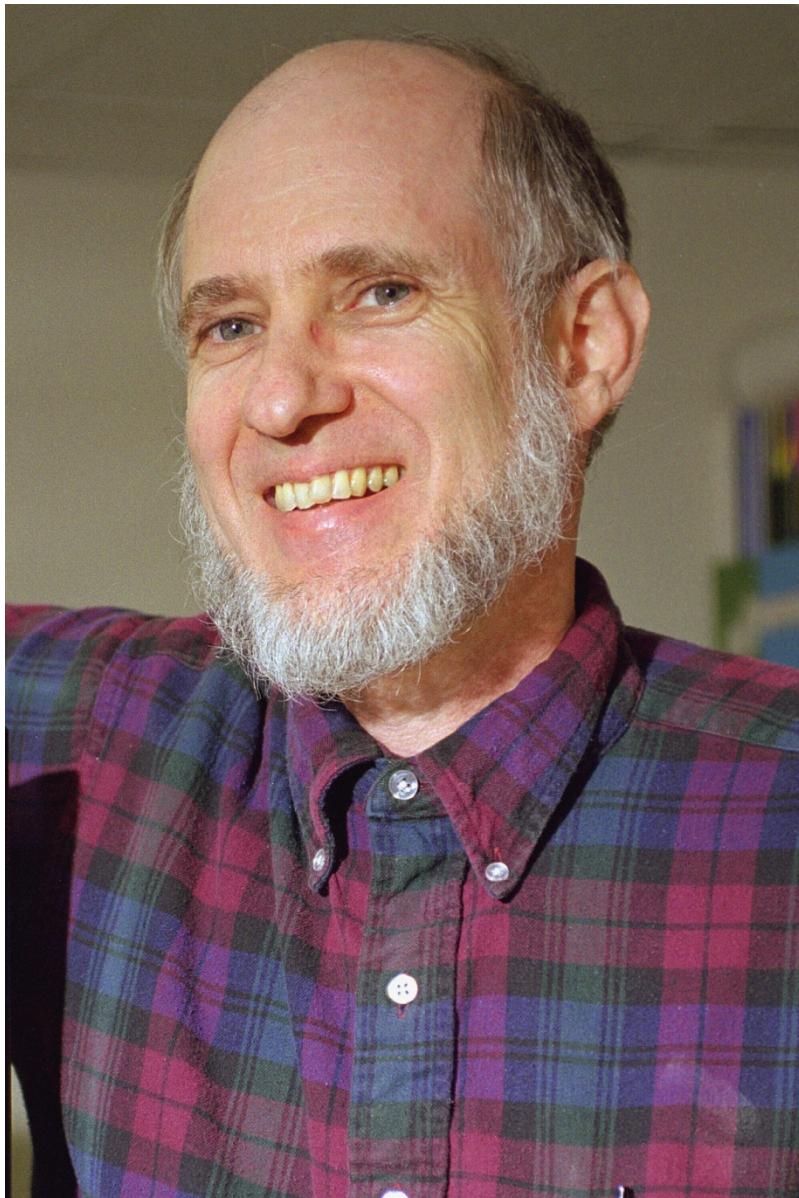


## Personal Reflections on 50 Years of Scientific Computing: 1967--2017

Journal:	<i>International Journal of Parallel, Emergent and Distributed Systems</i>
Manuscript ID	GPAA-2020-0009
Manuscript Type:	Survey
Date Submitted by the Author:	24-Jan-2020
Complete List of Authors:	Watson, Layne; Virginia Polytechnic Institute and State University, Computer Science, Mathematics, and Aerospace and Ocean Engineering
Keywords:	computer hardware, software, numerical analysis, scientific computing
Note: The following files were submitted by the author for peer review, but cannot be converted to PDF. You must view these files (e.g. movies) online.	
IJPEDS17.tex	

**SCHOLARONE™**  
Manuscripts



101x152mm (300 x 300 DPI)

1  
2  
3 Layne T. Watson ---  
4 Personal Reflections on 50 Years of Scientific Computing: 1967--2017  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60

For Peer Review Only

# Personal Reflections on 50 Years of Scientific Computing: 1967–2017

Layne T. Watson\*

*Departments of Computer Science, Mathematics, and  
Aerospace and Ocean Engineering  
Faculty of Health Sciences  
Virginia Polytechnic Institute and State University  
Blacksburg, VA 24061 USA*

Computer hardware, software, numerical algorithms, and science and engineering applications are traced for a half century from the author's perspective.

**Keywords:** computer hardware, software, numerical analysis, scientific computing

## 1. Introduction and Disclaimer

This essay constitutes purely personal and inevitably biased reflections and opinions, with likely faulty memory and a lack of technical references, due to vendor technical hardware and software user manuals for machines that existed for only a couple years (or less) being held mostly in private personal libraries. I feel privileged to have witnessed and participated in computing from close to its beginning, and as Newton said, “If I have seen further, it is by standing on the shoulders of giants.” Many of my quotes of professors and colleagues are from my course and research notes, and I apologize if the quotes are not exact.

## 2. The 1960s

My first exposure to computing was a (quarter long) sophomore level mathematics course (since computer science departments did not exist then) on computing in 1966 at Evansville College, a small private school in southern Indiana. The machine was an IBM 1620 with (modern at the time!) transistors and a magnetic core memory of exactly 20,000 decimal digits. It was programmed in machine language (not assembler!), using two decimal digit op codes, decimal addresses, and fields defined by flagging the content of a memory location. It had no registers and no hardware decimal arithmetic instructions (arithmetic was done by software, storing the multiplication table and doing lookups!). All I/O was via 80-column IBM punched cards, a line printer, and a rotating magnetic disk that only the college administration could access. I had no problem coding with the low level machine language, but struggled mightily with the IBM 026/029 card punches that regularly jammed, mangled cards, and randomly failed to print at the top of the card the characters that were punched in the card. I recall spending more time punching the program onto cards than writing the program!

The next quarter computing course was on numerical methods (more mathematics than computer science), used a very early (NCE FORTRAN) compiler, and more resembled modern computing, with one huge difference — the 20,000 decimal digit memory and the only input device being a card reader. Any serious numerical computation had to be done in stages, with intermediate results punched into cards and then read back in for the next stage. Note that loading all

---

\* Email: [ltwatson@computer.org](mailto:ltwatson@computer.org)

1  
2  
3 program stage card decks and intermediate punched result decks into the card reader was done  
4 manually, so you had to be a machine operator in addition to a scientist and programmer.  
5

6 Apparently I had a knack for writing and debugging machine language code, because the  
7 computing course instructor (Professor C. W. Buesking) asked me to become a lab assistant for  
8 the spring quarter of 1967. That experience led to an engineering co-op (cooperative education,  
9 a program whereby you alternated one quarter in school with one quarter working for industry)  
10 position in the summer of 1967 with U. S. Naval Ammunition Depot, Crane in Crane, Indiana  
11 (really, there is a naval ammunition depot in the center of the state of Indiana, nowhere near an  
12 ocean or river). Crane had an early model IBM 360/30 (a modern CISC architecture, base 16) and  
13 was getting a Honeywell 2200, and the engineers and accountants were trying to figure out how to  
14 use them. Honeywell ran a short course on COBOL, and I mainly worked with IBM FORTRAN  
15 and Honeywell COBOL. The machines were restricted, and I had little access to them (other than  
16 some number theoretic computations I ran, and managed to justify to my superior as relevant to  
17 USNAD Crane).  
18

19 Computing got much more interesting in the summer of 1968, when I worked for Nathan I.  
20 Lieberman, CPA (I remember his 'NIL' ink stamp) writing machine language code for an RPC  
21 4000. This machine had paper tape I/O, 8192 32-bit words of memory, hardware base 2 arithmetic,  
22 several registers, and a revolving drum main memory. The interesting challenge on the RPC was  
23 to not just get the program code correct, but to place the data and instructions on the drum so as  
24 to minimize the rotational latency. My lifelong career in optimization began here, when I realized  
25 that a greedy, locally optimal data placement would likely not be globally optimal for the entire  
26 program.  
27

28 Paper tape is considerably worse than cards except in one aspect — tape occupies less physical  
29 space per character. The contents of a punch card can be read from the printing at the top  
30 (assuming a working card punch!) — paper tape punches do not print on the tape. Assuming  
31 sequence numbers are punched into cards (standard procedure, e.g., FORTRAN reserves columns  
32 73–80 for sequence numbers), a dropped card deck can be quickly reassembled by a collating  
33 machine. A dropped spool of paper tape (think fishing line tangled in a fishing reel) is, well, a  
34 disaster. During my time at NIL, every programmer dropped a 100 ft spool of tape exactly *once* —  
35 and after spending the night, on their own time, untangling the 100 ft of paper tape, *never again*.  
36 Colored paper tape chad, by the way, makes great wedding confetti — I still have a box of it.  
37

38 The RPC machine code was for an interpreter called ABLE (Account Balance and Ledger  
39 Entry system), which was extremely profitable for NIL. We routinely did quarterly balance sheets  
40 and weekly payrolls for companies with 1,000 employees, all with a 8192-word RPC and paper  
41 tape data files. When I tell students what can be done with 8K words, they refuse to believe me.  
42 Current commercial software, some requiring gigabytes to run, is obscenely bloated.  
43

44 In 1969, after graduating from Evansville with majors in mathematics and psychology, I  
45 worked for another accounting firm (Wright, Griffin, and Davis, CPAs) in Ypsilanti, Michigan  
46 while a graduate student in mathematics at The University of Michigan. That work involved  
47 the niche IBM System 3, a machine with BCD (binary coded decimal) integer arithmetic, 32-bit  
48 words, a rotating disk, line printer, an assembler, and 96-column punched cards about 1/3 the size  
49 of IBM 80-column cards for I/O. I led a team (of mathematics graduate students, with no prior  
50 experience with computing) to design and build an interpreter called PROFIT (Periodic Reporting  
51 Of Financial Input Transactions) with functionality similar to ABLE, including a monitor and all  
52 the physical I/O device drivers. The input was on punched cards, the financial reports were printed,  
53 and the financial data was stored on a hard disk drive. The disk data structures involved binary  
54 structures.  
55

1  
2  
3 trees, linked lists, and garbage collection, all created by students with no prior knowledge of any  
4 of these concepts. Years later, I was contacted about why PROFIT, rewritten in COBOL (which  
5 had become available for the System 3 by that time), might be running so much slower than the  
6 original assembly language version.  
7

### 8 9 3. The 1970s

10 Michigan was a numerical analysis graduate student's dream — an IBM 360/67 (later 370/168)  
11 with special virtual memory hardware built just for Michigan, the Michigan Terminal System  
12 (MTS) operating system (the most sophisticated, user friendly OS I have ever seen — bar none),  
13 the Merit network, and star faculty such as Jim Wilkinson, Cleve Moler, Carl de Boor, David  
14 Kahaner, and George Fix. The 360/67 took time shared batch job and interactive computing to a  
15 new level, and is more impressive in retrospect than it seemed at the time.  
16

17 MTS was an extremely usable operating system, embodying human-computer interaction  
18 (HCI) principles long before HCI was a recognized specialty within computer science and industrial  
19 engineering. Many of the most powerful features of Multics and UNIX were foreshadowed in MTS  
20 (sophisticated pattern matching, long and short versions of command names and options, scripting  
21 and composition of commands, powerful editors and utilities, and on and on). As a scientific com-  
22 puting user of MTS, the feature I liked the best was the “principle of least surprise:” a command  
23 or option should do what a reasonably competent user would expect it to do. So when in doubt  
24 make a reasonable guess, and 90% of the time in MTS you'd be correct. Contrast this with current  
25 commercial software (Microsoft, Adobe, etc.), where 90% of the time a reasonable guess is wrong.  
26

27 Later (1975), on the faculty at Michigan State University, which had a CDC 6500 (one in a  
28 long line of outstanding scientific computing machines — 60-bit floating point numbers, excellent  
29 compilers and mathematical libraries), I made extensive use of the Merit network connecting major  
30 universities in the state of Michigan. In terms of file access and remote computing, Merit in 1975  
31 was every bit as functional as today's internet. E-mail would come a few years later.  
32

33 When EISPACK (a precursor of LAPACK) was under development by a consortium of uni-  
34 versities and national laboratories, Jim Wilkinson would visit Michigan for a summer school on  
35 numerical linear algebra, hosted by Cleve Moler. As a graduate student and aspiring numerical  
36 analyst, I attended all of Wilkinson's lectures. At the conclusion of the school, Moler hosted a party  
37 for all the attendees at his house. My wife and I, on the way to the party, encountered Wilkinson  
38 on the Michigan diag (campus landmark), looking a bit lost. We offered him a ride to the party,  
39 and in thanks he gave me an autographed copy of *Handbook for Automatic Computation* that he  
40 was using for his lectures. That is one of my most treasured books.  
41

42 At Michigan, I was privileged to have learned from Wilkinson and Moler, took the core nu-  
43 matical analysis courses from Carl de Boor, studied under David Kahaner, and officially finished  
44 the Ph.D. under George Fix (my thesis topic and early supervision actually came from Katta G.  
45 Murty, Department of Industrial and Operations Engineering, before Fix arrived). From my notes  
46 during those years, I have some quotes that I find especially trenchant. Cleve Moler, on the inverse  
47 of a matrix: “there is nothing you can do with the inverse that you can't do better without it.”  
48 Carl de Boor, on splines: “the only good thing about the natural cubic spline is its name.” George  
49 Fix, on theorems concerning random matrix properties, in the context of PDEs: “there is no such  
50 thing as a random matrix.” Larry Shampine, on comparing ODE algorithms: “the implementation  
51 of an algorithm is more important than the algorithm itself.” Richard Hamming, on everything:  
52 “the purpose of computing is insight, not numbers.”  
53

1  
2  
3 At this point in time, numerical analysis and scientific computing were searching for an aca-  
4 demic home. Numerical analysis was regarded by traditional mathematics departments as “not  
5 real mathematics” and by computer science departments as “too mathematical.” Moving from the  
6 cold north to the old south, I found a welcoming Department of Computer Science in Blacksburg.  
7 Virginia Polytechnic Institute & State University (now known as Virginia Tech for publicity pur-  
8 poses) in 1978 had an Amdahl 470 (IBM 360 clone) and only punched card batch access (it wasn’t  
9 Michigan), but sufficient to do meaningful solid and fluid mechanics applications with numerous  
10 engineering and mathematics collaborators.  
11  
12

#### 13 4. The 1980s

14  
15 Things improved quickly at VPI&SU; the Amdahl was replaced by a Honeywell 68/60 running  
16 Multics, enabling interactive computing from my office via a DECwriter terminal (132 column  
17 fan feed paper, and a dot matrix print head). I was still using this DECwriter as a printer  
18 in 2000, and gave it up only because the ink ribbons were no longer available. This decade  
19 included a severe economic recession, and disruptive computing technologies: personal computers  
20 and minicomputers. The interaction of the recession and the new technologies had profound  
21 consequences.  
22  
23

24 The first VAX 11/780 minicomputer at VPI&SU went to the Department of Geology to  
25 run canned software, the next one to Computer Science, and the third one to the Spatial Data  
26 Analysis Laboratory. The latter two ran DECnet between them for the first computer network  
27 (1980) at VPI&SU. Getting approval from the university to run this DECnet cable through the  
28 campus steam tunnels was an eye-opening experience in how entrenched bureaucracies respond  
29 to new technology. Suddenly VAXes and other minicomputers were sprouting like weeds, and  
30 the central campus computing center saw an elegant solution to their budget crisis: computing  
31 would henceforth be “decentralized” with departments, colleges, and laboratories running their  
32 own minicomputers, and the center supporting only administrative computing and one token VAX  
33 11/780. Concurrent with this was the adoption of a “personal computer requirement” for computer  
34 science and engineering incoming freshmen, essentially forcing them to provide their own computing  
35 resources (VPI&SU was the first public university to have such a requirement).  
36  
37

38 The profound result of this confluence, replicated throughout the country, was that now sci-  
39 entific computing either had to be done on personal computers, minicomputers, or at national  
40 supercomputing centers (requiring competitive proposals). Once the “big iron” was gone locally,  
41 it would take a long time for it to return to university computing centers. Ironically, “small”  
42 scientific computing (a la Matlab, Mathematica) got easier, but “large” scientific computing got  
43 much harder (unless your laboratory or company had a supercomputer).  
44  
45

46 This decade also saw the emergence of parallel computing in almost every conceivable flavor,  
47 a truly exciting period for computer and computational scientists. Alas, the market could not  
48 sustain this cornucopia of parallel architectures, and almost all the companies and/or products  
49 failed. I lost this decade working on the Elxsi 6400 (the first parallel machine at Sandia National  
50 Laboratories in 1984), Intel iPSC-32 (the “hypercube”, serial number 2, at General Motors Re-  
51 search Laboratories in Warren, Michigan), NCUBE-10 (the company sold numerous machines but  
52 went out of business without ever producing a working FORTRAN compiler), Sequent Balance  
53 21000, Sequent Symmetry S81, Cray 2, Cray Y-MP C90, Intel Paragon XP/S, SGI Origin 2000.  
54 Despite the commercial failure of parallel computing, the community learned a lot about parallel  
55 programming, interconnection topologies, communication technologies, and the tradeoffs between  
56 parallel and sequential computing.  
57  
58

1  
2  
3 shared and distributed memory systems. The community also learned that parallel programming  
4 was, and still is, *very* hard.  
5

## 6 5. The 1990s 7

8 After the hardware architecture wild west of the 80s, mainframe hardware evolution settled down  
9 to conventional, incremental, and predictable. Computing has always combined hardware, soft-  
10 ware/algorithms, and the human usability of those two for productivity. Usability and scientific  
11 productivity saw quantum leaps in the 1990s. First, there were notable advances in numerical  
12 algorithms and mathematical software. Homotopy algorithms for solving nonlinear systems of  
13 equations, using new differential geometry theory, were widely applied in science and engineering,  
14 e.g., aerospace engineering, mechanical engineering, systems biology, chemistry, solid mechanics,  
15 fluid mechanics, computer vision, robotics, material science, and many other disciplines. The deep  
16 mathematics, and computational experience of numerical analysts, were made accessible and usable  
17 via mathematical software (HOMPACK90, 1997, [11]). New algorithms and software tools became  
18 available for multidisciplinary design optimization (MDO), which fundamentally changed the way  
19 large systems (aircraft, ships, automobiles) were designed [3]. By definition, MDO requires super-  
20 computing, and MDO became both a user and a driver of high performance computing. There were  
21 major advances with new approaches to old problems:  $H_2$  control theory (versus the standard  $H_\infty$   
22 approach) [12], wavelets for signal processing [4], the DIRECT algorithm for global optimization  
23 (Don Jones, 1993) [8, 5, 6, 7].  
24

25 Perhaps the largest strides occurred in improving the usability of computing for scientists  
26 and engineers, through problem solving environments (PSEs), related to older concepts such as  
27 expert systems, decision support systems, automated assistants, and recommender systems. A  
28 PSE is narrowly domain specific, employs HCI principles for the user interface, and often trans-  
29 parently uses distributed or supercomputing for the solution of problems [10]. (Technically Matlab  
30 and Mathematica are PSEs for mathematical models already formulated, and do not assist in the  
31 *formulation* of those models.) Examples of PSEs at VPI&SU alone include Vizcraft (high per-  
32 formance aircraft design), WBCSim (analysis and design of wood based composite materials), L2W  
33 (analysis of the effects of land use changes on watersheds and wildlife: involved computer science,  
34 hydrology, economics, biology, ecology, urban planning, civil engineering) [10], JigCell (chemical  
35 kinetic modeling of biological systems, e.g., the eukaryotic cell cycle) [1]. The advance in PSEs is  
36 having systems that match the users' terminology, thought processes, and workflow, rather than  
37 forcing the user to learn "how the computer works."  
38

39 While distributed computing is as old as networking (DEC famously advertised that with  
40 the VAX and DECnet, "the network is the computer"), new concepts and technologies emerged,  
41 generically called grid computing [9]. The new technologies were both high level (e.g., automatic  
42 acquisition of a needed service) and low level (e.g., remote authentication), and intended to expand  
43 computing capability (e.g., solve a larger problem) or improve reliability (e.g., provide redundant  
44 sources for some service) or increase machine utilization (e.g., by routing a task to a lightly loaded  
45 system). The major players in grid computing were Globus (Argonne National Laboratory), Legion  
46 (University of Virginia), Condor (University of Wisconsin), and SORCER (National Bureau of  
47 Standards, now known as the National Institute of Science and Technology).  
48

49 Last but not least, and in the same vein of usability, comes HTML and the World Wide  
50 Web (WWW, 1990). Certainly computational scientists extensively use the WWW, but whether  
51 the WWW has significantly improved computational science is moot. Computer e-mail, from the  
52

1  
2  
3 previous decade, and the open internet (circa 1990) for file transfer and remote login have more  
4 directly aided scientific computing collaboration.  
5

## 6. The 2000s

7  
8 The ballyhooed Y2K disaster (when all the world's computer databases would turn into pumpkins  
9 at the stroke of midnight marking the turn of the century) never happened. However, the .com bust  
10 and ensuing deep recession, a real life version of the tale of Icarus, did happen. In 2003, an NSF  
11 equipment grant to Srinidhi Varadarajan, myself, and others was leveraged, in a rare confluence  
12 of departmental, college, and university priorities, to build System X ("ten," for ten teraflops,  
13 and for running Apple OS/X). System X consisted of 2200 Apple G5 processors, connected with  
14 Infiniband, and running OS/X. At 12.5 teraflops, System X came in at number three in the 2003  
15 world Top 500 computer rankings, behind two national laboratory machines that cost at least  
16 three times as much. After a story in the New York Times about System X, and some national  
17 laboratory managers were asked to justify the budgeted cost of their proposed comparable new  
18 machines, there was a shocking backlash against VPI&SU's accomplishment. I was in the audience  
19 at a conference at Oak Ridge National Laboratory when a speaker from a national laboratory  
20 (not ORNL) said System X was built by "pizza fueled undergraduates," was "a fake," and "not a  
21 serious machine." I assured the speaker that System X was real and being used for serious work,  
22 and invited him to come up the road to Blacksburg and see for himself. He did not respond.  
23

24 System X was the beginning of the return of "big iron" to university campuses, but the return  
25 was slowed by another deep recession and depressed computer science and STEM enrollments.  
26 Computer companies that were once the heart of the industry went under, bought up for their  
27 reputation but not their product research and development. DEC, SUN, and SGI disappeared.  
28 Computer chip manufacturing underwent a similar global consolidation.  
29

30 As the end of Moore's Law appeared nearer and nearer, chip manufacturers considered other  
31 materials and other manufacturing processes, and computer architects pursued parallelism to main-  
32 tain the steady increase in performance. Multicore (2, 4, 8 cores) architectures (shared memory  
33 parallelism) with compiler support became common, although even a modest speedup of seven  
34 with eight cores on a general purpose and highly variable instruction stream was rarely achieved.  
35 Everyone knew this day would likely come (the end of increases in single thread performance), but  
36 now it was uncomfortably close.  
37

38 An emerging major consumer of supercomputing was bioinformatics (the intersection of some  
39 parts of computer science with some parts of biology), as academic departments hired bioinformati-  
40 cians and universities created bioinformatics research centers (the Virginia Bioinformatics Institute  
41 at VPI&SU was one of the first such centers, built, ironically, with Tobacco Master Settlement  
42 Agreement money). The mix of scientific computing jobs run on computing center machines be-  
43 gan to shift from almost exclusively the physical sciences and engineering toward the life sciences  
44 (biology, bioinformatics, medicine, pharmacology) and data analysis from all disciplines, including  
45 business and the social sciences.  
46

## 47 7. The 2010s

48 The current frontier is exascale computing, but the path is far from certain. Exascale power  
49 requirements, and power consumption of large machines in general, have become major issues.  
50 Both automated power management systems, based on data mining and profiling, and individual  
51 programs controlling their own power usage are being studied. Chip manufacturers are increasing  
52 the number of cores (128, so far) on a single chip, bringing shared memory programming to the  
53

1  
2  
3 forefront again. Distributed memory programming and communication topologies are still the  
4 backbone of high performance scientific computing.  
5

6 As computing systems move toward exascale, the mean time to failure (MTF) of some component  
7 becomes less than the run time of a typical scientific computation, making hardware/software  
8 resilience a *sine qua non* for production computing at the exascale. As the complexity of scientific  
9 computations grows, and the inherent variability of hardware performance also grows (the latter  
10 has been confirmed recently by a NSF CSR project), the reproducibility of both scientific results  
11 and system (hardware, operating system, application code) behavior becomes a serious concern,  
12 affecting the confidence in the very computational science enterprise itself.  
13

14 Large segments of existing computer science, statistics, mathematics, and electrical engineering  
15 have been co-opted under the name of “machine learning” (ML, also known as “data mining”)  
16 to analyze “big data;” presumably all analyzed data heretofore had been “small.” More erudite  
17 sounding terms for the general activity are “data analytics” and “knowledge discovery,” as if finding  
18 a pattern in data constitutes an explanation or understanding. As the statisticians say, “correlation  
19 does not imply causality.” One is reminded of genetic algorithms, evolutionary computation,  
20 and “biologically inspired” algorithms, as if a mere analogy with something biological confers  
21 superiority. Nature is, in fact, often far from optimal, as any biologist will attest. To paraphrase  
22 my earlier de Boor quote: “The only thing good about genetic algorithms or artificial neural  
23 networks is their name.”  
24

25 Many phenomena that are stochastic are modeled as deterministic (e.g., chemical reactions by  
26 ordinary differential equations), or a stochastic quantity is replaced by a deterministic functional  
27 of that quantity (e.g., a random variable by its mean), for theoretical or computational convenience.  
28 Theoretical and algorithmic progress has been made in dealing directly with the stochastic  
29 quantities without assuming something is deterministic. MDO under uncertainty (known as ro-  
30 bust design) or stochastic systems biology problems (e.g., parameter estimation for a stochastic  
31 cell cycle model) can now be attacked directly without making ensembles of runs to estimate a  
32 deterministic functional (e.g., a variable’s mean). One such new algorithm, supported by rigor-  
33 ous convergence theory, is called QNSTOP (quasi-Newton stochastic optimization), for which high  
34 quality mathematical software exists [2]. Similar important advances have been made in the theory  
35 and numerical solution of stochastic ordinary and partial differential equations, which are often  
36 more faithful representations of reality.  
37  
38

## 40 41 6. Epilogue

42 These past 50 years of scientific computing have been quite a run, and I wish the following genera-  
43 tions an equally exciting and enjoyable 50 years. Pergite Veneti!

## 44 45 46 References

47  
48 [1] N. A. Allen, C. A. Shaffer, M. T. Vass, N. Ramakrishnan, and L. T. Watson, *Improving the development*  
49 *process for eukaryotic cell cycle models with a modeling support environment*. *Simulation* 79 (2003),  
50 pp. 674–688.  
51 [2] B. D. Amos, D. R. Easterling, L. T. Watson, W. I. Thacker, B. S. Castle, and M. W. Trosset, *Algorithm*  
52 *XXX: QNSTOP—quasi-Newton algorithm for stochastic optimization*. Technical Report TR-14-02,  
53 Dept. of Computer Sci., VPI & SU, Blacksburg, VA, 2014.  
54 [3] S. Burgee, A. A. Giunta, V. Balabanov, B. Grossman, W. H. Mason, R. Narducci, R. T. Haftka, and  
55 L. T. Watson, *A coarse-grained parallel variable-complexity multidisciplinary optimization paradigm*.  
56 *Internat. J. Supercomputer Appl. High Performance Comput.* 10 (1996), pp. 269–299.  
57  
58  
59

1  
2  
3 [4] I. Daubechies, *Ten Lectures on Wavelets*, SIAM, Philadelphia, 1992.  
4 [5] J. He, A. Verstak, L.T. Watson, and M. Sosonkina, *Performance modeling and analysis of a massively*  
5 *parallel DIRECT: Part 1.* Internat. J. High Performance Comput. Appl. 23(1) (2009), pp. 14–28.  
6 [6] J. He, A. Verstak, M. Sosonkina, and L.T. Watson, *Performance modeling and analysis of a massively*  
7 *parallel DIRECT: Part 2.* Internat. J. High Performance Comput. Appl. 23(1) (2009), pp. 29–41.  
8 [7] J. He, L.T. Watson, and M. Sosonkina, *Algorithm 897: VTDIRECT95: Serial and parallel codes for*  
9 *the global optimization algorithm DIRECT.* ACM Trans. Math. Software 36(3) (2009), Art. 17, pp.  
10 1–24.  
11 [8] D. R. Jones, C. D. Perttunen, and B. E. Stuckman, *Lipschitzian optimization without the Lipschitz*  
12 *constant.* J. Optim. Theory Appl. 79(1) (1993), pp. 157–181.  
13 [9] N. Ramakrishnan, L. T. Watson, D. G. Kafura, C. J. Ribbens, and C. A. Shaffer, *Programming*  
14 *environments for multidisciplinary grid communities.* Concurrency Comput.: Pract. Exper. 14 (2002),  
15 pp. 1241–1273.  
16 [10] L. T. Watson, V. K. Lohani, D. F. Kibler, R. L. Dymond, N. Ramakrishnan, and C. A. Shaffer,  
17 *Integrated computing environments for watershed management.* J. Comput. Civil Engrg. 16 (2002),  
18 pp. 259–268.  
19 [11] L. T. Watson, M. Sosonkina, R. C. Melville, A. P. Morgan, and H. F. Walker, *Algorithm 777: HOM-*  
20 *PACK90: A suite of Fortran 90 codes for globally convergent homotopy algorithms.* ACM Trans. Math.  
21 Software 23 (1997), pp. 514–549.  
22 [12] D. Žigić, L. T. Watson, E. G. Collins, Jr., and D. S. Bernstein, *Homotopy approaches to the  $H_2$*   
23 *reduced order model problem.* J. Math. Systems, Estimation, Control 3 (1993), pp. 173–205.  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60