



Energy-Efficient Abundant-Data Computing: The N3XT 1,000×

Mohamed M. Sabry Aly, Mingyu Gao, Gage Hills, Chi-Shuen Lee, Greg Pitner, Max M. Shulaker, Tony F. Wu, and Mehdi Asheghi, Stanford University

Jeff Bokor, University of California, Berkeley

Franz Franchetti, Carnegie Mellon University

Kenneth E. Goodson and Christos Kozyrakis, Stanford University

Igor Markov, University of Michigan, Ann Arbor

Kunle Olukotun, Stanford University

Larry Pileggi, Carnegie Mellon University

Eric Pop, Stanford University

Jan Rabaey, University of California, Berkeley

Christopher Ré, H.-S. Philip Wong, and Subhasish Mitra, Stanford University

Next-generation information technologies will process unprecedented amounts of loosely structured data that overwhelm existing computing systems. N3XT improves the energy efficiency of abundant-data applications 1,000-fold by using new logic and memory technologies, 3D integration with fine-grained connectivity, and new architectures for computation immersed in memory.

The rising demand for high-performance IT services with human-like interfaces is driving the quest for the next generation of energy-efficient computers. These computers will operate on abundant data that can be highly unstructured and often streamed in terabytes. Abundant-data workloads arise from social networks, e-commerce transactions, genome sequences, and multimedia analytics. Within 10 years, trillions of sensors will be connected to the Internet, creating a massive data deluge that could overwhelm communication bandwidths. Computers must be able to process, understand, classify, and organize relevant data in real time and in an energy- and cost-efficient manner.

The slowdown of silicon CMOS (Dennard) scaling has prompted comprehensive research on faster, more energy-efficient switches. However, better switches alone will not deliver the necessary leaps in performance. In particular, abundant-data applications expose gross inefficiencies in traditional architectures, where poor locality leads to excessive cache misses, causing massive and slow off-chip traffic to pin-limited DRAMs that face their own scaling challenges. Thus, only small fractions of time and energy of the system are responsible for computation itself, presenting an opportunity for major improvements.

N3XT: AN END-TO-END APPROACH

Our Nano-Engineered Computing Systems Technology (N3XT) approach capitalizes on several recent nanotechnology breakthroughs (see Figure 1). Instead of focusing solely on improving transistors or memory cells, N3XT adopts an integrated approach for a

new system technology that promises to breathe new life into computing. Key N3XT components include the following:

- › High-performance and energy-efficient field-effect transistors (FETs) based on atomic-scale nanomaterials, such as 1D carbon nanotubes (CNTs) and 2D layered semiconductors.
- › Massive amounts of nonvolatile storage such as low-voltage resistive RAM (RRAM) and magnetoresistive memories such as spin-transfer torque magnetic RAM (STT-MRAM). These diverse technologies offer complementary tradeoffs among high density, quick access, long data retention, and read/write endurance. Their advantages can be successfully utilized and their drawbacks avoided through a carefully designed memory hierarchy and tight

enabled by low-temperature layer transfer techniques. This unique approach decouples high-temperature nanomaterial synthesis (to achieve high-quality materials) from low-temperature monolithic 3D integration.

- › Embedded cooling technologies targeting a range of application domains (for example, handheld versus servers) to overcome power density challenges. Examples include conduction using 2D materials, management of thermal transients based on phase change, and convective copper nanomesh structures connected to chip periphery microfluidics.
- › New microarchitectures and system runtimes for scalable computation immersed in memory that lead to massive amounts of active data, enabled by the above technology components and

N3XT PROMISES MAJOR ENERGY-DELAY PRODUCT BENEFITS FOR WIDE-RANGING APPLICATIONS, ESPECIALLY ABUNDANT-DATA WORKLOADS.

integration with computation units.

- › Fine-grained (monolithic, for example) 3D integration of computing and memory elements with ultradense connectivity between layers. Such fine-grained monolithic 3D integration is natural to the N3XT transistor and memory technologies,

their fine-grained integration. Cross-layer resilience techniques overcome yield and reliability challenges.

We demonstrate the effectiveness of N3XT by using the system-level *energy-delay product* (EDP) metric—the product of a software program's total energy consumption and total

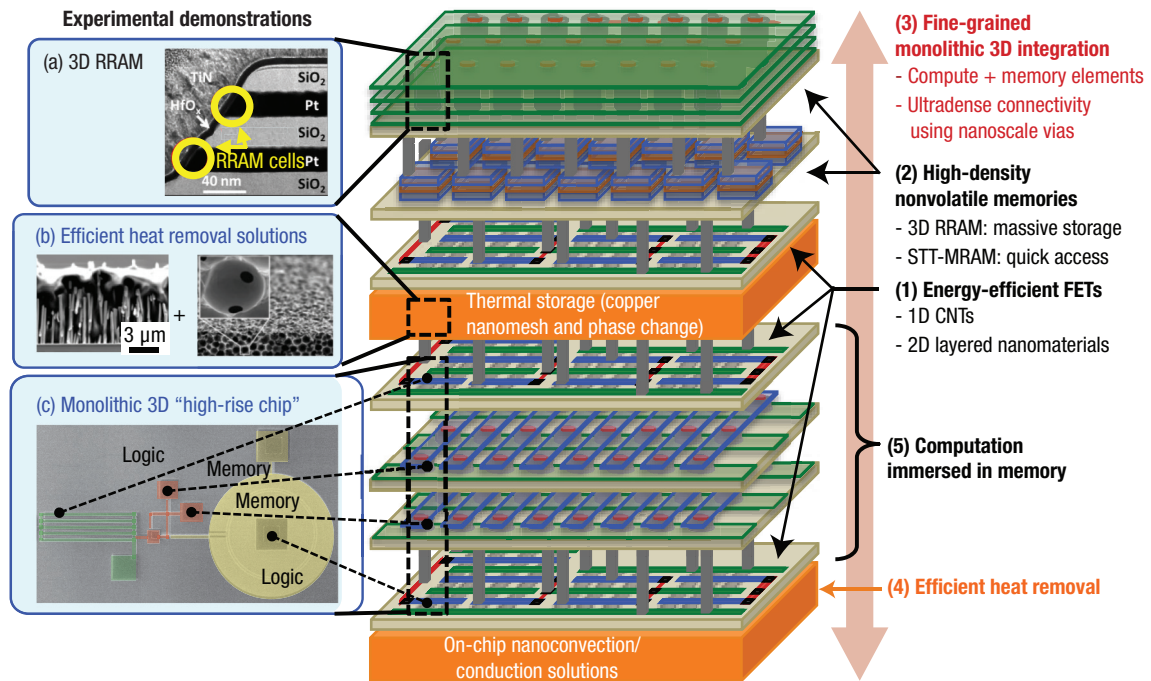


FIGURE 1. Monolithically integrated 3D system enabled by Nano-Engineered Computing Systems Technology (N3XT). On the right are the five key N3XT components. On the left are images of experimental technology demonstrations: (a) transmission electron microscopy (TEM) of a 3D resistive RAM (RRAM) for massive storage, (b) scanning electron microscopy (SEM) of nanostructured materials for efficient heat removal (left: microscale capillary advection; right: copper nanomesh with phase-change thermal storage), and (c) SEM of a monolithic 3D chip for high-performance and energy-efficient computation. CNTs: carbon nanotubes, FETs: field-effect transistors, and STT-MRAM: spin-transfer torque magnetic RAM.

execution time—subject to power density constraints. Given that speed can be traded for energy and vice versa, the EDP metric is important in quantifying computing system performance.¹ To enable new frontiers of abundant-data applications for both mobile devices and the cloud, we target EDP improvements by 1,000×. For traditional multiprocessor workloads, N3XT targets 10×–100× EDP benefits. As we show, N3XT experimental prototypes can be built today.

Such significant benefits are generally rare, and cannot be achieved with evolutionary improvements in architectures, transistors, or memory cells alone—an end-to-end approach such as N3XT is essential. Take, for example, the total delay of a processor pipeline or the total energy of processor cores and memories, where each component must show comparable improvement. N3XT improves each component and finds symbiotic relations to enhance

key performance metrics among components. Additional synergies arise; for example, faster memory accesses cut core idle times, reducing energy consumption and overall execution time. Additional improvements arise from ultradense monolithic 3D integration with fine-grained connectivity and increased memory bandwidth, enabling many concurrent memory accesses and significantly reducing memory access contention; and from nonvolatile memories, which dramatically reduce idle energy consumption and simplify memory access mechanisms.

N3XT TECHNOLOGY FOUNDATIONS

Table 1 summarizes the primary nanotechnologies that form the foundations of N3XT. They work synergistically to overcome the limitations of existing approaches while meeting application-level thermal constraints.

Atomically thin logic devices

N3XT logic devices capitalize on the unique properties of atomic-scale nanomaterials including 1D CNTs and 2D layered semiconductors (for example, black phosphorus, WSe₂). These nanomaterials are ideal for building highly scaled FETs that can deliver large drive currents at low supply voltages. Such FETs exhibit excellent electrostatic control (resulting from atomically thin 1D CNTs with approximately 1-nm diameter and 2D layered semiconductors) while simultaneously achieving excellent carrier transport.

CNTs are hollow cylindrical nanostructures of carbon atoms with exceptional electrical, thermal, and mechanical properties. A carbon nanotube FET (CNFET) consists of multiple CNTs connected in parallel to form the transistor channel (see Figure 2a). CNFETs promise an order-of-magnitude better EDP versus silicon CMOS at the digital

TABLE 1. Nano-Engineered Computing Systems Technology (N3XT) technology foundations.

Technology		Impact on		
		Computation	Storage	Memory access
Field-effect transistors: 1D carbon nanotubes and 2D layered semiconductors		Highly energy-efficient digital systems (including logic and interconnects)	NA	Energy-efficient memory controllers and peripheral circuits
Emerging nonvolatile memory	Spin-transfer torque magnetic RAM	NA	Quick access, high endurance	No refresh; simple control; energy-efficient management by turning off unused banks
	3D resistive RAM	NA	Very high density, long retention	
Fine-grained (monolithic) 3D integration		Computation immersed in memory High computation density for a given footprint	Massive on-chip storage Integration of heterogeneous memory technologies	High bandwidth and low latency
Thermal solutions		High-performance computing on all tiers	Minimized temperature-induced degradation	NA

system level, including interconnect parasitics.² Nevertheless, until recently, imperfections and variations inherent to CNTs (for example, mispositioned CNTs and semiconducting versus metallic CNTs) posed major obstacles and prevented demonstrations of large-scale digital systems.

Considerable progress has been made recently toward full wafer-scale CNFET-based digital systems. A combination of CNFET circuit design and CNT processing techniques, the *imperfection-immune paradigm*, overcomes the challenges of CNT imperfections and variations in a VLSI-compatible manner.³ This enabled the first experimental demonstration of the CNT computer⁴ (see Figure 2c) and, more generally, arbitrary CNFET digital systems. These are the first nano-system demonstrations among various promising emerging nanotechnologies for high-performance and energy-efficient digital systems. Recent work has also demonstrated exceptionally scalable CNFETs with sub-10-nm channel lengths,² complementary n-type and p-type CNFETs,⁵ approaches to overcome contact resistance challenges,⁶ and high-performance CNFETs with CNT densities of >100 CNT/ μm (see Figure 2b).

FETs based on 2D layered semiconductors are presently less advanced than CNFETs, but the potential and challenges are evident.⁷

Sub-nanometer-thin 2D layered semiconductors could enable similar gate scaling as CNFETs, and offer more degrees of freedom to optimize edge versus surface carrier injection at contacts. Two-dimensional layered materials have also been synthesized over large substrates but encounter the challenge of coexisting monolayer and few-layer domains. For both 1D CNTs and 2D semiconductors, the high-temperature synthesis process (to achieve high quality) can be decoupled from low-temperature layer transfer, thus enabling dense monolithic 3D integration.

N3XT applies to other logic switch candidates, such as tunneling FETs or negative-capacitance FETs, as long as they provide high drive currents and low leakage, are scalable to device pitches like CNFETs, and can be integrated in a fine-grained fashion akin to monolithic 3D.

Emerging nonvolatile memories

STT-MRAM (see Figure 2d) can be programmed at low voltages (<0.5 V) with tens of microamperes, can attain read/write access times in few tens of nanoseconds (with potential for another 10 \times speedup), and can offer almost infinite endurance.⁸ Moreover, STT-MRAM cells (approximately 6–20 F²) are substantially smaller than static RAM

(SRAM) cells, resulting in increased memory capacity for the same footprint. A novel device concept for magnetic memory and logic, the m-Cell (an access transistorless spintronic memory cell; see Figure 2e), has demonstrated potential for sub-100-mV operation.⁹ An even lower energy of operation might be enabled using spin-Hall effect switching. These characteristics make emerging magnetic memories promising candidates for ultra-low-power embedded memory layers very close to computing layers (see Figure 1).

For high-capacity storage, metal-oxide RRAM is a leading candidate.⁸ It can be programmed at 1–2 V with currents from nanoamperes to tens of microamperes and 10-year retention. Researchers have achieved endurance through 10¹² cycles and demonstrated a sub-10-nm RRAM device (see Figure 2g). Recently, researchers also demonstrated bit-cost scalable 3D RRAM architectures that are fabricated akin to 3D NAND flash (see Figure 2f). For a future half-pitch of 5 nm, a 128-tier 3D RRAM is projected to yield 64 Tbits, programmed at 1 V and 1 μA current, with 5-ns access time and 10⁹ write cycles, thus enabling ultra-high capacity on-chip storage. Various research groups have achieved these specifications at the

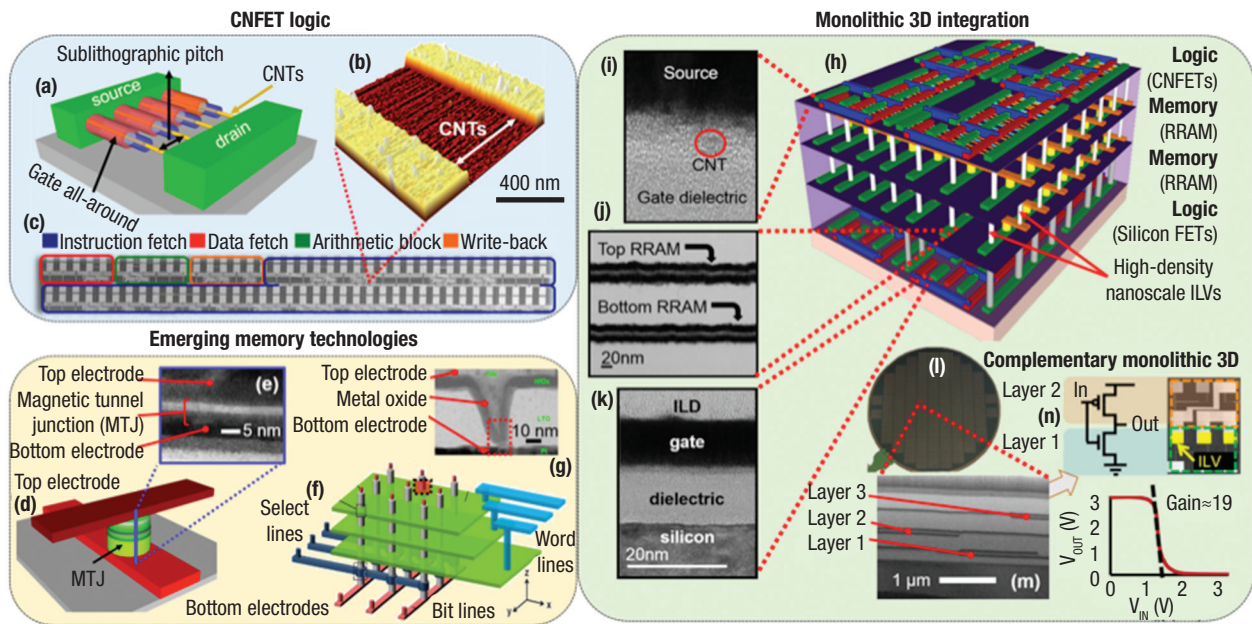


FIGURE 2. N3XT technology foundations. (a) Gate-all-around CNT field-effect transistor (CNFET). (b) Atomic force microscopy (AFM) image of a high-performance CNFET with >100 CNTs/ μm . (c) Turing-complete microprocessor built entirely using CNFETs. (d) Spin-transfer torque magnetic RAM. (e) TEM cross-section image of a fabricated m-Cell (an access transistorless spintronic memory cell). (f) High-density 3D RRAM. (g) TEM cross-section image of an RRAM cell with sub-10-nm feature size. (h) Monolithic 3D integrated circuit (IC) with four vertical layers (logic, memory, memory, and logic). (i) TEM cross-section image of a CNFET on the fourth layer of a monolithic 3D IC described in (h). (j) TEM cross-section image of the middle two layers of RRAM from (h). (k) TEM cross-section image of the bottom layer of silicon FETs from (h). (l) Microscopy image of a wafer with three stacked layers of CNFET-based logic. (m) TEM cross-section image of (l), showing the three vertically stacked layers of CNFET logic. (n) Schematic with an SEM image, and measured waveform of one such fully complementary CNFET 3D logic circuit (p-channel metal-oxide semiconductor CNFET on layer 2 above the n-channel metal-oxide semiconductor CNFET on layer 1).

single-device level. Future challenges include developing appropriate selectors, reducing device variations, and productizing integration technologies for 100-plus 3D RRAM layers.

Fine-grained 3D integration

To achieve the massive EDP benefits offered by N3XT, we must densely interweave computation elements and memory. Such integration is realized by monolithically stacking tiers of logic and memory. Consecutive tiers are connected using nanoscale interlayer vias (ILVs; used for wire routing), which contrasts sharply with traditional 3D integration using through-silicon vias (TSVs). ILVs enable 1,000-fold denser vertical connectivity than TSVs, which is key to greater energy efficiency. To maximize the benefits of monolithic 3D integration, logic and memory layers must be vertically interleaved to build

computing and memory-access circuits adjacent to memory arrays. Thus, memory access latency and energy are reduced. Moreover, device density per unit footprint increases with additional layers despite 2D scaling difficulties.

Monolithic 3D integration requires low-temperature fabrication for the upper tiers ($<400^\circ\text{C}$). This is generally difficult for silicon technologies but comes naturally with N3XT technologies. Monolithic 3D integration with vertically interleaved layers of logic and memory in arbitrary order has been experimentally demonstrated (see Figures 2h-2k), leveraging CNFETs for logic layers and RRAM for memory layers.¹⁰ Importantly, these hardware prototypes have been fabricated directly over a starting silicon substrate, demonstrating that the N3XT approach is compatible with today’s silicon technologies.

Thermal solutions

Effective thermal solutions are essential for reasons ranging from prevention of thermal runaway to maintenance of low skin temperature for mobile and wearable systems. System-level temperature management requires careful electrothermal codesign. Thermal solutions for high-performance computing platforms will require unique micro/nano heat convection and conduction solutions. Embedded cooling technologies might combine solid-state energy storage¹¹ and conduction media, including novel 2D materials.¹² Thermal solutions can also leverage novel micro/nanofluidic cooling, both chip-internal and chip-external, depending on the heat flux densities handled (for example, mobile versus server applications). Advanced convective structures such as copper nanomeshes and tree-like structures (see Figure 1b) can handle

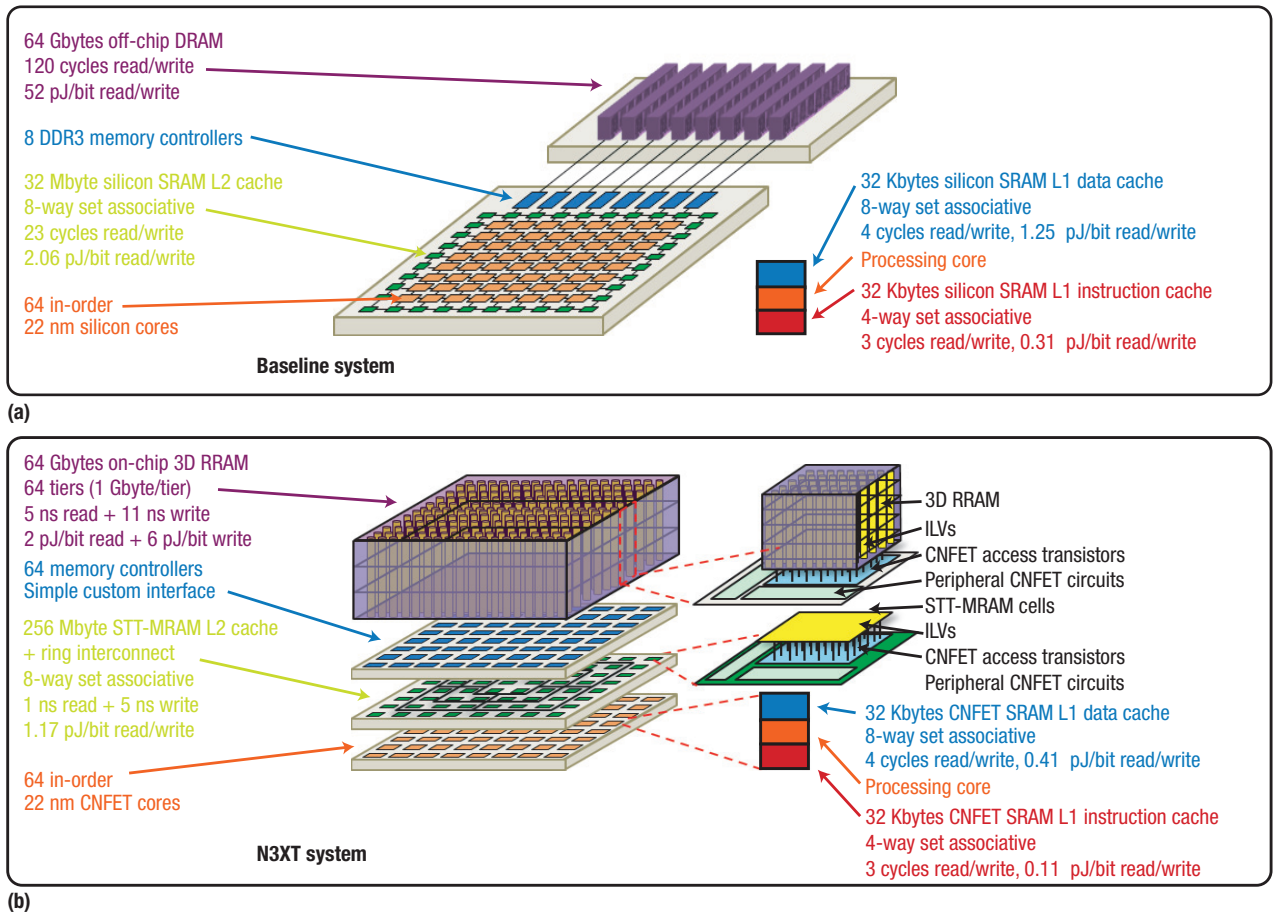


FIGURE 3. Baseline (a) and N3XT system (b) configurations. ILVs: interlayer vias.

heat flux densities from 10 W/cm² to 5 kW/cm². The copper matrix could also encapsulate thermal phase-change materials like paraffin to suppress thermal transients and maintain system temperature constraints.¹¹

N3XT BENEFITS

Our N3XT approach enables massive EDP benefits for a wide range of applications. To demonstrate this, we simulated baseline and N3XT system configurations (see Figure 3). The baseline system is similar to the many-core Intel Xeon Phi. The physical parameters in Figure 3 are derived from industrial data sheets, CNFET SPICE models (<https://nano.stanford.edu/stanford-cnfet2-model>) calibrated using experimental CNFET measurements, as well as energy and delay estimation tools. Ongoing investigations of the key N3XT technologies—for

example, 1D and 2D FETs, STT-MRAM, RRAM, monolithic 3D, and cooling—might produce more accurate values of physical parameters used for our simulations. In the meantime, we used the most accurate values available, validated by hardware experiments where possible.

We performed detailed physical design using place-and-route tools and carefully checked the routability, timing, and power for both implementations. We used Zsim (<https://github.com/s5z/zsim>) and 3D-ICE (<http://esl.epfl.ch/3D-ICE>) for architectural and thermal simulations, respectively. Although the specific technology and architecture selections in Figure 3 allow us to validate the N3XT principles through comprehensive simulations, these selections are not exclusive and the N3XT principles remain general.

Abundant-data multicore workloads

We examined a range of multicore workloads—for example, PARSEC, Powergraph, and IBM Graph analytics benchmarks (<http://systemg.research.ibm.com/analytics.html>)—to thoroughly assess N3XT’s EDP benefits. The observed EDP gains ranged from 10× for computation-bound applications in traditional multicore benchmarks to more than 1,000× for abundant-data applications—our main target. Our analysis relied on uncustomized software implementations and compilers. Further gains might be achieved through careful software and compiler optimizations but at the cost of increased software-development effort.

Consider the traditional PageRank application, a key representative workload for abundant-data applications that is used extensively in Web search

TABLE 2. N3XT energy-delay product (EDP) benefits for simple single-core workloads.

Workload	C-style code	System	Execution time*		Energy consumption		N3XT EDP benefit
			CPU	Memory access	CPU	Memory access	
Computation-dominated	<pre> For (t = 0; t ≤ 1; t +=0.05) { x1 = pow(t, x-1); y1 = pow(1-t, y-1); z+ = x1*y1; } </pre>	Baseline	0.999	0.001	0.960	0.040	9.96×
		N3XT	0.333	0.0006	0.297	0.0031	
Sequential read/writes	<pre> For (i = 0; i < MAX_ITER; i++) y[i] = x[i]; </pre>	Baseline	0.400	0.600	0.280	0.720	66×
		N3XT	0.130	0.060	0.070	0.007	

*Execution time and energy consumption values are normalized to the corresponding total values in the baseline case.

and social networks (benchmarked using Stanford Network Analysis Project’s 16-Gbyte input dataset; <http://snap.stanford.edu>). We used a gather-apply-scatter graph-parallel model with edge-centric streaming implementation, which arranges the edges consecutively in memory and optimizes for sequential memory accesses. For PageRank, N3XT improves EDP by 850×: a simultaneous 23× application speedup and 37× application energy reduction (see Figure 4). The average power density for the N3XT system is 67 W/cm², and the peak temperature is 63°C (versus 65 W/cm² and 61°C for baseline).

The N3XT EDP benefits for PageRank can be further improved to 1,105× (simultaneous 65× application speedup and 17× application energy reduction), but at the cost of increased power density and peak temperature. Hence, N3XT thermal solutions are essential in this context.

To put the N3XT benefits into perspective, TSV-based stacked 3D processor-in-memory (PIM) architecture with 22-nm silicon CMOS, with 8 3D wide I/O interface (www.jedec.org/standards-documents/docs/jesd229) memory channels, provides only 16× EDP benefits for PageRank. We confirmed similar N3XT benefits across other graph-processing workloads from IBM Graph analytics benchmarks.

The observed N3XT EDP benefits arise from the close proximity of computation and memory via monolithic 3D integration, in addition to

energy-efficient logic implemented using CNFETs. Multicore workloads create even more opportunities for additional benefits. For multicore workloads, the processor cores must compete for memory accesses. These access contentions cause major energy-efficiency and performance bottlenecks. N3XT overcomes this by providing much greater memory bandwidth through fine-grained 3D integration; we utilize this opportunity by using 64 memory controllers (see Figure 3) that enable many concurrent memory accesses in the N3XT system. Such concurrency is essential to efficient execution of abundant-data applications with nonlocal data accesses.

We also estimated the impact of I/Os (we focus on off-chip data transfer from a socket to a different processing socket) on N3XT EDP benefits for abundant-data applications. We distributed the input (graph) data evenly across memories in different sockets, connected by the Intel QuickPath Interconnect interface. For N3XT configurations without sufficient on-chip memory, we observed EDP benefits of 613× for PageRank. With growing data volumes and processing rates, proportionately scaling conventional systems’ resources and expanding them in two dimensions can be costly and risky. N3XT offers a strikingly different path to large-scale computation tightly integrated with high-capacity memory: vertically stacking computing

cores and storage units to improve resource density and communication bandwidths. Further research opportunities include effective integration of multiple N3XT chips, thermal management, and corresponding system-architecture optimizations.

Simple single-core workloads

We illustrated N3XT benefits using two simple workloads executed on a single processor core of the baseline and N3XT systems: a computation-dominated kernel (beta function evaluation) and sequential memory accesses (see Table 2). These simple workloads provide insights into the sources of the previously mentioned benefits in the N3XT system. Even for simple workloads executing on a single core, the N3XT system shows significant benefits. The computation-dominated workload spends very little energy or time on memory accesses and improves EDP 10×, mainly owing to CNFETs. For the memory-access workload, memory accesses dominate execution time and energy consumption for the baseline system. However, for the N3XT system, the processor core dominates execution time and energy consumption. This is due to the improved memory system resulting from monolithic integration of 3D RRAM.

SYSTEM-LEVEL CONSIDERATIONS

We now discuss three important N3XT aspects: ensuring acceptable yield and

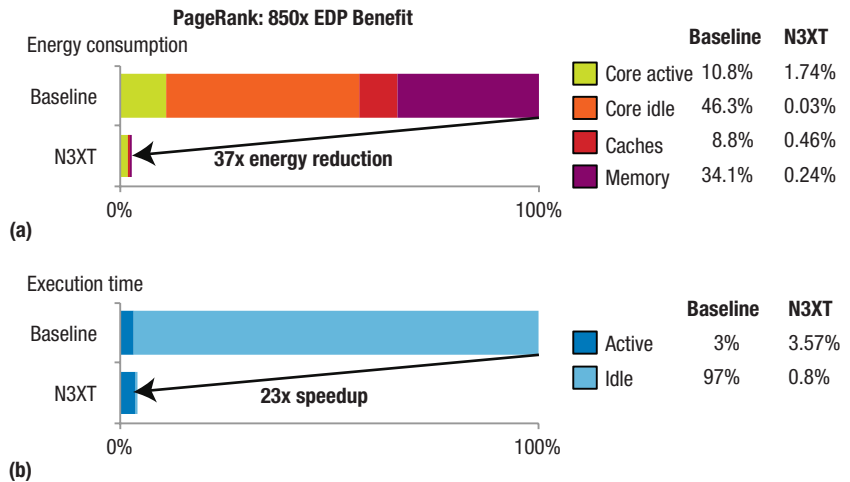


FIGURE 4. PageRank benchmark workload on the baseline and N3XT systems: (a) energy consumption and (b) execution time. N3XT improved energy-delay product (EDP) by 850 \times : a simultaneous 23 \times application speedup and 37 \times application energy reduction.

reliability, positioning N3XT in the context of key energy-efficient computing concepts, and the program-mability and hardware-software co-optimization of the N3XT architecture.

Yield and reliability

Fine-grained 3D integration requires a deep understanding of variability, yield, and reliability, as well as techniques to manage them at the device, circuit, and architecture levels. For example, the imperfection-immune paradigm overcomes substantial imperfections inherent in CNTs.³ Additionally, effectively exploring the interplay between CNT variations and circuit-level energy, delay, noise margin, and functional yield enables co-optimization of CNT processing and CNFET circuit design techniques that overcome CNT variations with <10 percent circuit-level EDP impact.¹³ Massive integration of nonvolatile memories requires similar strategies and new error-correction techniques that are aware of failure modes. Various integration technologies for dense 3D can offer promising yield-improvement opportunities, for example, through intermediate testing of various substrates before integrating them. At the system level, the error-tolerant nature of abundant-data applications and algorithms and the distributed nature of large-scale architectures create powerful opportunities for tolerating hardware failures using techniques at the application, architecture, and circuit levels.

Energy-efficient computing perspective

N3XT is compatible with various techniques for energy-efficient computing. For example, runtime adaptive techniques such as dynamic voltage and

frequency scaling (DVFS) or power gating can also be applied to N3XT systems. Application-specific integrated circuits (ASICs) and accelerator-rich heterogeneous computing architectures can utilize energy-efficient device concepts in N3XT to achieve further benefits. Although hardware specialization through accelerators enhances computing systems' energy efficiency (compared with programmable processors), inadequate data accessibility (for example, few memory access ports or small memory capacity) limits their systemwide effect.¹⁴ Fine-grained accesses to many memory arrays in N3XT overcome such limitations, which in turn boosts the performance benefits of hardware specialization.

Hardware–software co-optimization

Whereas the results reported earlier are for uncustomized implementation of abundant-data applications, careful codesign and co-optimization of N3XT software and hardware systems can enable even higher energy efficiency and performance. The key is achieving this objective at reasonable development costs and time. For example, significant gains are possible through algorithm–architecture codesign that explores a very large space of candidates. Domain-specific

languages (DSLs) with proper compiler support provide effective approaches for such co-optimization,¹⁵ as well as efficient mapping of abundant-data applications onto N3XT hardware architecture. Key elements of such co-optimization include

- › DSLs that provide high-level software abstractions for data transformation (data wrangling), data querying, data feature generation, machine learning, graph analysis, and visualization; and
- › compilers that translate the high-level DSL abstractions into optimized code.

DSL compilers optimize computation and improve memory locality, and the optimized code can then be managed by software and hardware techniques. Such runtime support can manage task distribution, communication, synchronization, power consumption, and fault tolerance in N3XT nanosystems. Along with extensive user-level programmability, software optimizations, and reuse of standard software and hardware modules, DSL compilers also offer automatic microarchitecture selection, as well as comprehensive word-level and some bit-level optimizations.

ABOUT THE AUTHORS

MOHAMED M. SABRY ALY is a postdoctoral scholar in electrical engineering at Stanford University. Contact him at msabry@stanford.edu.

MINGYU GAO is a PhD candidate in electrical engineering at Stanford University. Contact him at mgao12@stanford.edu.

GAGE HILLS is a PhD candidate in electrical engineering at Stanford University. Contact him at ghills@stanford.edu.

CHI-SHUEN LEE is a PhD candidate in electrical engineering at Stanford University. Contact him at chishuen@stanford.edu.

GREG PITNER is a PhD candidate in electrical engineering at Stanford University. Contact him at gpitner@stanford.edu.

MAX M. SHULAKER is a PhD candidate in electrical engineering at Stanford University. Contact him at maxms@stanford.edu.

TONY F. WU is a PhD candidate in electrical engineering at Stanford University. Contact him at tonyfwu@stanford.edu.

MEHDI ASHEGHI is a consulting associate professor of mechanical engineering at Stanford University. Contact him at masheghi@stanford.edu.

JEFF BOKOR is the National Semiconductor Distinguished Professor of Engineering in the Department of Electrical Engineering and Computer Sciences at the University of California, Berkeley. Contact him at jbokor@berkeley.edu.

FRANZ FRANCHETTI is an associate research professor of electrical and computer engineering at Carnegie Mellon University. Contact him at franzf@ece.cmu.edu.

KENNETH E. GOODSON is the Bosch Chairman and the Davies Family Provostial Professor of Mechanical Engineering

at Stanford University. Contact him at goodson@stanford.edu.

CHRISTOS KOZYRAKIS is an associate professor of electrical engineering and computer science at Stanford University. Contact him at kozyraki@stanford.edu.

IGOR MARKOV is a professor of electrical engineering and computer science at the University of Michigan, Ann Arbor. Contact him at imarkov@eecs.umich.edu.

KUNLE OLUKOTUN is the Cadence Design Systems Professor of electrical engineering and computer science at Stanford University. Contact him at kunle@stanford.edu.

LARRY PILEGGI is the Tanoto Professor of Electrical and Computer Engineering at Carnegie Mellon University. Contact him at pileggi@cmu.edu.

ERIC POP is an associate professor of electrical engineering at Stanford University. Contact him at epop@stanford.edu.

JAN RABAEY is the Donald O. Pederson Distinguished Professor of Engineering in the Department of Electrical Engineering and Computer Sciences at the University of California, Berkeley. Contact him at jan_rabaey@berkeley.edu.

CHRISTOPHER RÉ is an assistant professor of computer science at Stanford University. Contact him at chrismre@cs.stanford.edu.

H.-S. PHILIP WONG is the Willard R. and Inez Kerr Bell Professor in the School of Engineering and a professor of electrical engineering at Stanford University. Contact him at hspwong@stanford.edu.

SUBHASISH MITRA is an associate professor of electrical engineering and computer science at Stanford University. Contact him at subh@stanford.edu.

N3XT promises major EDP benefits for wide-ranging applications, especially abundant-data workloads presented by big-data processing and the myriad sensors that produce a massive data deluge. N3XT enables unprecedented computing capabilities. It is a major IT leap and is crucial for addressing several of the National Academy of Engineering's

21st century grand challenges. N3XT is urgent because existing system technologies and architectures have hit major obstacles; as a result, the future of computing faces formidable challenges, especially with the slowdown of traditional integrated circuit scaling.

N3XT is an integrated approach spanning emerging logic switches

and memories, fine-grained 3D integration, cooling solutions, computer architecture, and software. Although significant research is required to realize the massive N3XT benefits we outlined, experimental hardware prototypes and detailed simulations using physical layouts and hardware-calibrated models clearly indicate that N3XT's technical challenges are

tractable. The specific N3XT implementation details presented here are not exclusive, but they can guide the design of more compelling systems.

We primarily focused on von Neumann-style computing platforms to capitalize on the large body of software technologies and their upcoming enhancements. However, we expect that several N3XT features, including energy-efficient logic, massive memory capacity, and densely integrated computation and memory, will significantly influence architectures targeting other computation models, such as brain-inspired architectures. **■**

ACKNOWLEDGMENTS

This work was supported in part by the National Science Foundation, the STAR-Net Systems on Nanoscale Information fabriCs Center (one of six SRC STARnet Centers sponsored by Microelectronics Advanced Research Corp. and DARPA), the Swiss National Science Foundation Early Postdoc.Mobility Fellowship 151965 to Mohamed Sabry Aly, the Hertz/SGF to Max Shulaker, and the Stanford University SystemX Alliance. We also acknowledge IBM Research investigators for their collaboration on CNFET modeling.

REFERENCES

1. R. Gonzalez and M. Horowitz, "Energy Dissipation in General Purpose Microprocessors," *IEEE J. Solid-State Circuits*, vol. 31, no. 9, 1996, pp. 1277-1284.
2. A.D. Franklin et al., "Sub-10 nm Carbon Nanotube Transistor," *Nano Lett.*, vol. 12, no. 2, 2012, pp. 758-762.
3. J. Zhang et al., "Carbon Nanotube Robust Digital VLSI," *IEEE Trans. Computer-Aided Design of Integrated Circuits and Systems*, vol. 31, no. 4, 2012, pp. 453-471.
4. M.M. Shulaker et al., "Carbon Nanotube Computer," *Nature*, vol. 501, 2013, pp. 526-530.
5. L. Suriyasena et al., "VLSI-Compatible Carbon Nanotube Doping Technique with Low Work-Function Metal Oxides," *Nano Lett.*, vol. 14, no. 4, 2014, pp. 1884-1890.
6. Q. Cao et al., "End-Bonded Contacts for Carbon Nanotube Transistors with Low, Size-Independent Resistance," *Science*, vol. 350, no. 6256, 2015, pp. 68-72.
7. G. Fiori et al., "Electronics Based on Two-Dimensional Materials," *Nature Nanotechnology*, vol. 9, 2014, pp. 768-779.
8. H.-S.P. Wong and S. Salahuddin, "Memory Leads the Way to Better Computing," *Nature Nanotechnology*, vol. 10, 2015, pp. 191-194.
9. D.M. Bromberg et al., "All-Magnetic Magnetoresistive Random Access Memory Based on Four Terminal mCell Device," *J. Applied Physics*, vol. 117, 2015; doi:10.1063/1.4913279.
10. M.M. Shulaker et al., "Monolithic 3D Integration of Logic and Memory: Carbon Nanotube FETs, Resistive RAM, and Silicon FETs," *Proc. IEEE Int'l Electron Devices Meeting (IEDM 14)*, 2014, pp. 27.4.1-27.4.4.
11. M. Fuensanta et al., "Thermal Properties of a Novel Nanoencapsulated Phase Change Material for Thermal Energy Storage," *Thermochimica Acta*, vol. 565, 2013, pp. 95-101.
12. E. Pop, V. Varshney, and A.K. Roy, "Thermal Properties of Graphene: Fundamentals and Applications," *MRS Bull.*, vol. 37, no. 12, 2012, pp. 1273-1281.
13. G. Hills et al., "Rapid Co-Optimization of Processing and Circuit Design to Overcome Carbon Nanotube Variations," *IEEE Trans. Computer-Aided Design of Integrated Circuits and Systems*, vol. 34, no. 7, 2015, pp. 1082-1095.
14. J. Hegarty et al., "Darkroom: Compiling High-Level Image Processing Code into Hardware Pipelines," *ACM Trans. Graphics*, vol. 33, no. 4, 2014, article 144.
15. K.J. Brown et al., "A Heterogeneous Parallel Framework for Domain-Specific Languages," *Proc. Int'l Conf. Parallel Architectures and Compilation Techniques (PACT 11)*, 2011, pp. 89-100.



Selected CS articles and columns are also available for free at <http://ComputingNow.computer.org>.