CURRICULUM VITAE WITH RESEARCH ON FOUR PAGES

Victor Y. Pan

Distinguished Professor
Department of Mathematics and Computer Science
Lehman College - City University of New York
Bronx, New York 10468
E-Mail: victor.pan @ lehman.cuny.edu
Website: http://comet.lehman.cuny.edu/vpan/

RESEARCH AREAS: Symbolic and Numerical Algorithms. General, Structured, Random Matrices. Polynomials. Low Rank Approximation. Parallel and Graph Algorithms in 1985–2000.

EDUCATION: *Ph.D. in Math* from Moscow State University (Advisor: A. G. Vitushkin)

EMPLOYMENT (since 1977): 1988 to present: Professor, then Distinguished Professor (since 2000), Math and Computer Science Depts. of Lehman College and Graduate Center, CUNY; 1979-1991: Professor, Computer Science Dept., SUNYA, Albany, NY.

Visiting Positions: INRIA, France 1996-97; ICSI 1991-1992 and MSRI 1998, both in Berkeley, CA; Columbia Univ., NY 1989-90; Univ. of Pisa, Italy 1984 and 2002; Stanford Univ. 1981; Institute for Advanced Study, Princeton 1979-80; IBM Research Center, Yorktown Heights, NY 1977-80.

GRANTS AND PROFESSIONAL HONORS INCLUDE: NSF Grants (collaborative): CCF-1563942, \$608,205, 2016-2021 and AitF, CCF-133834, \$448,086.00, 2017-2021; (individual): \$1,483,057, 1980-2016; Special Creativity Extension Award from the Numeric, Symbolic, and Geometric Computation Program of the CCR Division in the Directorate CISE of NSF (1993); Best Paper Award, Journal of Complexity (2000); Designation of Fellowship in American Math Society "For Contributions to the Mathematical Theory of Computation" (2013).

Membership in Professional Societies (since 1977): SIAM, ACM, AMS.

Other Professional Activities: Associated Editor of 3 Journals in Applied Math and Computer Science; Managing Guest Editor of 4 Special Issues of the Theoretical Computer Science (TCS) on Symbolic–Numerical Algorithms; Member of 17 Program Committees of International Conferences; Local Arrangements Chair of the ACM SIGSAM ISSAC 2018, Organizer of Minisymposia.

Research reports in arXiv, lectures at the universities, research centers and professional conferences worldwide, dissemination of research results via the Internet and personal communication.

Adviser, Mentor and Defense Committee Chair of 27 defenses of PhD degree in Math and Computer Science. Guiding five PhD students in 2018-19. Extensive joint publications with PhD students, with some for more than a decade after their defense.

RESEARCH HIGHLIGHTS: OVERVIEW AND BY SUBJECTS (CV on my Website lists my conference talks since 1991 and all my papers, including dozens of my papers in SICOMP, SIMAX, SISC, and proceedings of SIAM sponsored conferences; its Sect. 12 details my research).

I solved a number of long-standing problems in Computational Math and Computer Science, introduced new insights and novel methods, revealed unexpected links among some seemingly distant subjects, and proposed new research directions and new areas of study. My research helped creating fields of Algebraic Computational Complexity, Symbolic-Numerical Computing, and Algebraic Multigrid and establishing synergistic links among Symbolic Computations, Numerical Computations, Theoretical Computer Science, and Numerical Linear Algebra.

Some concepts and definitions introduced in my papers, my techniques and insights are widely adopted and commonly used, sometimes as folklore. Google Scholar and DBLP list my four books (1623+LXXIV pages overall), over 20 surveys in journals and book chapters, over 170 research articles in journals, over 90 in refereed conference proceedings, and over 11,000 citations of my work.

My book with Dario Bini "Polynomial and Matrix Computations", Birkhäuser, Boston, 1994, over 700 Google Scholar citations, covers its title subjects in depth and presents a wealth of our new research results. Reviewers in ACM SIGACT News (June 1995) and ACM SIGACT Bulletin (Sept. 1996) found it "outstanding" and recommended it as "must-have book" and "a reference for a *graduate course* in symbolic computations or computer algebra." Polynomial and matrix computations are also the subjects of my three other books, a dozen of my book chapters, and most of my research and survey papers, including my long surveys in SIAM Reviews in 1984, 1992, and 1997.

I. In 1962, while a sophomore, I solved A. M. Ostrowsky's problem of 1955 by proving that Horner's classical polynomial evaluation is optimal. My technique of active operations/linear substitution was extended by Volker Strassen, S. Winograd, and other renowned researchers to proving optimality of classical algorithms for some fundamental matrix computations and has led to emergence of a new field of *Algebraic Computational Complexity*, now popular. I became known to experts as *"polynomial Pan"*. This work and my results on polynomial evaluation with pre-processing have been surveyed in Russian Math Surveys 1966 by myself, in V. Strassen's two chapters (both called *"Pan's method"*) – in Proc. Intern. Congress of Mathematicians, Vancouver, 1974 (R.D. James, ed.), **1**, 497-501, 1974 and "Analytic Complexity of Computations" (J. Traub, ed.), pp. 1-10, Plenum Press, NY 1972, and in D. E. Knuth's "The Art of Computer Programming", v.2, 1981 and 1997.

II. My next research breakthrough in 1978 has ended a stalemate of a nearly decade on fast matrix multiplication (see my papers in IEEE FOCS 1978, SICOMP 1980, and SIREV 1984). Strassen's decrease of the classical cubic arithmetic time $2n^3 - n^2$ for $n \times n$ matrix multiplication to $O(n^{\log_2 7})$ has become a worldwide sensation of 1969. Then, according to D.E. Knuth, "literally all leading experts worldwide worked on" decreasing the exponent $\log_2 7 \approx 2.81$. I was happy to break Strassen's record in 1978, by attacking this 2-dimensional challenge from three dimensions: I first reduced the task to the decomposition of a special trilinear form or equivalently 3-dimensional tensor into a small number of terms and then nontrivially exploited some symmetry in the tensor of matrix multiplication. I called my technique *trilinear aggregation* in 1978 but have actually published it already in an unnoticed paper of 1972 (in Russian)[1], which presented *the first nontrivial acceleration of matrix computations by means of tensor decomposition*, now a thriving area.

With my participation my record exponent was further decreased based on combination of various advanced methods, trilinear aggregation being indispensable ingredient according to D. Coppersmith and S. Winograd 1990, page 255. The new advanced algorithms, however, involved numerous recursive steps, each squaring the input size, and A. Schönhage concluded introduction to his seminal paper in SICOMP 1981 with *"Pan's estimates of 1978 for moderate values of n are still unbeaten"*.

This is still true today, except that in 1982, 1984 and with J. Laderman and X. Sha in 1992 I improved my design of 1978, by relying exclusively on trilinear aggregation. The well-known implementations of some of these algorithms by I. Kaporin in 1999 and 2004 use much less memory and are more stable numerically than the other known fast algorithms, including Strassen's of 1969.

III. In LAA 1980, jointly with W. L. Miranker, I introduced *hierarchical aggregation*. This was a basic step in the emergence of the field of Algebraic Multigrid, now popular.

IV. In 1985-2000, in book chapters and many papers in leading journals and proceedings of respected workshops, by myself and with coauthors, I published new *fast and processor-efficient parallel algorithms* for computations with matrices, polynomials, and graphs. Processor efficiency is critical for practice of computing but was a novelty at that time in Theory of Computing. My answers to some well-known nontrivial technical challenges in this area include nearly optimal parallel polynomial division with Bini in J. Complexity 1986, FOCS 1992 and SICOMP 1993, efficient symbolic solution of linear systems of equations in abstract fields with E. Kaltofen in ACM SPAA 1991 and FOCS 1992 and over integers by myself in TCS 1987 and SICOMP 2000, and a proof of NC equivalence of planar integer programming and integer polynomial GCD with D. Shallcross and Yu Lin in ACM-SIAM SODA 1992, FOCS 1993 and SICOMP 1998. Magazines "Science", "Science

[1]Its first translation into English appeared only in 2014 – in arXiv:1411.1972

News", and "Byte" covered part of my work in 1985–1986.

V. My paper in Math. of Computation 1990 *enhanced the power of using the displacement structure of matrices.* Matrices with Toeplitz, Hankel, Vandermonde and Cauchy displacement structures are omnipresent in modern computations, but form four different matrix classes, each allowing different computational benefits and long studied independently of each other. My transformations of their structures into each other, however, enabled extension of any successful inversion algorithm for matrices of one of these classes to the inversion of the matrices of the other classes, and similarly for solving linear systems of equations.[2] Cauchy-to-Toeplitz reduction at once implied theoretical decrease of the complexity; since 1995 Toeplitz-to-Cauchy reduction has led leading experts to devising efficient numerical algorithms of users' choice for Toeplitz linear systems of equations.

In 2013-2017 by combining my approach with Fast Multipole Method (FMM) I significantly accelerated numerical multipoint polynomial evaluation. As by-product I formally supported long-known empirical observation that Vandermonde matrices are ill-conditioned unless their knots lie on or near the unit circle centered in the origin and are nearly equally spaced on it. (See my papers in LAA 2015, SIMAX 2016, and Math of Computation 2017.)

Efficient *computations with structured matrices* and their applications were major subjects of my books of 1994 (with D. Bini) and 2001, both published by Birkhäuser, my book chapters in 1992, 1999 and 2009, and dozens of my research articles. Besides displacement transformation, subjects of my research in this area include compression of displacement (SIMAX 1993), inversion of displacement operators (SIMAX 2003), acceleration of Newton's iterations for structured matrix inversion by means of displacement re-compression in 1992-2010,[3] and fast solution of structured linear systems of equations by means of lifting and divide and conquer techniques (1996-2017).

VI. In my review in SIAM Review of 1992 (38 pages), my books of 1994 (with Bini) and 2001, and a number of research papers I established and explored links between computations with structured matrices, polynomials, and rational functions. In 1997-2005 in a series of joint papers with B. Mourrain and with I.Z. Emiris in leading journals and workshop proceedings I proposed a number of novel, record fast, and now popular structured matrix methods for solving a *multivariate polynomial system of equations*. With Mourrain I shared the *Best Paper Award of Journal of Complexity 2000.*

VII. Univariate *polynomial root-finding* has been the central problem of Mathematics and Computational Mathematics throughout four millennia and is still important for the theory and practice of computing. My algorithms in STOC 1995, CAMWA 1996 and J. Symbolic Comp. 2002 run in record low Boolean time, which is *optimal up to polylogarithmic factors.* The algorithms approximate all roots almost as fast as one can read input coefficients.

My algorithm in J. of Complexity 2000, extending Weyl 1924, Henrici 1974 and Renegar 1987, was in turn extended into another nearly optimal root-finder by Becker et al. in J. Symbolic Comp. 2018. Already its implementation reported in ICMS 2018 by R. Inbach, myself, and C. Yap and supported by NSF Grant AF: Medium: Collaborative Research: "Numerical and Algebraic Differential Equations", CCF-1563942, \$608,205, 2016-2021 (joint with A. Ovchinnikov and Yap) was fastest known for root-finding in a complex disc with fewer roots, but in CASC 2019 and arxiv:1805.12042 I proposed significant improvements of this algorithm as well as numerical multipoint polynomial evaluation, which is fundamental for polynomial root-finding and for many other subjects of numerical algebraic computations. My new root-finders efficiently handle polynomials defined by subroutines for their evaluation rather than by their coefficients, which is important novel feature for both theory and practice. Implementation of my work is in progress, see arxiv:1906.04920.

My other polynomial root-finders have been well-recognized as well. A few examples are my papers in SICOMP 1995 by myself, in J. of Complexity 1996 and CAMWA 2004 with D. Bini, in ETNA 2004 and Numer. Math. 2005 with him and L. Gemignani, in ISSAC 2013 and J. Symbolic Comp. 2016 jointly with E.P. Tsigaridas, and in Symbolic-Numerical issue of TCS 2017 with my student L. Zhao. This includes nearly optimal real root-finding, nontrivial use of matrix methods,

[2] One operates with these matrices fast by working with their displacements of low rank. The 4 matrix classes differ in 4 ways of defining such displacements, but proper multipliers enable unification of these 4 ways.

[3] Newton's iterations recursively increase the rank of the displacement of a matrix. This slows down the computations, but in J. of Complexity 1992, IEEE Trans. Parallel Distr. Comp. 1993, and a number of subsequent papers I countered this deficiency with my re-compression techniques, which recursively recover low-rank representation.

and technical novelties commonly used afterwards (see section 12 of my CV).

My book of 2013 with J.M. McNamee provides *unique nearly complete coverage of the known polynomial root-finders* up to the date and is indispensable for researchers and algorithm designers in this field. I also surveyed the subject in SIAM Review 1997 (34 pages), the magazine American Scientist 1998, and my book chapters (with coauthors) in 1997, 1999, 2004, 2007, 2009 and 2014.

VIII. *Synergy of symbolic and numerical computations* can be highly beneficial for both of these important subjects, historically developed quite independently of one another. Since the early 1990s I have been promoting such benefits as an organizer of conferences, a member of their Program Committees, and the Managing Editor of four Special Issues of TCS on this subject in 2004, 2008, 2011 and 2013, but even more so with my books of 1994 (joint with D. Bini), 2001 (by myself), and 2013 (joint with J.M. McNamee), my surveys, in particular in SIAM Review 1992 and 1997, my book chapters (with co-authors) in four Handbooks on Algorithms and Computing of 1999, 2004, 2009, and 2014, and dozens of my research papers. E.g., the Special Issue of TCS on Symbolic-Numerical Algorithms in 2017 includes my three joint papers out of 13 papers overall in that Issue.

IX. Since 2010 I published in LAA a number of algorithms for the solution of nonsingular and homogeneous singular linear systems of equations with *randomized preprocessing*. E.g., in LAA 2013, 2015 and 2017 (with my students) I studied randomized multiplicative preprocessing as a means of numerical stabilization of Gaussian elimination instead of communication intensive pivoting. We proved that this produces accurate solution with a high probability (whp) for random inputs and proposed new classes of efficient sparse and structured multipliers.

X. Since 2016 I have extended that study to the hot research subjects of *Low Rank Approximation (LRA)* of a matrix, which is highly important for numerical linear and multilinear algebra and Big Data mining and analysis. Realistically one can only access a tiny fraction of all entries of input matrices representing such data, e.g., unfolding matrices of multidimensional tensors, and so one must devise *sublinear cost algorithms*, which use much fewer flops and memory cells than the input matrix has entries. The known algorithms, however, use superlinear time because no deterministic or randomized sublinear cost algorithm can output accurate LRA for the worst case input and even for a specified small family of matrices that admit rank-1 approximation.

In our papers in Procs. of MACIS 2019 and arxiv 1906.04223, 1906.04327, 1906.04929, and 1907.10481, submitted for publication, we, however, proposed and formally supported a remedy: achieve sublinear cost performance by means of trivializing the bottleneck stages of some known algorithms: pre-process input matrices with new sparse multipliers rather than known random dense (Gaussian, SRHT, or SRFT) ones or compute LRA by means of random sampling under trivial choice of sampling probabilities (called leverage scores), thus avoiding their costly computations.

We proved that the resulting sublinear cost algorithms still output reasonably accurate LRA whp for a random input matrix that admits LRA. We call such an LRA dual, but we proposed deterministic sublinear cost LRA algorithms for a Symmetric Positive Semidefinite matrix.

Furthermore we proposed novel extension of the popular sketching (aka subspace sampling) randomized LRA algorithms. Originally they output LRA at superlinear cost, but our recursive extension of these algorithms converges at sublinear cost under some mild assumptions about an input. We also apply that algorithm to iterative refinement of a crude but reasonably close LRA, which is a task of independent importance.

The results of our extensive tests with both synthetic and real world inputs (some from SIAM Review paper of 2011 by Halko et al) are in good accordance with that formal study of LRA.

Our progress implies acceleration of various known algorithms linked to LRA[4] and should motivate further effort by ourselves and other researchers towards sublinear cost matrix computations and synergistic combination of the techniques developed in the Applied Linear Algebra and Computer Science communities. Our work on LRA was largely motivated by its expected contributions to *Deep Learning Networks*, on which I am working supported by my NSF Grant AitF CCF 1733834, $448,086, 2017-2021 (joint with Bo Yuan and Xue Lin).

[4]E.g., for the average input superfast LRA turns Fast Multipoint Method into Superfast Multipoint Method by enabling superfast (sublinear cost) computation of short generators involved in this method, which is frequently its bottleneck stage.