXiv:quant-ph/0011023)

[17] F. Le Gall. An efficient quantum algorithm for some instances of the group isomorphism problem. Proceedings of STACS 2010. (also, arXiv:1001.0608)

[18] M. Freedman, A. Kitaev, Z. Wang. Simulation of topological field theories by quantum computers. *Communications in Mathematical Physics, 227:587-603*, 2002.

[19] M. Freedman, M. Larsen, Z. Wang. A modular functor which is universal for quantum computation. *Comm. Math. Phys. 227(3):605-622*, 2002. (also, arXiv:quant-ph/0001108)

[20] H. Buning, T. Lettmann. Propositional logic: deduction and algorithms. *Cambridge University Press, Ltd*, 1999.

[21] A. Shende, A. Prasad, I. Markov, J. Hayes. Synthesis of Reversible Logic Circuits. *IEEE Trans. CAD 22(6):710 - 722*, 2003.

[22] T. Toffoli. Reversible Computing. newblock *ICALP: 632-644* , 1980

[23] M. Nielsen, I. Chuang. Quantum Computation and Quantum Information. *Cambridge University Press*, 2000.

[24] S. Clark, R. Jozsa, N. Linden. Generalized Clifford groups and simulation of associated quantum circuits. *Quant Inf Comp 8:106 - 126*, 2008.

[25] A. Barenco et al. Elementary gates for quantum computation. *Physical Review A 52(5):3457-3467*, 1995.

[26] P. Boykin et al. A new universal and fault-tolerant quantum basis. *Information Processing Letters, 75(3):101-107*, 2000.

[27] A. Bocharov, K. Svore. A Depth-Optimal Canonical Form for Single-qubit Quantum Circuits. *Physical Review Letters 109:190501*, 2012. (also, arXiv:1206.3223)

[28] V. Kliuchnikov, D. Maslov, M. Mosca. Asymptotically optimal approximation of single qubit unitaries by Clifford and T circuits using a constant number of ancillary qubits. `arXiv:1212.0822`, 2012.

[29] V. Kliuchnikov, D. Maslov, M. Mosca. Practical approximation of single-qubit unitaries by single-qubit quantum Clifford and T circuits. `arXiv:1212.6964`, 2012.

[30] B. Giles, P. Selinger. Exact synthesis of multiqubit Clifford+T circuits. `arXiv:1212.0506`, 2012.

[31] P. Selinger. Efficient Clifford+T approximation of single-qubit operators. arXiv:1212.6253, 2012.

[32] S. Bravyi, A. Kitaev. Universal Quantum Computation with ideal Clifford gates and noisy ancillas. *Physical Review A 71:022316*, 2005. also, arXiv:quant-ph/0403025, 2004.

[33] B. Reichardt. Improved magic states distillation for quantum universality. http://arxiv.org/pdf/quant-ph/0411036.pdf, 2004.

[34] A. Meier, B. Eastin, E. Knill. Magic-state distillation with the four-qubit code. arXiv:1204.4221, 2012.

[35] A. Bocharov, Y. Gurevich, K. Svore. Efficient Decomposition of Single-Qubit Gates into *V* Basis Circuits. arXiv:1303.1411, 2013.

[36] C. Dawson, M. Nielsen. The Solovay-Kitaev Algorithm. arXiv:quant-ph/0505030, 2005. (http://arxiv.org/abs/quant-ph/0505030).

[37] A. Y. Kitaev. Quantum computations: algorithms and error correction. *Russ. Math. Surv., 52(6):1191-1249*, 1997.

[38] F. Bloch. Nuclear induction. Phys. Rev. 70(7-8) (460), 1946.

[39] A. Harrow, B. Recht, I. Chuang. Efficient Discrete Approximations of Quantum Gates. *J. Math. Phys. 43:4445*, 2002. (also, arXiv:quant-ph/0111031), 2001

[40] B. Eastin. Distilling one-qubit magic states into Toffoli states . arXiv:1212.4872, 2012

[41] Cody Jones. Novel constructions for the fault-tolerant Toffoli gate. *Phys. Rev. A 87:022328* , 2013. (also, arXiv:1212.5069), 2013

[42] Cody Jones. Composite Toffoli gate with two-round error detection. arXiv:1303.6971, 2013

[43] J. Conway. A Course in Functional Analysis. Springer, 1990.

The H-index can be Easily Manipulated

Bart de Keijzer [*] Krzysztof R. Apt [†]

Abstract

We prove two complexity results about the H-index concerned with the Google scholar *merge* operation on one's scientific articles. The results show that, although it is hard to merge one's articles in an optimal way, it is easy to merge them in such a way that one's H-index increases. This suggests the need for an alternative scientific performance measure that is resistant to this type of manipulation.

1 Introduction

The *H-index* was introduced by the physicist J.E. Hirsch in [3] to 'quantify an individual's scientific research output'. Recall that it is defined as the largest x such that one's x most cited paper is cited at least x times. (An aside: Hirsch's original definition was ambiguous as pointed out in [4], where the current definition is proposed.) Its introduction led to an impressive literature. According to Google scholar; by 18th of April 2013 this paper was cited 3043 times. To mention just one example, [5] provided its axiomatic definition.

The H-index started to be used as a universal measure to assess and compare researchers in a given discipline. Hirsch suggested in his paper '(with large error bars) that for faculty at major research universities, $h \approx 12$ might be a typical value for advancement to tenure (associate professor) and that $h \approx 18$ might be a typical value for advancement to full professor'.

In fact, computer scientists seem to cite each other much more often. Jens Palsberg maintains at `http://www.cs.ucla.edu/~palsberg/h-number.html` a list of computer scientists with H-index 40 or higher (a value corresponding in Hirsch's article to Nobel prize winners). The list has more than 600 names and is based on the output generated by Google scholar.

Several people made obvious observations that the H-index can be boosted by such simple measures as adding your name to the articles written by members of

[*]Centre for Mathematics and Computer Science (CWI), `keijzer@cwi.nl`

[†]Centre for Mathematics and Computer Science (CWI) and ILLC, University of Amsterdam, The Netherlands, `apt@cwi.nl`

Bulletin of the EATCS no 110, pp. 79–85, June 2013

your group, splitting a long article into a couple of shorter ones, by citing one's and each other's work, etc. For example, [1] studies the problem of manipulability of the H-index by means of self-citations.

This brings us to the subject of this note. *Google scholar* allows one to perform some operations on the listed articles; notably, the *merge*-operation allows one to combine two versions of an article even if they have different titles. By means of the merge operation, you can obviously improve your H-index. Suppose for instance that your H-index is 20. Then you can increase it by merging two articles that are cited each 11 times.

This suggests two natural problems, where in each case we refer to the improvement of the H-index by means of the merge operation.

- Is it possible to improve your H-index?
- Given a number k, determine whether your H-index can be improved to at least k.

2 Two results

To deal with these questions, we introduce first some notation. A researcher's output is represented as a multiset of natural numbers, each number representing a publication and its value representing the number of its citations. For example the multiset $\{1, 1, 2, 3, 4, 4, 5, 5, 5\}$ represents an output consisting of 9 publications with the corresponding H-index 4. Given a multiset T of numbers we abbreviate $\sum_{x \in T} x$ to $\sum T$. So $\sum T$ is the number of citations resulting from the merge of the publications in T into one.

To deal with the outcomes of merges we need to consider partitions of such multisets.

Fix a finite multiset S of numbers from $\mathbb{N}_{>0}$. We denote by $\bar{S}$ the singletons partition $\{\{x\} \mid x \in S\}$. Given a partition $\mathcal{T}$ of S, we define

$$v(\mathcal{T}) = \max\{|\mathcal{T}'| \mid \mathcal{T}' \subseteq \mathcal{T}, \forall T \in \mathcal{T}' : \sum T \geq |\mathcal{T}'|\},$$

where, as usual, $|\mathcal{T}'|$ denotes the cardinality of the multiset $\mathcal{T}'$ (which is a submultiset of a partition of S in this case). In words, call a subset $\mathcal{T}'$ of the partition $\mathcal{T}$ *good* if each element T of $\mathcal{T}'$ after merge into a single publication yields at least $|\mathcal{T}'|$ citations. So if one allows the merge operation, then a good partition $\mathcal{T}'$ ensures that the H-index can be set to at least $|\mathcal{T}'|$. Then $v(\mathcal{T})$ is the cardinality of the largest good subset of $\mathcal{T}$, hence $v(\mathcal{T})$ is the largest H-index one can obtain by means of the merge operation, while $v(\bar{S})$ is the H-index corresponding to the input multiset S. To put it more directly,

$$v(\bar{S}) = \max\{|T| \mid T \subseteq S,\ \forall x \in T\ x \geq |T|\},$$

where we refer to the submultisets.

We call a partition $\mathcal{S}$ of S an *improving partition* if $v(\mathcal{S}) > v(\bar{S})$. We can now formalize the above two problems as follows, given as input a finite multiset S of numbers in $\mathbb{N}_{>0}$.

H-index improvement problem Does there exist an improving partition? If yes, find it.

H-index achievability problem Given a number k, does there exist a partition $\mathcal{T}$ of S, such that $v(\mathcal{T}) \geq k$?

In Section 3, we present the proofs of the following two results.

Theorem 1. *The H-index improvement problem can be solved in polynomial time.*

Theorem 2. *The H-index achievability problem is strongly* NP*-complete.*[1]

In particular, it is strongly NP-hard to compute the maximal H-index that can be achieved through the merge operation.

From the viewpoint of manipulability, Theorem 1 is bad news. Ideally, we would like to have a performance measure that is computationally difficult to manipulate. One can see a parallel with the search for voting methods that are difficult to manipulate, see, e.g. [6]. Our conclusion is that the H-index is not the last word in the ongoing quest to find a credible way to quantify one's scientific output.

3 Proofs of the theorems

In what follows, we assume that a multiset is represented as a list of possibly duplicate numbers. A different way of representing a multiset would be the more compact one, where we list only the distinct numbers that appear in the multiset, along with their respective multiplicity. We consider the latter representation to be unnatural, given the context in which we study this problem.

Proof of Theorem 1. Let S be the given multiset. Let S' be the smallest submultiset of S such that $v(\bar{S}) = v(\overline{S'})$. For instance, if $S = \{5, 4, 3, 3, 3, 2\}$, then $S' = \{5, 4, 3\}$ and if $S = \{5, 3, 3, 3, 3, 2\}$, then $S' = \{5, 3, 3\}$. In both cases $v(\bar{S}) = 3$. Call a number $x \in S'$ *supercritical* if $x > v(\bar{S})$ and *critical* if $x = v(\bar{S})$. Let C_+

[1] A decision problem that involves numerical input is said to be *strongly* NP-complete if the problem is NP complete even if all the numbers in the input are represented in unary.

be the multiset of all supercritical numbers in S' and C the multiset of all critical numbers in S'. Note that C and C_+ partition S' and that $v(\bar{S}) = |C_+| + |C|$. Furthermore, let L denote the multiset of $|C|$ smallest numbers in S.

For instance, if $S = \{5, 4, 3, 3, 3, 2\}$, then $C = \{3\}$ and $L = \{2\}$, and if $S = \{5, 3, 3, 3, 3, 2\}$, then $C = \{3, 3\}$ and $L = \{3, 2\}$.

Note that below, we treat duplicate numbers in S as having "separate identities", so that for two numbers $x, y \in S$ that are equal in magnitude, it may hold that $x \in C$ but $y \notin C$ or $x \in L$ but $y \notin L$. We believe that this slight informality and definitional abuse will cause no confusion to the reader.

We first establish the following characterization result.

Lemma 1. *There exists an improving partition of S iff $L \cap C = \varnothing$ and $\sum S \setminus (C \cup C_+ \cup L) > |C| + |C_+|$.*

Proof. Suppose there exists an improving partition $\mathcal{S}$ of S.

We can assume without loss of generality that the following properties then hold:

1. Each supercritical number in S appears in a singleton set in $\mathcal{S}$. These are the only singleton sets in $\mathcal{S}$.

   Indeed, if a supercritical number $x \in S$ appears in a non-singleton set $T \in \mathcal{S}$, then take the partition $\mathcal{T}$ of S obtained from $\mathcal{S}$ by splitting T into singletons. Because $\mathcal{S}$ is an improving partition, there are at least $v(\bar{S})$ multisets $T' \in \mathcal{S}\setminus\{T\}$ such that $\sum T' > v(\bar{S})$. All multisets of $\mathcal{S}\setminus\{T\}$ are in $\mathcal{T}$. Also the number x is in a singleton set of $\mathcal{T}$ and $x > v(\bar{S})$. Therefore, there are in $\mathcal{T}$ at least $v(\bar{S}) + 1$ multisets T' such that $\sum T' > v(\bar{S})$. Hence, $\mathcal{T}$ is an improving partition.

   After we have repeatedly performed the above splitting steps we obtain an improving partition $\mathcal{S}'$ such that each supercritical number $x \in S$ appears in a singleton set in $\mathcal{S}'$.

   Since
   $$v(\mathcal{S}') > v(\bar{S}) = |C_+| + |C| \geq |C_+|,$$
   there exists in $\mathcal{S}'$ a non-singleton multiset $T \in \mathcal{S}$ that contains only non-supercritical numbers. Merging with it all singleton sets that contain a non-supercritical number yields the desired improving partition.

2. L is disjoint from C.

   By Property 1, the supercritical numbers form singleton sets in $\mathcal{S}$, and each remaining multiset has cardinality at least 2. If L were not disjoint from C, then we would have $|S| \leq |C_+| + |L| + |C|$, so $|S \setminus C_+| \leq |L| + |C| = 2|C|$,

hence the number ℓ of non-singleton multisets in $\mathcal{S}$ would be at most $|C|$. This yields a contradiction, since we would then have $v(\mathcal{S}) \leq |C_+| + \ell \leq |C_+| + |C| = v(\bar{S})$.

3. In $\mathcal{S}$, every critical number is in a set of cardinality 2.

   Indeed, by Property 1, critical numbers do not appear in singleton sets. Further, if a critical number $x \in S$ appears in a multiset $T \in \mathcal{S}$ of cardinality exceeding 2, then we can split T in any way so that x is put in a multiset T' of cardinality 2. It then holds that $\sum T' > v(\bar{S})$, so the resulting partition remains an improving partition.

4. There is a bijection $\pi : C \to L$ such that $\{x, \pi(x)\} \in \mathcal{S}$ (i.e., C is "matched" with L in $\mathcal{S}$).

   Indeed, by Property 3, every critical number is in a set of cardinality 2. Now, let x be a critical number and let $\{x, y\} \in \mathcal{S}$ be the multiset of cardinality 2 that contains x. If y is not in L, then $|C| = |L|$ implies that there is a number $y' \in L$ that occurs in a multiset T in $\mathcal{S}$ that does not contain a critical number. Because $y' \leq y$, the operation of swapping y' and y in $\mathcal{S}$ does not decrease the number of multisets that sum to at least $v(\bar{S}) + 1$. So the partition that results after this swap remains an improving partition.

We have $v(\mathcal{S}) > v(\bar{S}) = |C_+| + |C|$, so by Properties 1,2, and 4, there is a multiset $T \in \mathcal{S}$ not intersecting C_+, C, and L, such that $\sum T > v(\bar{S})$. Hence $\sum S \setminus (C \cup C_+ \cup L) \geq \sum T > v(\bar{S}) = |C| + |C_+|$. We conclude that if there is an improving partition, then $L \cap C = \emptyset$ and $\sum S \setminus (C \cup C_+ \cup L) > |C| + |C_+|$.

Conversely, if $L \cap C = \emptyset$ and $\sum S \setminus (C \cup C_+ \cup L) > |C| + |C_+|$, then there is an improving partition. It consists of

- the singletons, each containing an element of C_+,
- the sets of cardinality 2, each containing a pair of elements from C and L,
- the multiset $S \setminus (C \cup C_+ \cup L)$.

□

The proof of Theorem 1 is now immediate. It is straightforward to compute C_+, C and L in polynomial time. Using the above lemma we can therefore determine in polynomial time whether an improving partition exists, and find one in polynomial time if it does. □

Proof of Theorem 2. The problem is clearly in NP, so the proof will focus on establishing NP-hardness. We do this by means of a polynomial time reduction

from a strongly NP-complete problem. The reduction is from the 3-PARTITION problem. In the 3-PARTITION problem, we are given a multiset M of $3m$ positive integers, such that $\sum M = mb$ for some $b \in \mathbb{N}$. We have to decide whether it is possible to partition this set into m submultisets, such that the sum of the numbers in each submultiset is exactly b.

Garey and Johnson [2] prove that the 3-PARTITION problem is strongly NP-complete, even under the assumption that M is represented as above (i.e., non-concisely). This means that the 3-PARTITION problem is NP-complete even when b is bounded by some polynomial in m. Denote this polynomial by $p(m)$. From now on, with the SPECIAL 3-PARTITION problem we will mean the special case of the problem where b is bounded by $p(m)$.

Before proceeding, one note is in order. In the original definition of the 3-PARTITION problem, the additional requirement is imposed that all sets in the partition are of cardinality 3 (and this is also where the name of the problem originates from). For convenience, we do not impose this requirement here. The reason it is not necessary to impose this requirement is because in [2], it is shown that strong NP-hardness holds even when all numbers in the multiset are strictly between $b/2$ and $b/4$. This enforces that all sets in the partition will be of cardinality 3. Without the cardinality constraint, the problem thus becomes more general, and is automatically strongly NP-hard.

Given a SPECIAL 3-PARTITION instance (S', m, b), we reduce it to an H-index manipulation problem instance (S, k) as follows. First, obtain S'' from S' by adding m to each number in S'. Note that (S'', m, k), where $k = b + 3m$, is a YES-instance of 3-PARTITION if and only if (S', m, b) is a YES-instance of SPECIAL 3-PARTITION. Note also that $k - m = b + 2m > 0$. Next, obtain the multiset S from S'' by adding $k - m$ copies of k to S'. This takes polynomial time, as k is bounded by $p(m) + 3m$.

We now show that (S, k) is a YES-instance of the H-index manipulation problem if and only if (S'', m, k) is a YES-instance of 3-PARTITION.

If (S'', m, k) is a YES-instance of 3-PARTITION, then let $\mathcal{T}$ be a certificate for that, so $\mathcal{T}$ is a partition of S'' into m multisets such that the sum of the numbers in each multiset is k. Then by adding to $\mathcal{T}$ exactly $k - m$ copies of the set $\{k\}$, we obtain a certificate that (S, k) is a YES-instance of the H-index achievability problem, because $k = k$.

Conversely, if (S, k) is a YES-instance of the H-index achievability problem, then let $\mathcal{T}$ be a certificate for that. We can assume without loss of generality that the partition $\mathcal{T}$ contains exactly $k - m$ copies of the set $\{k\}$. Indeed, otherwise we can split each non-singleton set in $\mathcal{T}$ that contains a copy of k into singleton sets. This will result in a desired certificate.

By removing all singleton sets $\{k\}$ from $\mathcal{T}$ we obtain a partition $\mathcal{T}'$ of S''. By the choice of (S, k) this new partition $\mathcal{T}'$ contains m multisets, each of which sums

up to k. $\mathcal{T}$ does not contain any additional multiset besides these m multisets, as then we would have $\sum S'' > mk$, which is not the case by construction. Therefore, $\mathcal{T}'$ is a certificate that (S'', m, k) is a YES-instance of 3-PARTITION. □

References

[1] C. Bartneck and S. Kokkelmans. Detecting h-index manipulation through self-citation analysis. *Scientometrics*, 87(1):85–98, 2011.

[2] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.

[3] J. E. Hirsch. An index to quantify an individual's scientific research output. *Proceedings of the National Academy of Sciences of the United States of America*, 102(46):16569–16572, 2005.

[4] F. Steutel. Persoonlijke ervaringen met de Hirsch-index. *Nieuw Archief voor Wiskunde*, 5:194, 2008. In Dutch.

[5] G. J. Woeginger. An axiomatic characterization of the Hirsch-index. *Mathematical Social Sciences*, 56:224–232, 2008.

[6] M. Zuckerman, P. Faliszewski, Y. Bachrach, and E. Elkind. Manipulating the quota in weighted voting games. *Artif. Intell.*, 180-181:1–19, 2012.

Report on StringMasters 2013

20-23 February 2013, Verona, Italy

Zsuzsanna Lipták

The **11th edition of the StringMasters workshop** took place from 20 to 23 February 2013 in the beautiful northern Italian city of Verona. In total, 27 participants from 11 countries came to Verona for this workshop, in order to spend a few days together working intensively on open problems in stringology. The website of the workshop is:

`http://stringmasters.di.univr.it`.

StringMasters has taken place since 2007, once to three times per year. Previous venues were Hamilton (Canada), Stellenbosch (South Africa), Perth (Australia), London (UK), Rouen (France), and Palermo (Italy). The workshop aims at bringing together researchers of different levels (senior, junior, graduate students), reinforcing existing collaborations, and forming new ones. Topics include all areas of stringology, spanning **algorithms on strings, combinatorics on words, and applications of strings (such as in bioinformatics, musical analysis, internet issues, data mining, natural language processing, and others)**.

The format emphasizes free collaboration. There is no formal program; the workshop starts with the presentation of open problems, and in the following participants are encouraged to form small groups, in which to attack a chosen problem. There are several rooms available for working, and participants are free to

Bulletin of the EATCS no 110, pp. 89–91, June 2013

switch between groups at any time. If groups have found their collaboration fruitful, they continue after the end of workshop via email and other means of communication. Special issues of high level journals are regularly devoted to results stemming from these workshops. StringMasters' steering committee consists of **Maxime Crochemore, Costas Iliopoulos**, and **William F. Smyth**. See also the central StringMasters website:

`www.stringmasters.org`.

This time the workshop kicked off with a short introductory round, and then 11 open problems were presented by participants. Most of these elicited a lot of interest. Of course, scheduling time slots for problems optimally, according to participants' expressed interest, under room and time constraints, is an NP-hard problem. So after a first tentative schedule offered by the organizers (with no approximation guarantee), from day 2, the workshop's credo was: "Self-organize!" And indeed, participants succeeded in forming groups according to problems; during the three and a half days of the workshop, most if not all of the initial problems were attacked (and some solved). These included problems on jumbled pattern matching, abelian factors, periods and squares and runs in words, indeterminate words, compression, and several others.

The workshop took place at the **Department of Computer Science, University of Verona**, and was organised by **Giuditta Franco** and **Zsuzsanna Lipták**, of the Verona Bioinformatics group. Four rooms were at free disposal of the workshop, in one of which coffee, cookies, and fruit were served in the late mornings. The department lies in the southern part of town, an easy bus ride away from the medieval centre, where most participants lodged. The social dinner was held in a desanctified church in the old part of town. Verona is arguably one of the most beautiful towns in Italy, with a Roman amphitheatre (the Arena), splendid churches and towers, medieval bridges giving breathtaking views of the river Adige—which forms a loop (some say: heart) around the old town. On Saturday morning there was a museum visit with a guided tour organised for participants (thanks to Travis Gagie).

Among the 27 participants, there were researchers of all levels: senior, junior (postdoc), as well as 5 PhD students and 1 masters student. The countries of origin of the participants were:

> Australia (1), Canada (1), Czech Republic (2), Finland (1),
> France (2), Germany (2), Hungary (1), Italy (6), Poland (2),
> South Africa (1), UK (8).

The organisers wish to thank the Department of Computer Science, University of Verona, for the financial and logistic support of StringMasters 2013.

There will be a **special issue of the Journal of Discrete Algorithms** devoted to results stemming from the StringMasters workshops in 2012 and 2013, edited by Maxime Crochemore, Jacqueline Daykin, and Zsuzsanna Lipták.

The next StringMasters will take place in **Prague from Sept 5-7, 2013**, right after the Prague Stringology Conference (PSC 2013).

Abstract of PhD Thesis

Author: Jose Gaintzarain

Title: Invariant-Free Deduction Systems for Temporal Logic

Language: English

Supervisor: Paqui Lucio

Institute: University of the Basque Country, Spain

Date: 13 July 2012

Abstract

In this thesis we propose a new approach to deduction methods for temporal logic. Our proposal is based on a non-customary inductive definition for eventualities, that allows us to provide dual systems of tableaux and sequents for Propositional Linear-time Temporal Logic (PLTL). Then, we extend our approach to the resolution framework and we present a clausal temporal resolution method for PLTL. Finally, we make use of this new resolution method for establishing logical foundations for declarative temporal logic programming languages.

The key element in the deduction systems for temporal logic is to deal with eventualities and "hidden" invariants that may prevent the fulfillment of eventualities. Traditional tableau systems for PLTL generate an auxiliary graph in a first pass. Then, in a second pass, the fulfillment of eventualities is checked and unsatisfiable nodes are pruned. The one-pass tableau introduced by S. Schwendimann requires an additional handling of information in order to detect cyclic branches that contain unfulfilled eventualities. In traditional sequent calculi for PLTL, the issue of eventualities and hidden invariants is mainly tackled by using invariant-based rules or infinitary rules. A remarkable consequence of using the above mentioned approaches in the tableau and sequent frameworks, is that temporal logic fails to carry out the classical correspondence between tableaux and sequents. In this thesis, we first provide the one-pass tableau method TTM that instead of a graph obtains a cyclic tree to decide whether a set of formulas is satisfiable. In TTM tableaux are classical-like. For unsatisfiable sets of formulas, TTM produces tableaux whose leaves contain a formula and its negation. In the case of satisfiable sets of formulas, TTM builds tableaux where each fully expanded open branch characterizes a collection of models for the set of formulas in the root. The tableau method TTM is complete and yields a decision procedure for PLTL. This tableau method is directly associated to the one-sided sequent calculus TTC and also to the two-sided sequent calculus GTC. Likewise TTM, the two calculi TTC and GTC

are sound and complete and also are free from all the structural rules that hinder the mechanization of deduction, e.g. weakening, contraction and cut (including invariant-based cut). Therefore, we show that the classical correspondence between tableaux and sequent calculi can be extended to temporal logic.

The most fruitful approach in the literature on resolution methods for temporal logic, due to M. Fisher, deals with PLTL and requires to generate invariants for performing resolution on eventualities. In this thesis, we present a new approach to resolution for PLTL. The main novelty of our approach is that we do not generate invariants for performing resolution on eventualities. Our temporal resolution method TRS is based on the methods of tableaux and sequents mentioned above and involves an effective translation of any formula into its clausal normal form. TRS is sound and complete. This method is also terminating, hence it gives rise to a new decision procedure for PLTL.

Finally, we present the declarative propositional temporal logic programming language TeDiLog that is a combination of the temporal and disjunctive paradigms in Logic Programming. We formally define operational and logical semantics for TeDiLog and prove their equivalence. The operational semantics of TeDiLog is based on TRS. TeDiLog is very expressive, in particular, it allows both eventualities and always-formulas to occur in clause heads and also in clause bodies.

Since the tableau method presented in this thesis is able to detect that the fulfillment of an eventuality is prevented by a hidden invariant without checking for it by means of an extra process, since our finitary sequent calculi do not include invariant-based rules and since our resolution method dispenses with invariant generation, we say that our deduction methods are invariant-free.

Table of Contents

Author's correspondence address Jose Gaintzarain
Lenguajes y Sistemas Informáticos
EUITI de Bilbao (UPV/EHU)
Paseo Rafael Moreno Pitxitxi, 3
48013-Bilbao
Spain
Email: *jose.gaintzarain@ehu.es*

Abstract of PhD Thesis

Author: Sergi Oliva

Title: On the Complexity of Resolution-based Proof Systems
Language: English

Supervisor: Albert Atserias
Institute: Universitat Politècnica de Catalunya – BarcelonaTech

Date: 2 May 2013

Abstract

Propositional Proof Complexity is the area of Computational Complexity that studies the length of proofs in propositional logic. One of its main questions is to determine which particular propositional formulas have short proofs in a given propositional proof system. In this thesis we present several results that answer a particular case of this question or are intimately related to it, all on proof systems that are extensions of the well-known resolution proof system.

The first result of this thesis is that TQBF, the problem of determining if a fully-quantified propositional CNF-formula is true, is PSPACE-complete even when restricted to instances of bounded tree-width, i.e. a parameter of structures that measures their similarity to a tree. Instances of bounded tree-width of many NP-complete problems are tractable, e.g. SAT, the boolean satisfiability problem. We show that this does not scale up to TQBF. We also consider Q-resolution, a quantifier-aware version of resolution. On the negative side, our first result implies that, unless NP = PSPACE, the class of fully-quantified CNF-formulas of bounded tree-width does not have short proofs in any proof system (and in particular in Q-resolution). On the positive side, we show that instances with bounded respectful tree-width, a more restrictive condition, do have short proofs in Q-resolution. We also give a natural family of formulas with this property that have real-world applications.

The second result proposes a semantic way to compare first-order principles with respect to the length of the proofs of their propositional translations in a particular propositional proof system. To do that, we use the classical logic concept of interpretability. Informally, we say that a first-order formula can be interpreted in another if the first one can be expressed using the vocabulary of the second, plus some extra features. We show that first-order formulas whose propositional translations have short R(const)-proofs, i.e. a generalized version of resolution with DNF-formulas of constant-size terms, are closed under a weaker form of in-

terpretability (that with no extra features), called definability. Our main result is a similar result on interpretability. Also, we show some examples of interpretations and show a systematic technique to transform some Σ_1-definitions into quantifier-free interpretations.

The third and final result is about a relativized weak pigeonhole principle. This says that if at least $2n$ out of n^2 pigeons decide to fly into n holes, then some hole must be doubly occupied. We prove that the CNF encoding of this principle does not have polynomial-size DNF-refutations, i.e. refutations in the generalized version of resolution with unbounded DNF-formulas. For this proof we discuss the existence of unbalanced low-degree bipartite expanders satisfying a certain robustness condition.

Table of Contents

Author's correspondence address Sergi Oliva
Edifici Ω – Office S108
Jordi Girona Salgado, 1-3
08034 – Barcelona
Spain
Email: *oliva@lsi.upc.edu*