

# THE FUTURE OF COMPUTING: VLSI, STORAGE, AND NETWORKS REFLECTIONS ON A LEGACY OF INNOVATION

*Bart Blaner*

IBM Corporation  
Essex Junction, Vermont, USA  
Email: blaner@us.ibm.com

## ABSTRACT

Building on the legacy of innovation in computing by Stamatis Vassiliadis and others, the future of computing in the principal areas of CMOS VLSI technology, storage, and networks will be explored, with extrapolations to large-scale systems. Applications and the societal impact of these technological advancements will be considered.

## 1. LEGACY

Human history, whether recorded in the layers of sediment in an archeological dig or in the pits of a DVD, records a continuous advancement of technology, from the primitive tools and implements found in caves, to modernity's most powerful supercomputers. Seldom is such advancement easy, but the human spirit is characterized not by giving up in the face of adversity or limits, but rather seeking ways around or over them. Throughout the human project, certain gifted men and women have had the critical insights, the energy of mind, the ingenuity, to make the next lurch forward. Such a person was Stamatis Vassiliadis. Through his lifetime of engagement with students, colleagues and friends, he left a legacy of critical thinking and advancement of computer technology. Extrapolating on this legacy will let us probe the future of computing and conclude that the advancement will continue in spite of the many difficult challenges that lay ahead.

Stamatis and I became co-workers in 1986 and quickly became friends. We were collaborating on a System/370 (called zSeries today) mainframe computer development project at IBM in Endicott, New York. I was a fairly new engineer working on the instruction sequencing and fixed-point arithmetic units, he leading the floating-point arithmetic unit. After this project, we began working closely together exploring ideas for overcoming what we perceived as barriers to microprocessor performance, in particular cycle-time limitations to the further exploitation of instruction-level parallelism.

Stamatis, whose forte, like his fellow Samian Pythagoras, was arithmetic, had innovative ideas about increasing instructions per cycle (IPC) while at the same time not increasing processor cycle time by using multi-operand arithmetic units to execute instructions in parallel. He called such putting-together of multiple instructions into conceptually a single instruction "compounding". Initially, our focus was on three operand adders but later we broadened the scope to include logical and other operations [1] and other frequent combinations of instructions. This innovative concept has been used to improve the performance of a number of microprocessors, including IBM's Power2 [2] and more recently Intel microprocessors where the technique is known as "fusion" [3].

Stamatis envisioned a software preprocessor to bind together instructions with tag information to indicate compounding. One of my key contributions was to move this function to a hardware pre-decoder on the input side of the instruction cache, marking compounded instructions with tag bits stored in the instruction cache, thus making the mechanism applicable to the "dusty decks" of old object code our S/370 machines would have to run. Such pre-decoding or instruction tagging on the input side of the instruction cache has become routine [4, 5, 6]. The technique had the additional benefit of simplifying parallel instruction issuing on the output side of the instruction cache, reducing the complexity of the decision-making logic and keeping such from impacting the cycle time or adding additional decode cycles.

Much effort was devoted to efficiently processing branch instructions: our vision of instruction execution was a brick wall of computation or load/store instructions rarely interrupted by a mis-predicted branch. Such branchless

execution streams are common today, for example [4], but not so common in 1990. Many innovative techniques were developed [7] and found their way into other processors such as the AS/400 described by Jeremiah in [5].

All of these ideas we put under the appellation of SCISM [8]—Scalable Compound Instruction Set Machine—which we felt an apt title, a homophone for schism and for us a tongue-in-cheek quip about what we believed to be a departure or schism from what we considered to be traditional microarchitecture of the time.

Departures from tradition are characteristic of innovation. In the following sections, we will consider several such departures in fundamental areas of computer engineering: CMOS VLSI technology, storage, and networks. This will give us insight about where these technologies may be headed in the future. We will then consider some hypothetical applications of these advancements and comment on their societal impact.

## 2. CMOS VLSI TECHNOLOGY

IBM CMOS VLSI technology has an established record of years of innovation to maintain its trajectory on the curve predicted by Moore’s Law. This curve is shown in Fig. 1. With the exception of super-linearity in the late 1990s, the trajectory is mostly linear with time on logarithmic scales for logic gate capacity and lithography dimension. This technology path was fueled by key innovations at critical junctures as shown. With lithographic dimensions approaching tens of atomic diameters, this will surely continue to be even more important in the future.

The first mainframe on which I worked with Stamatios used the IBM CMOS2 technology, which would be the leftmost node in the figure. With maximum capacities of 60k gates and a mere two levels of metal, a four-chip multi-chip module (MCM) was required to contain instruction sequencing and fixed-point arithmetic units alone. More MCMs contained the floating-point unit, instruction and data caches, L2 cache, and I/O. By the time of our collaboration on SCISM, maximum wireable gates increased 27 fold due in part to the innovations of chemical mechanical planarization (CMP) and shallow trench isolation (STI), which facilitated the addition of more metal layers, greatly increasing the wireable gates available on a silicon die. Suddenly, we had the realization that things like adders, formerly area expensive, had become “commodities”, i.e., if a 32-bit adder was needed locally for some function, we put one in instead of trying to share one located elsewhere on the chip. This opened the door to increased instruction level parallelism, since we quickly realized we could in many cases replicate entire execution units without escalating the cost of the product.

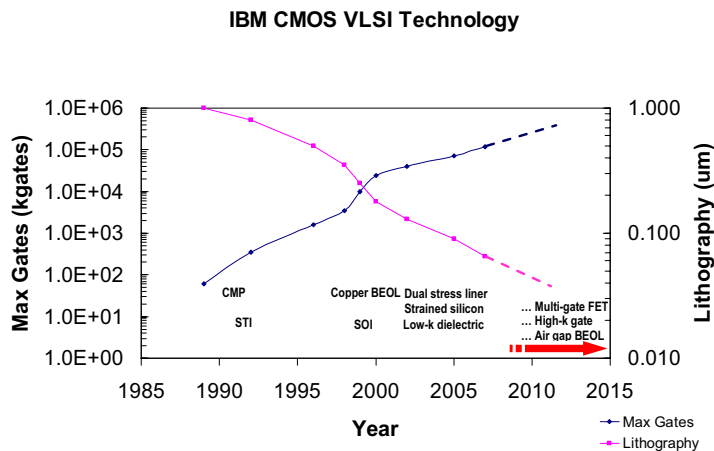


Fig. 1. IBM CMOS VLSI technology [courtesy T. Bednar and N. Rohrer, IBM Corp.].

The CMOS VLSI technology innovations have continued steadily from those days. Laymen seem to have the impression that Dennard scaling [9] comes simply from a “turn of the crank”, but this has never been the case. As we have seen, CMP and STI were fundamental to being able to use the increased gate count of chips. Other innovations have allowed the performance, power, and area scaling to continue, as shown in the figure:

- Silicon on insulator (SOI) reduced junction capacitance, provided a performance boost, and helped with power consumption [10].
- Strained silicon improved carrier mobility, reducing circuit delay [11]. Dual stress liners allowed this to be applied to both pFETs and nFETs.
- Low-k dielectric material in the back-end of line (BEOL, i.e., metal layers) reduced the capacitance of the metal layers thereby reducing signal propagation delay in the metal.

Nevertheless, new challenges have arisen. Vdd scaling has either slowed significantly or ceased, while circuit density continues to scale, and since power is quadratic with voltage, power density has now become one of the most formidable barriers to further scaling. Passive or leakage power dominates in technologies beyond the 130 to 65nm regimes [12]. Vdd has not scaled much below 1.0V. As [12] points out power density limited functions may have to be redesigned just to avoid increasing power density with no gain in performance whatsoever. Technologists offer several innovations to help surmount the power density problem, including high-k gate insulator to reduce leakage [13], and multiple-gate devices, which provide better control over the channel and offer performance improvement for a few more technology generations. Several multiple-gate device topologies offer promise, including the double-gate device and FinFET [12]. These transistor topologies will cause disruptions to both circuit design and manufacturing that will require further innovation to see broad use.

Circuit innovations such as Limited Switch Dynamic Logic circuits [14] have shown the capability to allow continued power supply scaling to some extent. In [14], Montoye et al. show a multiplier design in 90nm using LSDL that scales in frequency linearly over 0.9V to 1.2V range and shows 2GHz operation down to 0.75V. Recently an 8T SRAM cell was demonstrated that scales linearly down to 0.41V without the use of either dynamic or secondary power supplies [15].

If increased circuit delay can be traded for circuit parallelism and energy efficiency, sub-threshold-voltage scaling can dramatically reduce energy consumption. Many challenges to this approach remain to be grappled with, including coping with variability and noise [16].

Increasing wire delays remain a challenge to scaling. As wire cross sections shrink, resistivity of the metal increases and the propagation delay increases. Innovations in the metal stack such as substituting copper metallurgy for aluminum to reduce resistivity and low-k dielectric material to reduce propagation delay have helped put the metal delay back in line, but more innovation is needed for the 45nm node and beyond. Recently, for example, IBM announced so-called air gap technology, where voids are created in between metal traces to further reduce the dielectric constant, reducing wire delay by 35% compared to wires with carbon silicate dielectric [17].

While these evolutionary technology steps have enabled continued performance and density scaling, in the future a more revolutionary approach is needed that attacks the problem from multiple fronts. An exciting technology on the horizon that does this is 3D integration [18], where multiple layers of chips can be stacked in the vertical dimension and interconnected with a high density of very short metal vias. One promising form demonstrated by IBM [19] is illustrated in Fig. 2. In this SOI-based so-called face-to-back process, a wafer of active devices and metal layers is manufactured using traditional means. A glass handle is attached to the wafer to support further processing. The wafer backside is thinned in the extreme, with all the bulk silicon removed. The resulting wafer can be bonded on top of another wafer consisting of more active devices and a metal stack. The glass handle is removed and the wafers are interconnected in the  $z$  dimension by through-silicon vias (TSV). TSV lengths are a few microns, and their density is seen increasing an order of magnitude over the next decade, from 100,000 to over 100,000,000 TSV/cm<sup>2</sup>.

The technology yields benefits on multiple fronts, including

- significant reduction in wire lengths, some by an order of magnitude or more
- power reduction due to fewer long wires needing buffers to drive them
- increased device density due to use of the vertical dimension for active device integration




Fig. 2. IBM 3D integration SOI face-to-back process abstract [20].

Exploiting all of these, one can envision morphing today’s two-dimensional microprocessor floorplan into three dimensions, for example, multiple cores comprising control circuit planes and computation planes, and one or more cache storage planes. Such a structure would reduce what are today millimeter-long wires to tens of microns, significantly reducing delay on wire-dominated critical paths. The potential for greatly increasing bandwidth between cache levels and processor cores is significant, providing an architectural benefit as well by greatly reducing trailing edge effects [20], giving more performance to bus-limited multi-core and multi-threaded systems.

Technical hurdles remain to be cleared to make 3D integration a reality. In the manufacturing process are challenges in wafer thinning, alignment, and bonding. In the EDA realm are issues such as

- floorplanning in three dimensions with the addition of thermal constraints imposed by the presence of active devices in deeply buried layers
- computer architecture design that takes into consideration the positioning of units or portions of units on different levels in the vertical dimension
- system level design considering integration of heterogeneous technologies such as analog/mixed signal, DRAM, etc.

### 3. STORAGE

Magnetic recording has been growing in storage density at a rate of 60 to 100% annually for many years, with an accompanying plummet in cost per megabyte [21]. The amount of storage one can carry in the pocket today would require over 20 hectares (over 50 acres) of magnetic recording material surface area in IBM RAMAC technology of the late 1950s. Like silicon technology, magnetic storage faces looming fundamental limits. The so-called superparamagnetic effect, where thermal energy from ambient temperatures is sufficient to disrupt the state of a magnetically stored bit, is thought to occur in the neighborhood of 250Gbit/in<sup>2</sup> [22]. To sidestep this effect, researchers have proposed alternative storage materials and means for reading and writing the stored data. Building on pioneering work in Atomic Force Microscopy (AFM) and combining a wide spectrum of disciplines, including nanoscale machine fabrication, digital signal processing, and materials science, a research team at IBM Zurich has developed “Millipede”—a thermomechanical MEMS—based scanning-probe storage mechanism [23] that has demonstrated storage densities of 641Gbit/in<sup>2</sup> with 1Tbit/in<sup>2</sup> on the horizon [22].

A schematic of the Millipede concept is shown in Fig. 3. An array of nanometer sharp tips is controlled mechanically and electrically to create and sense pits in a polymer material. The presence of a pit indicates a binary “1” while absence of a pit is a binary “0”. In the 641Gbit/in<sup>2</sup> demonstration, each pit was spaced approximately 100 atoms from its neighbor and was a mere 10 atoms deep! Different from AFM, where the tip does not touch the material, the Millipede tip impinges on the polymer and is heated to form the pit and write a “1”, and scans the material with the tip dipping into a pit and sensing a small change in current to detect the presence or absence of a pit. Sophisticated servo control and digital signal processing techniques are required to both scan the bits in parallel from the storage medium and recover the minute data signals from noise.

The Millipede concept is interesting not only in its demonstration of a way over the barriers to continued growth in magnetic recording storage density, but also in its exhibition of the human capability to manipulate structures




Fig. 3. Schematic of the Millipede concept [23].

at the nanometer scale or atomic levels. Such capabilities will doubtless prove useful in surmounting obstacles in other areas and applications of computing technology in the future.

#### 4. NETWORKS

Advancement in off-chip networks has leapt forward in recent years, providing the infrastructure for both the global internet and telephone networks. For example, since its origin in the late 1970s, Ethernet bandwidth has grown from 10Mbps to 10Gbps, with 100Gbps recently demonstrated [24] and with 10Tbps speeds predicted in the coming decade [25].

More recently, VLSI densities have enabled multiple processor cores to be integrated on a single chip and attention has turned to bringing off-chip networking concepts on-chip to tackle the challenge of non-scaling wire delays. Peering into what kinds of networks would be required for a 3D integrated multi-core device, a team at Penn State University [26] proposes a combination of mesh-connected compute-node/packet-router pairs for interconnection in the  $x$ - $y$  plane, with a TDMA bus structure for communicating in the much shorter  $z$  direction. This research also addresses some of the 3D integration thermal issues by designing a network topology that allows for staggered placement of the heat-generating cores.

Agrawal's Raw multicore processor project at MIT [27] takes a further step by elevating the status of the on-chip network to first class citizen so to speak, making it accessible to the instruction set and thereby enabling, if not forcing, the compiler and system software to directly consider the network and its latencies in a multi-core system. A tiled, mesh-connected multi-core processor design ensures scalable one-cycle latency to nearest-neighbor core with static routes known at compile time, and a dynamic multi-hop network provides the more traditional dynamically routed network such as servicing caches misses, message passing, etc. Significant speedups have been demonstrated over a broad range of applications and scalability has been explored over a range of tiles counts [28].

The limitations of electrical signaling have driven digital signal transmission media to optical networks, but interfacing silicon-based I/O to optical networks has presented many difficulties. Nevertheless, IBM recently demonstrated an optical transceiver [29] that achieves 160Gbps bandwidth in a mere 17mm<sup>2</sup>. Because it is built using a standard CMOS process plus gallium arsenide and indium phosphide for the optical components, it is not difficult to imagine a future device integrated into the top layer of a 3D integrated module as its network and storage interface.

#### 5. TYING IT TOGETHER

As an example of where these technologies might lead, let us hypothesize an extrapolation from the IBM Blue Gene supercomputer. The Blue Gene/L system installed at Lawrence Livermore National Laboratory has held the top position among supercomputers for some time, delivering 280.6TFlops (tera-floating-point-operations per second) [30]. In a Blue Gene/L system [31], a rack comprises 1024 compute nodes each having two PowerPC cores and each of them capable of 4 Flop per cycle, all built from 130nm technology clocked at 700MHz. Five networks, each designed to best fit its particular type of communication task, interconnect the compute nodes, racks, storage, and other I/O to make a complete system. Taking the next logical step on the technology scaling curve, the recently

announced Blue Gene/P [32] achieves over 1Pflops (1000Tflops) sustained performance, integrating four PowerPC cores per node. Where would we expect Blue Gene to go if 3D integration, on-chip networks, ultra-dense storage, and integrated optical transceivers all come together in the next decade? Consider a 3D integrated module comprising chips built using 22nm CMOS VLSI technology—two generations from the state of the art today. Keeping the chip size and frequency constant, 64 cores could fit on a chip and a 6-layer 3D stack would deliver over 1TFlops. A single 1024-node rack consisting of these modules would deliver 1PFlops, nearly the equivalent of a 216-rack Blue Gene/P system. This ignores further advantages to performance gained by reducing latency and increasing bandwidth that arise from integrating a 3D network as proposed by [26], reducing or eliminating some trailing edge effects, adding a layer to the module for an optical off-module transceiver, and integrating ultra-dense storage capacity in close proximity to the memory subsystem. ExaFlops performance is in view for a large system.

At the other end of the spectrum, one can hypothesize about pocket-sized computing-communications devices of incredible capacity and capability. Imagine, for example, a 3D integrated device with performance like today’s supercomputer, enabled with high bandwidth wireless communication, GPS capable, having terabytes of storage capacity, and so on. One could expect such a device to finally have useful speech recognition, understand natural language commands, accurately translate languages, and provide storage for *all* of one’s digital data, including video and medical databases.

## 6. APPLICATIONS AND SOCIETAL IMPACTS

When Stamatis and I started our careers in computer engineering, the Cray-1, with its capability of a few hundred Mflops in 1976, was the world’s fastest supercomputer and was targeted at applications in physics, weapons research, and weather forecasting. By the late 1990s, with burgeoning compute power, supercomputers would start to play a key role in the life sciences, for example, designing drugs and mapping the human genome. The life sciences, especially the areas of computational biology and computational chemistry, have an insatiable demand for compute power, as topics such as protein folding, DNA sequence analysis and classification, organ simulation, and the like, are explored [33]. The computational demands for protein folding are enormous: at 1Pflops, where Blue Gene/P is today, three years are required to simulate a  $100\mu s$  protein fold, and yet some proteins require approximately 1 second to complete a fold! Clearly there is more work to be done in the interdisciplinary domain of computer engineering, biology, and computational chemistry. The potential benefit to the human species in accomplishing such simulations is great: understanding protein folding, and especially misfolds, may hold the key to unlocking cures for dread diseases that plague us, such as cancer, AIDS, cystic fibrosis, Cruetzfeldt-Jakob disease, and many others [33].

The unrelenting advance in computer technology is not without its dark side. Improvement in human behavior lags far behind the pace of technological advancement, if it has improved much at all since man began scratching pictures on cave walls. There seems to be no end to the stream of barbarians at the gate, exploiting advances in computation to steal identities, stalk individuals, gain access to private information, and commit other predatory acts. We can be assured that such miscreants will take full advantage of the technology produced by the advancements discussed, giving them access to awesome computational power. Matching advancements in computer security and ethics are needed to both defend the innocent against such attacks and stem the tide of such misanthropic behavior. As Aristotle said, “At his best, man is the noblest of all animals; separated from law and justice he is the worst”.

## 7. CONCLUSION

Aristotle also said, “The energy of the mind is the essence of life”. We friends, colleagues, and students of Stamatis Vassiliadis have all been touched by the energy of his great mind. Critical thinking and creativity defined much of the man he was. The innovations he made in his profession, the colleagues with whom he labored and debated, and the students whom he taught both how and what to think, have left an indelible legacy to humanity. We carry that legacy with us: the discoveries and advances we, his friends, will continue to make, will propel computer engineering and technology forward into the future. Formidable obstacles will doubtless continue to arise, but the human spirit is one of overcoming, and that, I believe, gives us reason to be optimistic about the future of computing.

## 8. ACKNOWLEDGEMENTS

I would like to extend thanks to my colleagues Tom Bednar, Kerry Bernstein, Evangelos Eleftheriou, Michael Gschwind, Bob Montoye, and Norm Rohrer for the information they provided, and to Audrey Holm-Hansen, Tom Jeremiah, and John Urda for reviewing the draft.

I will be forever grateful that Stamatis' path intersected mine and that we were able to enjoy a few years working together on fascinating and challenging things. He will be in my heart always.

## 9. REFERENCES

- [1] S. Vassiliadis, J. E. Phillips, and B. Blaner, "Interlock Collapsing ALUs," *IEEE Trans. Computers*, July 1993.
- [2] J. Barreh, S. Dhawan, T. Hicks, and D. Shippy, "The POWER2 Processor," *Compton Spring 1994 Digest of Papers*, pp. 389 – 398, 1994.
- [3] "Intel<sup>®</sup> 64 and IA-32 Architectures Optimization Reference Manual, Order Number: 248966 – 015," May 2007.
- [4] J. M. Tendler, J. S. Dodson, J. S. Fields Jr., H. Le, and B. Sinharoy, "POWER4 System Microarchitecture," *IBM J. of Res. and Dev.*, vol. 46, No. 1, pp. 5 – 25, January 2002.
- [5] T. L. Jeremiah, S. Vassiliadis, and B. Blaner, "Superscalar Branch Instruction Processor," in *Proceedings Vol. 2., 12th International Conference on Control Systems and Computer Science*, January 1999, pp. 163 – 168.
- [6] C. N. Keltcher, K. J. McGrath, A. Ahmed, and P. Conway, "The AMD Opteron Processor for Multiprocessor Servers," *IEEE Micro*, pp. 66 – 76, March-April 2003.
- [7] B. Blaner, S. Vassiliadis, and T. L. Jeremiah, "A Branch Instruction Processor for SCISM Organizations," in *Proceedings of the 21st EUROMICRO Conference*, September 1995.
- [8] S. Vassiliadis, B. Blaner, and R. J. Eickemeyer, "SCISM: A Scalable Compound Instruction Set Machine," *IBM J. of Res. and Dev.*, vol. 38, No. 1, pp. 59 – 77, January 1994.
- [9] R. H. Dennard, F. H. Gaensslen, V. L. R. H.-N. Yu, E. Bassous, and A. Leblanc, "Design of Ion-Implanted MOSFETs with Very Small Physical Dimensions," *IEEE. J. Solid-State Circuits SC-9*, pp. 256 – 268, 1974.
- [10] G. G. Shahidi, "SOI Technology for the GHz Era," *IBM J. Res. and Dev.*, vol. 46, No. 2/3, pp. 121 – 131, March/May 2002.
- [11] J. Chu, H. Chen, K. A. Jenkins, T. Kanarsky, K. Lee, A. Mocuta, H. Zhu, R. Roy, J. Newbury, J. Ott, K. Petrarca, P. Mooney, D. Lacey, S. Koester, K. Chan, D. Boyd, M. Jeong, and H.-S. P. Wong, "Characteristics and Device Design of Sub-100 nm Strained Si n- and pMOSFETs," in *Symp. VLSI Technol.*, 2002, pp. 98 – 99.
- [12] W. Haensch, E. J. Nowak, R. H. Dennard, P. M. Solomon, A. Bryant, O. H. Dokumaci, A. Kumar, X. Wang, J. B. Johnson, and M. V. Fischetti, "Silicon CMOS Devices Beyond Scaling," *IBM J. of Res. and Dev.*, vol. 50, No. 4/5, pp. 339 – 361A, September 2006.
- [13] E. P. Gusev, V. Narayanan, and M. M. Frank, "Advanced High – k Dielectric Stacks with PolySi and Metal Gates: Recent Progress and Current Challenges," *IBM J. Res. and Dev.*, vol. 50, No. 4/5, pp. 387 – 410, 2006.
- [14] W. Belluomini, D. Jamsek, A. K. Martin, C. McDowell, R. K. Montoye, H. C. Ngo, and J. Sawada, "Limited Switch Dynamic Logic Circuits for High-Speed Low-Power Circuit Design," *IBM J. Res. and Dev.*, vol. 50, No. 2/3, pp. 277 – 286, 2006.
- [15] L. Chang, Y. Nakamura, R. K. Montoye, J. Sawada, A. K. Martin, K. Kinoshita, F. H. Gebara, K. B. Agarwal, D. J. Acharyya, W. Haensch, K. Hosokawa, and D. Jamsek, "A 5.3GHz 8T-SRAM with Operation Down to 0.41V in 65nm CMOS," in *Symposium on VLSI Circuits*, 2007, pp. 252 – 253.
- [16] S. Hanson, B. Zhai, K. Bernstein, D. Blaauw, A. Bryant, L. Chang, K. K. Das, W. Haensch, E. J. Nowak, and D. M. Sylvester, "Ultralow-voltage, minimum-energy CMOS," *IBM J. of Res. and Dev.*, vol. 50, No. 4/5, pp. 469 – 490, September 2006.
- [17] <http://www-03.ibm.com/press/us/en/pressrelease/21473.wss>.
- [18] K. Banerjee, S. J. Souri, P. Kapur, and K. C. Saraswat, "3-D ICs: A Novel Chip Design for Improving Deep-Submicrometer Interconnect Performance and Systems-on-Chip Integration," *Proc. IEEE*, vol. 89, No. 5, pp. 602 – 633, May 2001.
- [19] A. W. Topol, D. C. La Tulipe, Jr., L. Shi, D. J. Frank, K. Bernstein, S. E. Steen, A. Kumar, G. U. Singco, A. M. Young, K. W. Guarini, and M. Jeong, "Three-dimensional integrated circuits," *IBM J. of Res. and Dev.*, vol. 50, No. 4/5, pp. 491 – 506, September 2006.

- [20] K. Bernstein, P. Andry, J. Cann, P. Emma, D. Greenberg, W. Haensch, M. Ignatowski, S. Koester, J. Magerlein, R. Puri, and A. Young, "Interconnects in the Third Dimension: Design Challenges for 3D ICs," in *Proc. DAC 2007*, June 2007.
- [21] D. A. Thompson and J. S. Best, "The Future of Magnetic Data Storage Technology," *IBM J. of Res. and Dev.*, vol. 44, No. 3, pp. 311 – 322, May 2000.
- [22] H. Pozidis, W. Häberle, D. Wiesmann, U. Drechsler, M. Despont, T. R. Albrecht, and E. Eleftheriou, "Demonstration of Thermomechanical Recording at 641 Gbit/in<sup>2</sup>," *IEEE Trans. Magnetics*, vol. 40, No. 4, pp. 2531 – 2536, July 2004.
- [23] E. Eleftheriou, T. Antonakopoulos, G. K. Binnig, G. Cherubini, M. Despont, A. Dholakia, U. Dürig, M. A. Lantz, H. Pozidis, H. E. Rothuizen, and P. Vettiger, "Millipede-A MEMS-Based Scanning Probe Data-Storage System," *IEEE Trans. Magnetics*, vol. 39, No. 2, pp. 938 – 945, March 2003.
- [24] <http://arstechnica.com/news.ars/post/20061115-8231.html>.
- [25] [http://www.directionsmag.com/article.php?article\\_id=372&trv=1](http://www.directionsmag.com/article.php?article_id=372&trv=1).
- [26] F. Li, C. Nicopoulos, T. Richardson, Y. Xie, V. Narayanan, and M. Kandemir, "Design and Management of 3D Chip Multiprocessors Using Network-in-Memory," in *Proc. 33rd Intl. Symp. Comp. Arch.*, June 2006, pp. 130 – 141.
- [27] M. B. Taylor, J. Kim, J. Miller, D. Wentzloff, F. Ghodrati, B. Greenwald, H. Hoffman, P. Johnson, J. W. Lee, W. Lee, A. Ma, A. Saraf, M. Seneski, N. Shnidman, V. Strumpfen, M. Frank, S. Amarasinghe, and A. Agarwal, "The Raw Microprocessor: A Computational Fabric for Software Circuits and General-Purpose Programs," *IEEE Micro*, pp. 25 – 35, March-April 2002.
- [28] M. B. Taylor, W. Lee, J. Miller, D. Wentzloff, I. Bratt, B. Greenwald, H. Hoffman, P. Johnson, J. Kim, J. Psota, A. Saraf, N. Shnidman, V. Strumpfen, M. Frank, S. Amarasinghe, and A. Agarwal, "Evaluation of the Raw Microprocessor: An Exposed-Wire-Delay Architecture for ILP and Streams," in *Proc. 31st Intl. Symp. Comp. Arch.*, June 2004, pp. 2 – 13.
- [29] <http://www-03.ibm.com/press/us/en/pressrelease/21278.wss>.
- [30] <http://www.top500.org/>.
- [31] A. Gara, M. A. Blumrich, D. Chen, G. L.-T. Chiu, P. Coteus, M. E. Giampapa, R. A. Haring, P. Heidelberger, D. Hoenicke, G. V. Kopcsay, T. A. Liebsch, M. Ohmacht, B. D. Steinmacher-Burow, T. Takken, and P. Vranas, "Overview of the Blue Gene/L System Architecture," *IBM J. of Res. and Dev.*, vol. 49, No. 2/3, pp. 195 – 212, March/May 2005.
- [32] <http://www.gizmag.com/go/7511/>.
- [33] W. C. Swope, "Deep Computing for the Life Sciences," *IBM Sys. J.*, vol. 40, No. 2, pp. 248 – 262, 2001.

