R&D100 WINNERS

IMPACTING GLOBAL TECHNOLOGY

LABORATORY DIRECTED RESEARCH AND DEVELOPMENT

Sandia National Laboratories

R&D100 Winners
Impacting Global Technology

LABORATORY DIRECTED RESEARCH AND DEVELOPMENT

For further information, contact:
Henry R. Westrich
LDRD Program Manager
hrwestr@sandia.gov
505-844-9092

ABOUT THE COVER:
Representative images from the R&D100-winning technologies in this brochure, all with their roots in Sandia Laboratory Directed Research and Development (LDRD).

Introduction

How LDRD Research Leads to an R&D100 Award

Rick Stulen, Vice-President and Chief Technology Officer

Each of the technologies presented in this brochure is an R&D100 Award winner, and furthermore, they have in common at least a portion of their roots in Sandia Laboratory Directed Research and Development (LDRD) funding. Each year, R&D Magazine's top 100 technologies recognize the year's most innovative technological ideas. Ultimately, a winning R&D100 entry has less to do with trappings of recognition, such as patents and corporate partners, as it does with creativity and ingenuity — the public's ability to utilize a uniquely resourceful solution for a national or global problem. Therefore, that an R&D100 award-winning technology has part of its funding base in Sandia's sole discretionary research and development funding source should not be surprising. The LDRD funding process at Sandia is a highly competitive one (10% funding success rate), resulting in a portfolio with greater than normal impact on national security issues. The fact that LDRD funding underlies almost 60% of those R&D100 awards speaks highly of the value of this program to Sandia and to the nation.

Innovations do not simply "happen" in a single magical step. More often, a creative LDRD idea is submitted to break new ground in science or technology relevant to our national security missions. Insights gleaned from the proposed R&D may lead to an "ah-ha" in the mind of a scientist or engineer as s(he) more-thoroughly comprehends a fundamental scientific or engineering principle. Often, the key step leading to innovation occurs when an LDRD project's principal investigator recognizes an opportunity to use the project's new discovery to solve a difficult technical problem identified within a direct program addressing an aspect of Sandia's national security mission. In this fashion, LDRD provides the opportunity to "connect the dots" between a scientific discovery and an innovative solution to a heretofore intractable mission problem. Within an LDRD project, one can test a proposed solution, demonstrate its value, patent it, and voila, an R&D 100 award-winning technology from an LDRD-funded project.

The following five case studies offer specific exemplars of LDRD investments over the last decade spanning a broad range of science and engineering disciplines, each of which was subsequently recognized as an R&D Magazine Top-100 technology. It is important to remember that while winning such an award is certainly gratifying for both institution and researchers, the most-important consideration is that these technologies demonstrate a high level of ingenuity in service to Sandia's national-security missions. All LDRD-funded research is so-positioned; R&D100 winners merely comprise a microcosm of that larger group.

Contents

The racks of individual processors that tend to characterize massively parallel computers.

More-sophisticated Silicon Brains

Sandia's Commitment to Massively Parallel Computing

In an era of challenges of immense complexity, it should come as no surprise that scientists and engineers must increasingly turn to faster, more-powerful computers to assist them both in processing experimental data and in constructing mathematical — and visual — models to allow them to manage this complexity; essentially to convert the complexity of an issue like the physics of materials properties into a form more manageable for the human brain. Silicon brain assists comprehension by carbon-based brain, which reciprocates by finding ways to increase the processing power and efficiency of the silicon brain. Sandia's commitment to improve the performance of its massively parallel (MP) computing resources is well-illustrated by the Laboratories' three R&D100 awards in a single decade, all within the area of software designed to improve the efficiency with which those parallel computing resources are used across different science and engineering regimes.

In a sense, with the advent of so-called "duo" and "quad" machines, many of us are now more familiar with parallel computing. In essence, the hardware accepts commands to initiate multiple jobs (software routines) simultaneously, assigning each to a different processor (chip) or different core within one processor. However with jobs as complex as, for example, modeling the physics of materials, a "single" model may encompass so many separate subroutines — each processing the evolution-in-time a subset of changes in the model, that what appears to be a single job can only be handled by multiple processors, each processing its sub-aspect of the overall job in parallel with other processors. For issues as complex as materials factors, the amount of parallel processing required makes this a job for massively parallel computers like Sandia's Red Storm, Los Alamos' (LANL)'s Roadrunner, or Lawrence Livermore National Laboratory's (LLNL)'s Blue Gene/L. Nanoscience and bio-nanoscience can also require massively parallel

Interior of the Sandia-built C Plant parallel computer, showing its many racks of processors

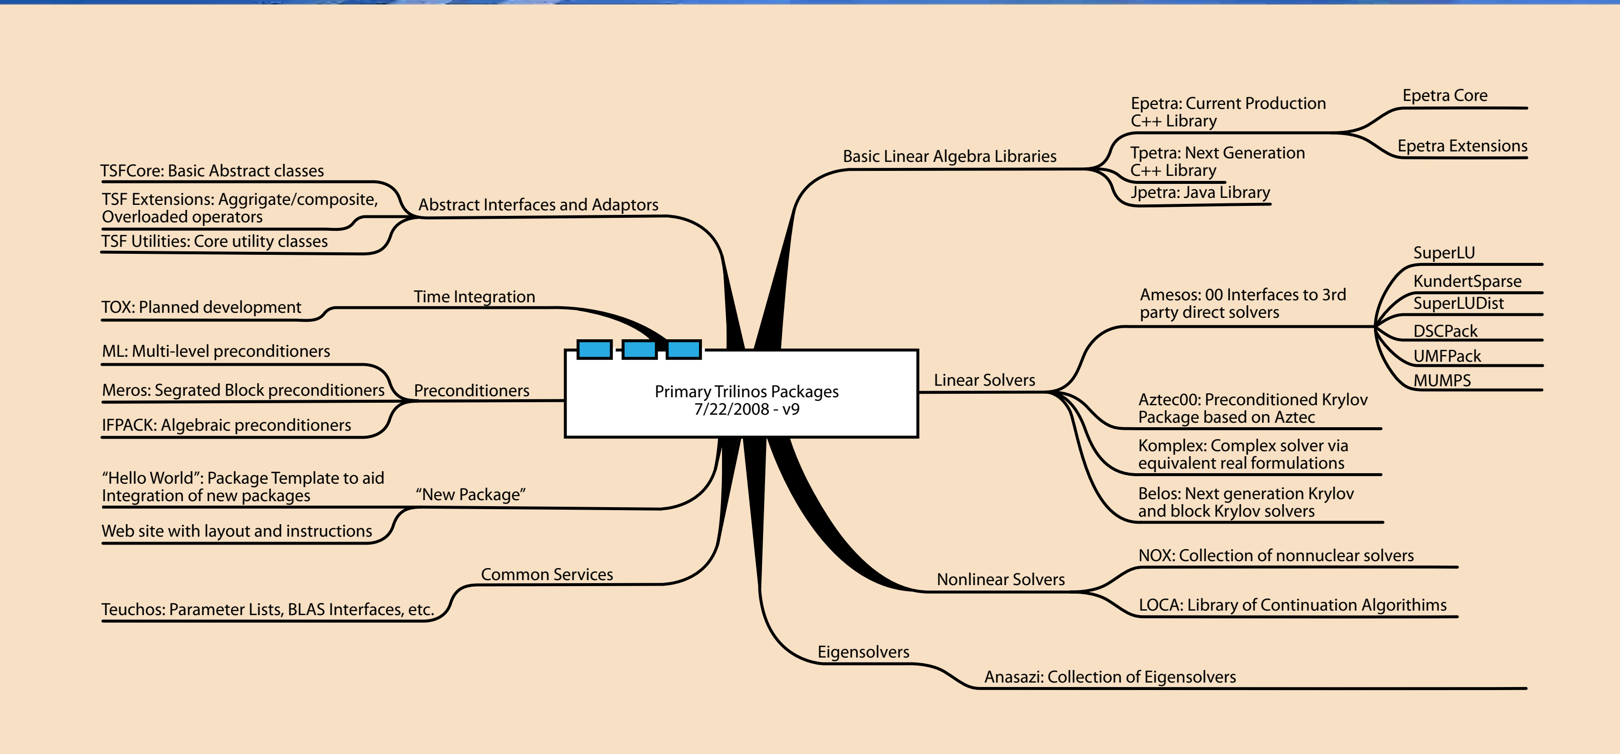

Overview of the many interrelated software solver packages within the Trilinos Library

> *"It can be safely said that Trilinos would probably not exist if not for that initial LDRD."*

computations; for example, the folding of protein molecules into their final three-dimensional molecular shape poses such a nontrivial problem encompassing a myriad of chemical interactions among the different parts of the protein (polypeptide chain[s]), between protein and solvent, and among solvent molecules (usually water) and dissolved salts — among others.

In addition to the obvious investment in increasing the power of such massively parallel machines, that is, by adding more hardware, more processors, there are other, more-subtle, and often less-expensive approaches to improving their performance. Included in these are the three R&D100 Award-winning Sandia initiatives: Aztec (1997), Trilinos (2004) and Compute Process Allocator (CPA) (2006), all discovering ingenious methods to advance and enhance the utility of MP computing without the significant investment costs of adding hardware.

As a general overview, one can view these Sandia approaches to improving the power and reliability of massively parallel computing as taking two complementary paths, both of which globalize the efficient use of massively parallel machines, and both of which originate in LDRD-funded projects. The first path facilitates the task of software developers in writing code to solve specific problems, the second optimizes the way that such software routines and subroutines actually run on these machines from the standpoint of hardware allocation. MP computers have their drawbacks, and unreliability can be one of them. For example, estimates are that in a machine that contains thousands of individual processors, one of these processors may fail as frequently as every 15 minutes. As machines get larger, with more total processors, this problem will persist, if anything, becoming even more of an issue — with more processors available to fail, the average interval between failures may decrease. Hence, solutions that can optimize usage without necessarily adding hardware are quite desirable.

Solver Libraries

A 1997 awardee, Aztec constituted a parallel solver, that is a software package designed to assist in finding the values for unknowns in a system of equations representing some large-scale scientific or engineering challenge, and running on a massively parallel machine. In addition to materials physics, examples of complex, real-world problems of this type would be transport of water-dissolved toxins through the environment, and human physiology studies, such as cardiac irregularities and physical description of electrical instabilities and defibrillation of a heart that has lost its intrinsic pacemaker rhythm. Aztec was originally designed as a solver for the application, MPSalsa, which approached the problem of simulating the dynamics of fluid flow, including applications such as the dynamics of flows occurring within a chemical vapor deposition reactor, (crucial to the process by which, for example, thin semiconductor films are created for computer

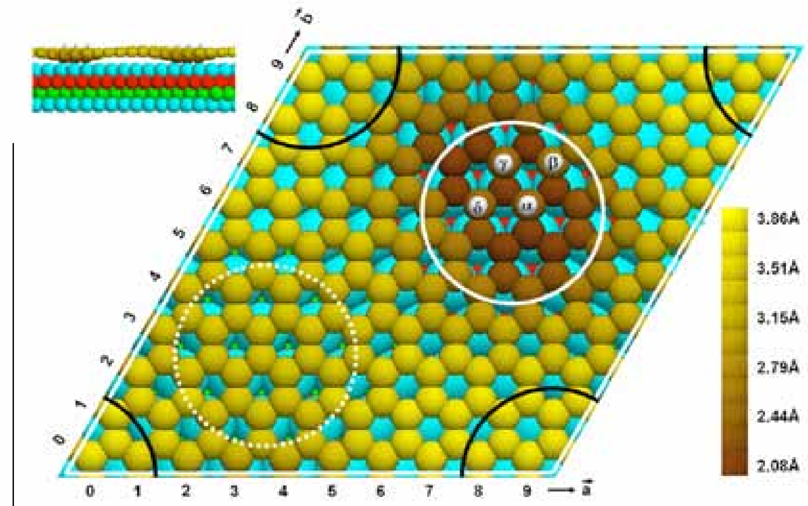

Parallel-computer simulation of the chemistry involved in the adherence of a graphene flake (single graphite layer) into an underlying layer of iridium atoms.

chips). The solver package was subsequently generalized to render it useful in a broader variety of applications.

This trend toward a software library of solvers was continued, refined, and rendered far more powerful in the 2004 R&D100 winner, Trilinos — a library of software whose components or "packages" as they are called — can be adapted by way of application programming interfaces (APIs, quite common in the C++ programming language) to running as solvers for a variety of specific software applications. "In the early 1970s, almost no libraries were used in applications, but over time, we've progressively come to embrace them," according to Trilinos principal developer Mike Heroux.

The general notion, in this context, is that, in writing software for a specific problem to be modeled on a parallel machine, a software developer will necessarily need to craft code that is germane to the specific issue under examination — materials physics or cardiac dynamics, for example — but there are also more-generic aspects of that software construction connected with solving groups of differential equations whose solvers can be found in one or more Trilinos packages. The developer who can make use of those

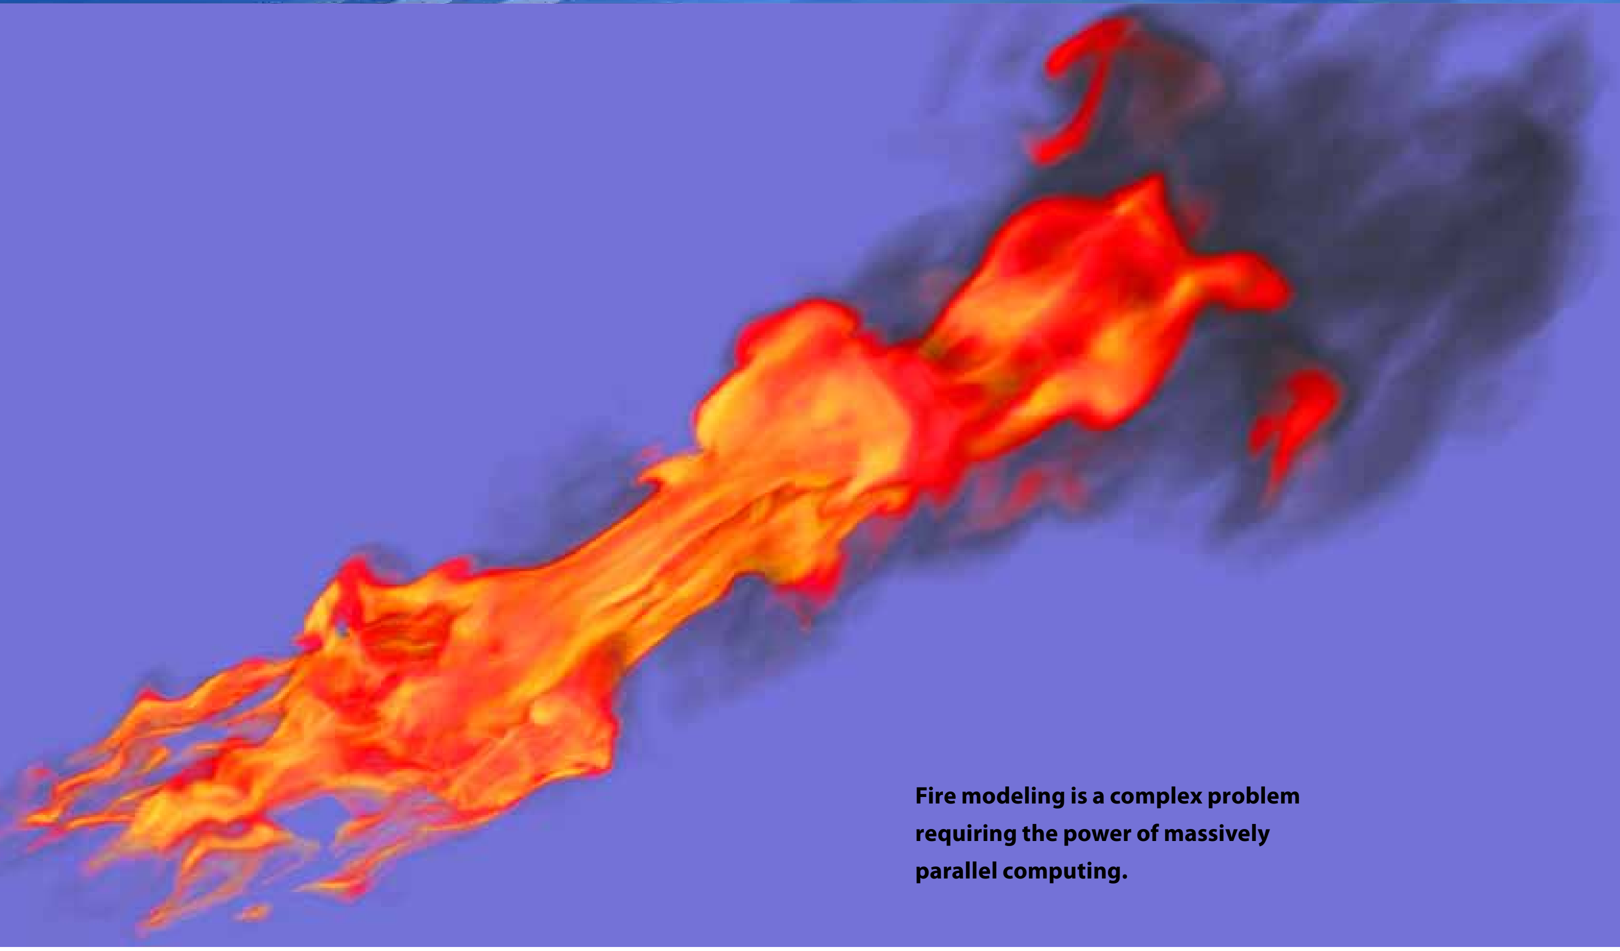

Fire modeling is a complex problem requiring the power of massively parallel computing.

> *“. . . collaboratively applied to industrial problems by companies such as Boeing and Goodyear.”*

generic, broadly applicable Trilinos packages — plugging them into the situation-specific code — has greatly facilitated his/her task. By analogy, your computer might accept a mouse manufactured by several different companies, as long as they all possess the same standard USB-plug interface. In the case of Trilinos, a software package modeling materials and software modeling cardiac rhythms might both have need of a particular type of Trilinos solver package that could be “plugged-in” to each via a standard API. Hence, according to Heroux, Trilinos offers the developer the opportunity to greatly reduce programming time. And as a corollary, in its ingenious way, Trilinos and its predecessor, Aztec, improve the efficiency with which Sandia’s massively parallel computing resources are applied to crucial national security issues, such as stockpile stewardship, climate change, energy security, and others.

Other than the pre-existence of Aztec, a version of which (Aztec 00) ultimately became a package within Trilinos, the developmental activity that spawned Trilinos began in 1998, with a LDRD proposed by Heroux shortly after his arrival at Sandia. Funded for three fiscal years, 1999-2001, the project produced Epetra, described by Heroux as what came to be the most important package in Trilinos, a cornerstone parallel linear algebra solver that provides the data structure for Trilinos but also can be utilized independently of other Trilinos packages. Because it is written in the most stable core of the C++ programming language, Epetra is quite accessible to C and Fortran programmers as well. Its importance as the tangible product of that first LDRD-funded project is underscored by Heroux’s own assessment: “It can be safely said that Trilinos would probably not exist if not for that initial LDRD.” He emphasizes the fact that although the Advanced Simulation and Computing (ASC) Program has served as a key and significant funding source for the ongoing development of Trilinos, particularly its adaptation to new hardware platforms as massively parallel computers become even faster and more powerful (approaching a new goal of sustained petaflop processing), ASC has “rarely had the luxury of high-risk funding.”

This observation succinctly defines for the nation and its taxpayers one key reason for the importance of the LDRD program. For in its highly competitive selection processes, LDRD seeks to fund those projects that show potential for novelty, for breakthroughs in the context of brilliant scientific reasoning. The fact that the route to those breakthroughs may entail significant failure risk, does not prevent LDRD funding, as it might in a more-mature funding setting. For the whole notion of the LDRD program is to encourage ingenuity, and ingenuity and risk are connubial partners in a marriage that can often produce award winners like Trilinos.

In this instance, the partnership successfully yielded not only Epetra, but another solver known as GOMA, a multiphysics fluid dynamics solver that eventually evolved into

The C Plant Computer, built at Sandia from commodity parts.

ARIA, a molecular dynamics code that assists scientists and engineers in understanding the interactions among the particles in fluids, a key issue in nanoscience. A second Trilinos-associated LDRD, which ran from 2001 through 2003, broached the issue of data repartitioning in parallel solvers, combining a set of data repartitioning tools in a package known as ZOLTAN and implementing it within Trilinos. Zoltan rebalanced the workloads of the different processors in a parallel simulation, more or less “on the fly” as processor requirements changed during the course of a simulation. In this same vein, the 2003 R&D 100 application, leading to Trilinos’ 2004 award was a logical outgrowth of the tremendous teaming inherent in the breadth of Trilinos, its many packages capable of operating either independently in application to a specific scientific domain requiring a set of solvers or interdependently when the complexity of a system being modeled requires such scope. Appropriately, the award included over 20 co-developers of the software, who, labored in parallel to develop Trilinos’ astounding diversity of solvers.

> *“In the early 1970s, almost no libraries were used in applications, but over time, we’ve progressively come to embrace them.”*

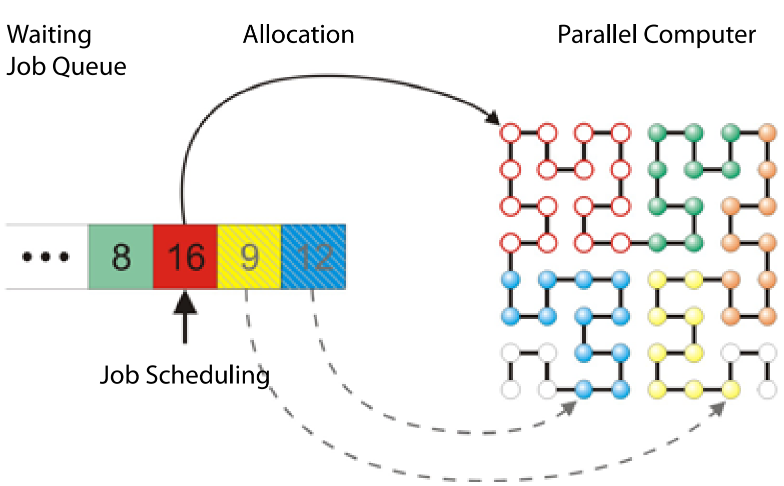

Schematic representation of CPA’s allocation of processors in assigning different jobs to a parallel processing machine.

Although the developers were confident in their entry’s quality, they were also somewhat surprised by the win, given the fact that, historically, the frequency of R&D 100 awards given to software libraries has been low, and the market impact of Trilinos is greatly tempered by the fact that the library was released in 2003 into the public domain and continues to be open source to this day. It hence represents a true collaborative effort leading to a great service to the national and worldwide scientific community in its attempts to model complex systems that require the power of massively parallel computing and the versatility of an enormous and adaptable solver library. For example, Trilinos has been essential to Xyce™, an initiative that models electrical circuitry at Sandia, and at the same time, has been collaboratively applied to industrial problems by companies such as Boeing and Goodyear Tire and Rubber. And development is ongoing: from the initial vision of three solver packages (hence the “tri” in the name), Trilinos now boasts 45 packages, and developers are looking ahead toward the next generation of massively parallel hardware. In yet another affirmation of this body of work, Xyce itself garnered a 2008 R&D100 Award.

Optimizing Hardware Allocation

In the context of the ongoing growth in hardware power, Compute Process Allocator (CPA), derived from a Sandia hardware initiative known as CPlant, captured a 2006 R&D 100 award for its ingenuity in hardware optimization. For as the hardware becomes more complex, the rate of individual processor failure will be unlikely to improve, and may, in fact, become even more of an issue. Based on his doctoral research at the University of California and simulation work, some of it conducted by IBM, CPA principal developer Vitus Leung had arrived at Sandia with the conviction that it might be possible to use

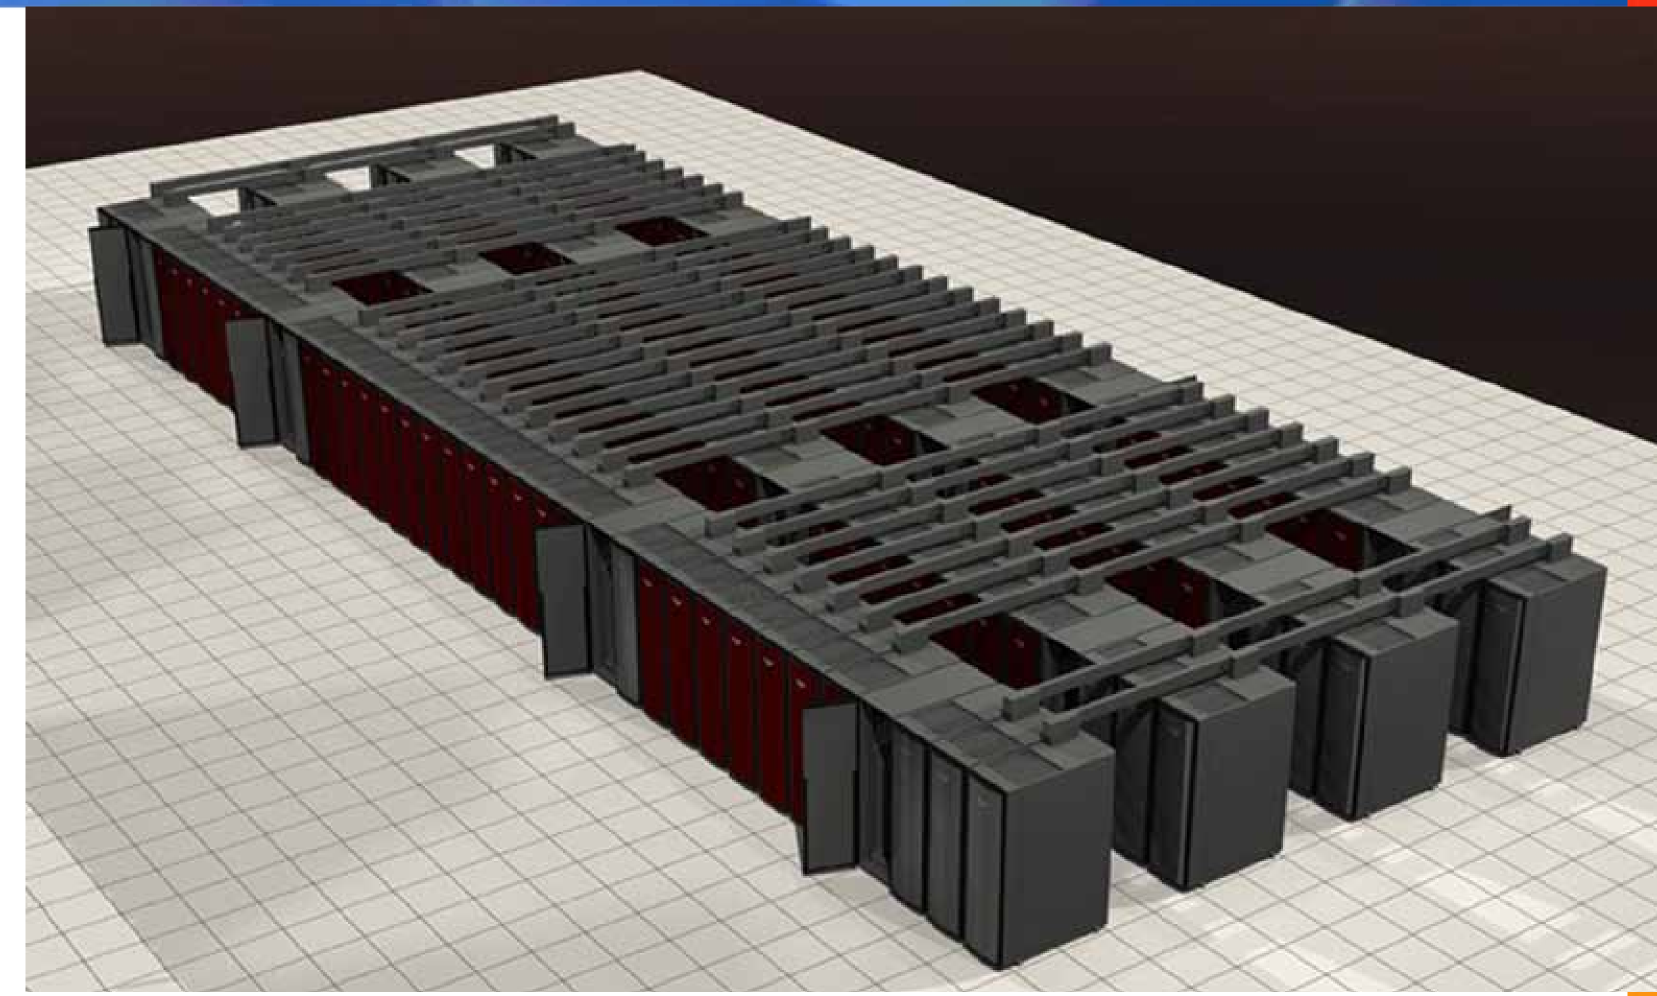

Artists’ drawing of Sandia’s Red Storm Computer, currently offering about 13,000 parallel processors.

algorithm theory to improve the throughput of jobs runs on massively parallel machines. The problem, Leung believes, was not only under-recognized at the time he began his investigations, but is also perhaps still under-appreciated, today. LDRD support spanning the years 2000-2003 funded the development of system software for the experimental CPlant computer, a parallel processing machine built at Sandia from commodity parts; and it was CPlant that strongly suggested the processing bottleneck deriving from inefficient processor allocation to multiple jobs running on the machine. Phased into the CPlant system software in 2002, conference presentation of the LDRD’s final results brought university collaborations, one of the essential aspects of LDRD-funded projects that is sometime under-appreciated.

Based on Sandia’s close affiliation with Cray Inc. in the implementation of Sandia’s Red Storm supercomputer, a Cray XT3 hardware platform, Sandia staff applied CPlant lessons to the Cray XT3 hardware in the form of CPA, a collaboration with computer scientists at the University of Illinois, Urbana and the State University of New York, Stony Brook. This team, which was ultimately honored by the R&D100 award, utilized mathematical algorithms to optimize processor allocation to given jobs simultaneously running on the hardware. Specifically, CPA uses both a space-filling curve and span minimization to ensure that it optimizes the allocation of processor running a particular job, both for processors that are physically contiguous in the machine and for those that are non-contiguous. The developers demonstrated that this combination improved system throughput by 23%, that is, 23% more jobs could be simultaneously run during a given time period. This ultimately represents not only a facilitation of research efforts requiring computer time, but also a

> “. . . discovering ingenious methods to advance and enhance the utility of MP computing without the significant investment costs of adding hardware.”

cost-savings, as well, given the high operating cost of such hardware from both electrical energy and staff-hour perspectives.

Although CPA is licensed to Cray, Inc. as part of the operating system for its XT3/XT4 architectures, it is actually more versatile, and can also be used to manage processor allocation in systems built by Dell, Hewlett Packard, IBM, and Silicon Graphics. It has impacted work in global climate change, nanoscience, astrophysics, and military applications.

A global overview of the past decade reveals that maximizing efficiency of resource utilization has been — and continues to be — an overarching theme in Sandia's multiply award-winning commitment to massively parallel computing. From LDRD seed funding, through international third-party recognition by the prestigious and recognizable R&D 100 award through release of software into the public domain, Sandia's commitment to national and global service via its MP investments is readily observable.

New And Faster Architectures: Ongoing LDRD Investment

As current designers anticipate the arrival of sustained PetaFLOPs (1000 trillion calculations per second) computing, those with an eye to the future see new challenges lurking. When one considers that over the course of only a decade, the amount of hardware necessary to achieve 1 TeraFLOPs (1000 billion calculations per second) has decreased from enough metal to fill a large room to a single chip smaller than the size of a dime, the implications are mind-boggling. For as nanoscience achieves the imprinting of more and more circuitry on semiconductor chips, the amount of sheer computing power in terms of computational operations per second is likely to continue to increase. In systems already composed of 10,000 or more processors, supercomputer manufacturers are both adding parallelism in the form of multicore processors and also pursuing hybrid architectures, in which processors initially designed for graphics such as those found in video games become computationally useful in the massively parallel environment, particularly as orchestrators of data movement within the computer.

This combination of multicore and heterogeneous (also known as “hybrid”) architecture approaches to processing power, within an already massively parallel context virtually ensures a continual upsurge into the realm of hundreds of petaFLOPs and beyond, into ExaFLOPs territory. Working with IBM, Los Alamos is pursuing such an architecture with its Roadrunner supercomputer, which, in conjunction with its parallel arrays of AMD Opteron processors, employs IBM's “Cell” processor, a modification of the processor powering the realistic graphics in the SONY Play Station 3. In the case of Sandia's Red Storm, several LDRD projects over the past five years have addressed the architecture issue for the next generation of massively parallel Cray machines.

Sandia developers have clearly recognized that the message passing interface (MPI) utilized with Red Storm — and which has served as a stable programming model for 15 years — is problematic in terms of its ability to scale up to the necessary level. Red Storm is currently a 13,000-processor environment, but the future holds in store a much higher number of parallel processing streams (threads). First, processors are being upscaled to multicore

A technician examines the interior of a cabinet within the Cray Inc. manufactured Red Storm supercomputer.

> “. . . developers face the real possibility . . . machines that can run hundreds of millions, even billions of parallel threads.”

units, where each core can execute a set of software instructions independently, thus making a single processor, itself an operational parallel processor. When one adds the fact that programming languages such as C permit even a single processor core to execute parallel sets of instructions (threads), developers face the real possibility that instead of a machine with 13,000 parallel software executions, they will eventually be faced with machines that can run hundreds of millions, even billions of parallel threads.

Memory Bottlenecks and Thermal Crowding

This incredible processing power brings to the forefront the problem that Roadrunner and other hybrid architectures are designed to address, a problem perhaps best termed the memory bottleneck. In most cases, computations on data can be executed as fast as the data can be moved from remote memory to local memory, and from local memory to their computational processors. A failure to move data with sufficient rapidity means that the processors, super-fast and numerous though they may be, may be starved for data upon which to compute. It is as though we have increased horsepower by adding cylinders to an auto engine, but failed to improve the fuel injection system that moves adequate fuel into the larger and more-powerful engine.

As Intel, Inc. (Red Storm's processor manufacturer) acquires graphics companies, it is also negotiating with Sandia to acquire the rights to LDRD research that is developing network methodologies to deliver data to processors, which would otherwise be “data starved,” that is able to process data faster than the next set of data can be retrieved from memory. There is a firm belief that the new network architecture for message passing (delivering data and instructions to the processors) developed in the current spate of LDRD projects will be incorporated into Cray, Inc.'s next iteration of this supercomputer technology, the XT-5. And in like fashion to IBM's solutions in LANL's Roadrunner, the evolved architecture must address the question of the memory bottleneck, else no amount of increase in processor speed will be effective at sustaining meaningful performance. Currently, the XT-5 employs AMD's Opteron processor, with an on-chip data cache and memory controller that reduces local memory latency.

Meaningful performance is essential to national security, both as Red Storm models Sandia's electrical-component stockpile work and in its implementation of models for energy security and climate. Manager Sudip Dosanjh projects that the model of on-chip memory, addressed in several LDRD projects will hopefully contribute to even more-robust solutions, as Red Storm's next-generation designers address the memory-bottleneck issue. And since on-chip memory will likely be of limited size, other LDRD projects, addressing message passing on a network level must solve the issue of message passing and parsing from larger central memory to on-chip memory, such that processors and their multiple cores and threads will continually have data upon which to act.

Finally, neither sheer processing power nor memory-bottleneck resolution solve another daunting problem that developers refer to as the “heat-dissipation issue.” Already consuming huge amounts of electrical power, increasing the number of processors and cores will require an even higher commitment of electricity. And with the ever more-crowded on-chip arrays — nanotechnology able to fit more circuitry per chip area — the flow of large quantities of electricity through such crowded environs generates heat, even in relatively efficient circuitry (a consequence of the Second Law of Thermodynamics). Drawing heat away from the chips at an adequate rate is crucial to maintaining the local temperature within certain limits such that performance is not compromised; more-efficient cooling therefore becomes an issue as multicore processors become the rule. (For an allusion to one Sandian solution, see the article on Mode-filtered Fiber Amplifier on page 19 of this publication.)

From developing software libraries that optimize the process of writing application-specific code, to optimizing the utilization of extant processors and job throughput, to solving the memory bottleneck and heat dissipation issues that loom over the next generation of supercomputers, LDRD projects have impacted every aspect of Sandia high performance computing. As such, indirectly, these LDRD projects have offered double-value: for from predicting the margin and uncertainty in weapons performance to developing new materials for space missions, warfighters and solid state electronics, to optimizing cellulose-degrading enzyme structures for bioethanol production, high performance computing is essential to a broad range of national security missions. Therefore, every LDRD dollar invested in optimizing computational performance is also an investment in facilitating optimized solutions to these other national exigencies. ♦

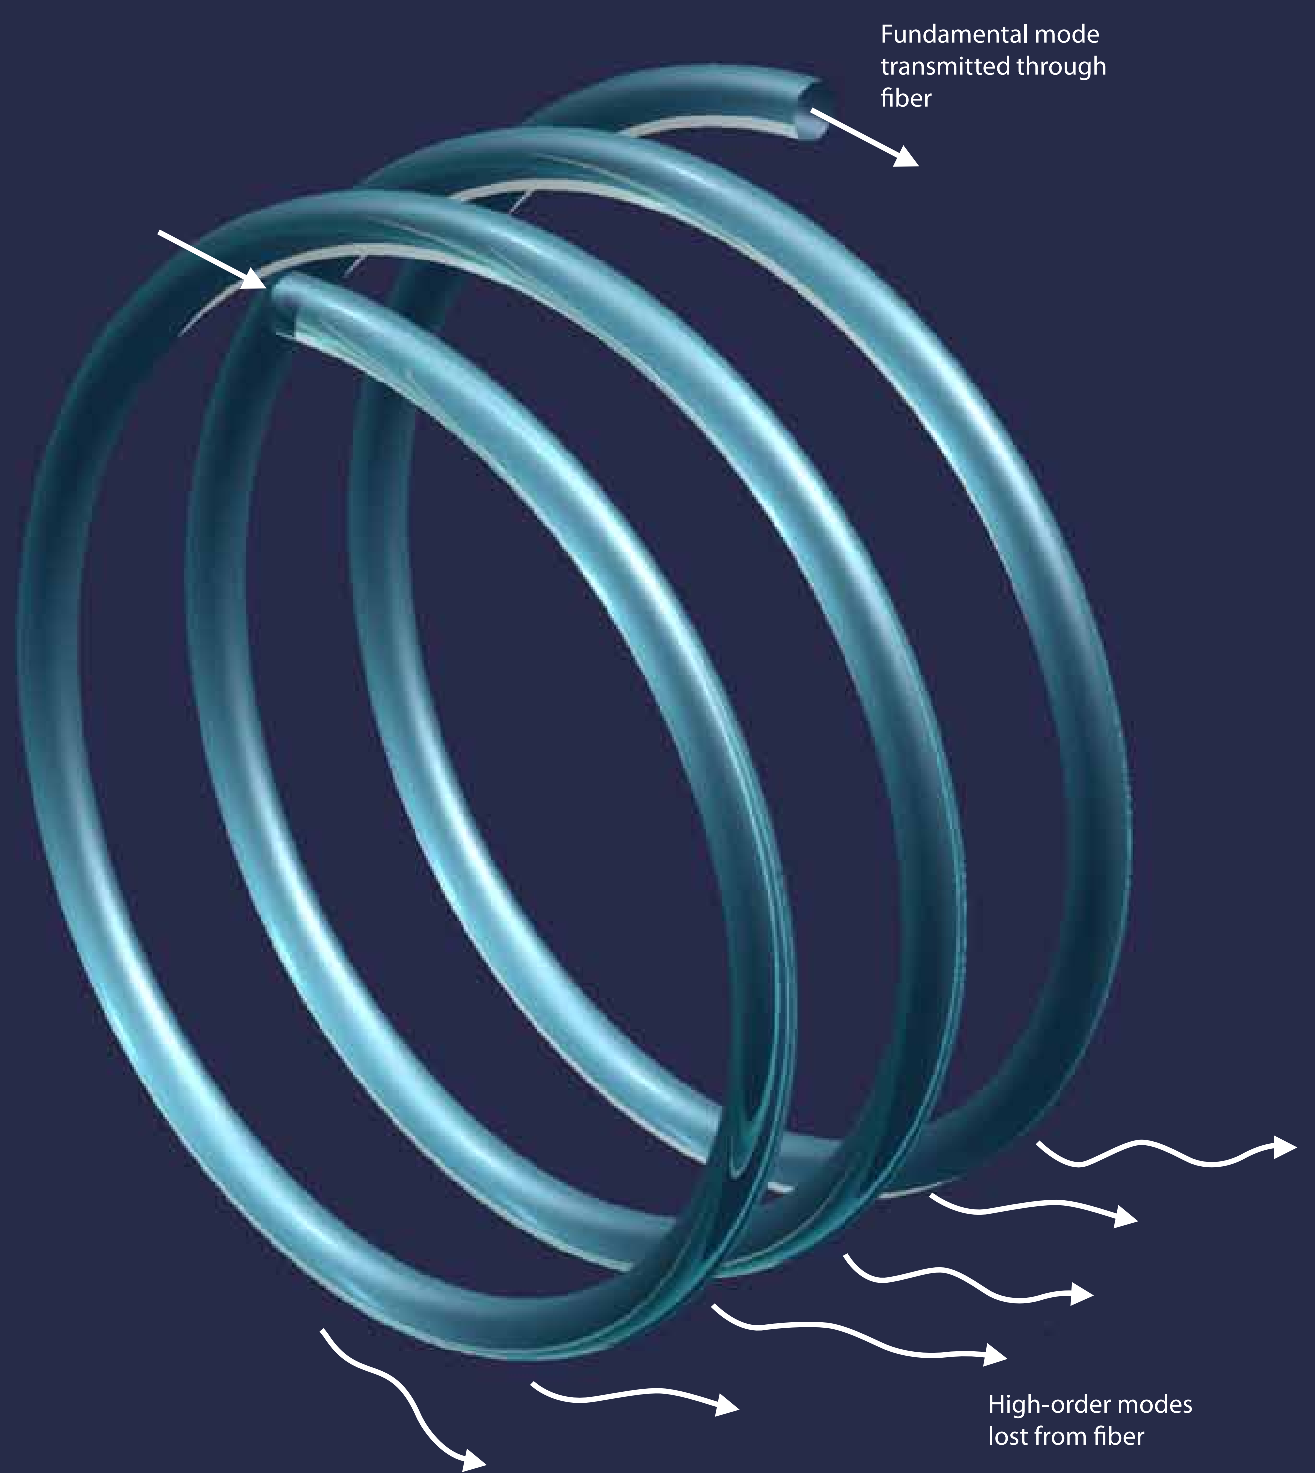

Mode filtering in a coiled fiber amplifier: The undesired higher-order modes are radiated from the side of the fiber along its entire length, while the desired fundamental mode of light is transmitted through the fiber.

Efficient, Dependable Light-power

Mode-filtered Fiber-Laser Amplifiers

For most individuals, the word "laser" has numerous connotations, many of them derived from popular culture, some perhaps inaccurate. "Powerful light" might perhaps be the most common of these, and this connotation is actually reasonably close to reality in many applications in which high-powered lasers play a key role such as drilling and patterning metals. Problematically, however, such high power has traditionally required the use of solid-state crystal lasers, whose requirements are demanding: Consuming enormous quantities of electrical power with low efficiencies, they are quite large, the antithesis of portability, they require elaborate cooling-water systems because of the very large quantity of heat generated, and they are quite sensitive to vibration and other changes in environmental conditions.

An alternative and quite different source of laser light has been known for nearly a half-century in the form of a special type of optical fiber, which instead of functioning as a passive, low-loss transmitter of light-encoded information (e.g., in telecommunications), actually possesses a metal-ion–doped core that can be excited to emit laser light. Such fiber lasers are compact, far more efficient than traditional solid-state lasers (30% or better versus about 5% or lower efficiency), and are far less sensitive to fluctuations in environmental conditions.

The problem has always been power — in order to generate higher-power laser light, the optical fiber must be given a thicker core (and one of altered composition), so that it has more available metal ions to excite for subsequent light emission. Unfortunately, it is far from that simple. Increasing the diameter of the fiber's core also provokes light emission in multiple modes — what can be thought of, for simplicity, as a disorganizing of the emitted laser light, as different standing waves are generated in the fiber and an overall undesirable, low-quality beam is generated. While a thinner fiber core emits less powerfully, it can also be prompted to emit in a single-mode, at the so-called fundamental mode, which is said to be diffraction-limited, governed by the solutions to Maxwell's equation for light propagation. Such a single-mode laser yields a high-quality light beam with a tight focus and small spot size, crucial for applications such as laser drilling in metallurgy, and for projection of the beam over long distances for applications such as free-air communications and remote sensing.

So the issue for fiber-laser researchers had become straightforward, albeit difficult: could a thicker-core fiber laser, which exhibits higher power as well as all other aforementioned fiber-laser advantages, be rendered single-mode, yielding a high-quality, high-power beam? Predicated on years of research at Sandia, much of it LDRD funded, ultimately at the Grand Challenge Level, the 2007 R&D100 award winner dubbed, "Mode-Filtered Fiber Amplifier" described a process for accomplishing that outcome, based on a deceptively simple technique, one

> *"LDRD allowed us to take the time to think coherently to take the longer view. . . think about fiber laser technology in its entirety."*

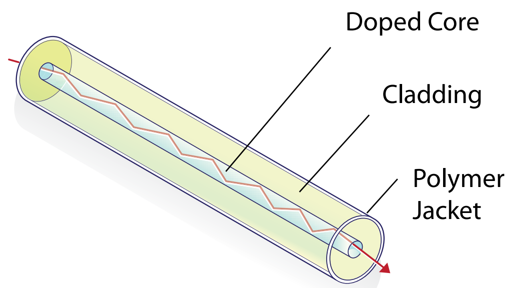

A fiber laser, consisting of a glass core doped with a rare-earth element (typically ytterbium ions in this research) surrounded by a glass cladding.

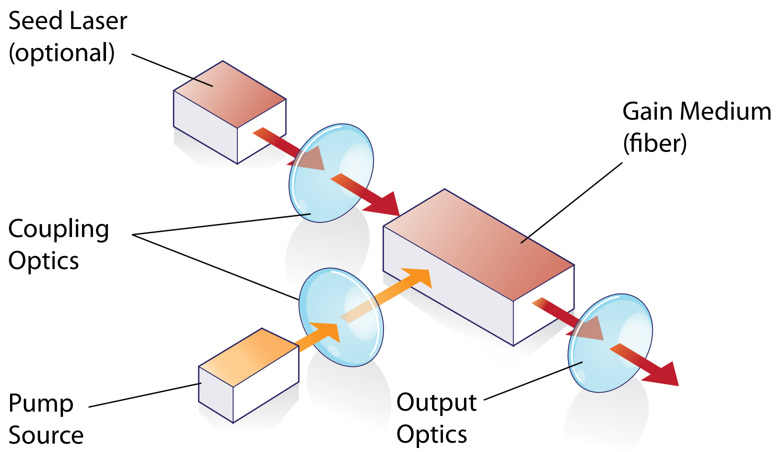

Basic components of a high-power laser system, where the gain medium is an optical fiber; light from the external "pump" source excites the rare-earth (ytterbium) ions, which emit light in specific bands when they drop back to the ground state, thereby providing "optical gain." Seeding the fiber amplifier with a low-power laser (left) is one method of producing coherent laser light.

that, however, presupposes a profound understanding of the physics involved.

A Decade of Research

As is often the case in developmental research, this solution did not arrive overnight. Sandia staff had been investigating applications for fiber laser technology for at least a decade. In the mid-to-late Nineties, a major interest was the monitoring of atmospheric chemistry, for smog-related pollutants such as sulfur dioxide (SO_2) and nitric oxide (NO). The emphasis was the demonstration of the success of a technique, using fiber laser technology, in the detection of an atmospheric species of chemical that was important to scientists and physicians studying both climate change and the impacts of such chemical species on human respiratory health. The technique employed was fiber-laser-based laser-induced fluorescence. In this instance, the lower-power requirements, made it possible to use a single-mode fiber-laser, which conferred the advantages of portability, reliability, efficiency, and lower cost; in general, these initiatives employed a UV excitation wavelength to excite molecular species in question, followed by detection of fluorescence as the detected molecules relaxed to the ground state. NASA funding followed the LDRD support for continuation of the SO_2 research.

This LDRD-supported project dovetailed with another line of research, supported by DOE, to detect propane and methane leak gases in refineries. The attraction for fiber lasers in this initiative was, again, their portability. The National Energy Technology Laboratory (NETL) of DOE affirmed the Department's conviction in the feasibility of fiber-laser-based detection in its support of another project to detect mercury emissions, a growing concern, given the known neurotoxic effects of methyl and ethyl mercury, and the suspicion that certain childhood learning disorders might be linked to mercury exposure.

This track record of high-sensitivity detection of atmospheric aerosols provided a logical motivator for LDRD to again fund a slightly new direction in a project to detect trace amounts of biological warfare agents. If the fiber laser detection (by way of the natural UV absorption and fluorescence of the amino acid tryptophan) could demonstrate high sensitivity, the portability of the instrument could make it an indispensable device in military field operations.

It is not difficult to understand how this extraordinary track record of accomplishment — drawing from both relatively small-sum LDRD funding and external work-for-others funding — might provide a firm foundation for the argument that fiber lasers deserved consideration for Grand Challenge–level support. With the obvious advantages of the technology, a proven track record in detection (and previously in communication), and the remarkable order-of-magnitude improvement in efficiency over conventional laser systems, the only factor lacking in this story was output power — enough power to serve, for example, in metallurgy and in directed energy applications for military systems.

Grand Challenge

Hence, Sandia's LDRD selection committee funded a three-year Grand Challenge (FY2005–FY2007) — with a fourth year later appended — the proposal promising to "provide a fundamentally enabling laser technology with previously unattainable performance characteristics."

Led by former Sandian, Dahv Kliner, the head-on accomplishment of this bold challenge was just as logical a choice for an R&D100 Award, honoring the methodology by which this technology achieved megawatt power while still retaining a high-quality single-mode beam. Additional recognition came in the form of the "editor's choice"

Sandia fiber-laser amplifier system.

award, that is, of the 100 technologies honored in 2007, R&D Magazine's editors chose Mode-filtered Fiber Amplifier as the most enabling technology, giving it top spot among those 100 as the single advance most likely to impact changes in its field.

In what sounds like an "easy fix," the team of Sandia scientists and engineers obviously drew from its years of experience with fiber lasers, as well as from a fundamental and deep understanding of the physics of the issue. Beginning with the design of a novel doped core for the optical fiber, one incorporating the rare earth element, Ytterbium (Yb), and continuing with the design and engineering of pump sources and optics, the key discovery entailed what is dubbed "bend loss." What is lost are the so-called higher-order modes, the modes of light output that essentially contaminate the desired fundamental output mode. This purging of undesired laser modes is accomplished by coiling the fiber at a specific radius, such that as light propagates through the coil, the undesired higher-order modes are radiated and lost from the fiber (see figure on page 18). Only the fundamental mode is transmitted through the coiled fiber and emitted as output, without significant attenuation. This filtering-out of higher-order modes is the price paid for higher power, but

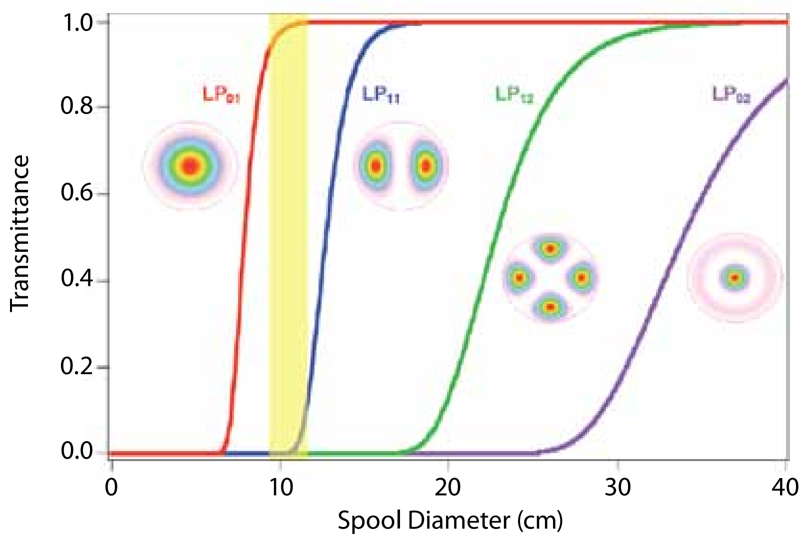

Calculation illustrating mode-filtering. The spatial patterns of four of the 14 modes supported by this 30-μm diameter fiber are shown, the one at far left (LP01, red curve) being the desired fundamental. The yellow bar shows that coiling the fiber on a spool with diameter of 10–12 cm will provide high loss for the other three undesired modes shown, but low loss for the desired fundamental.

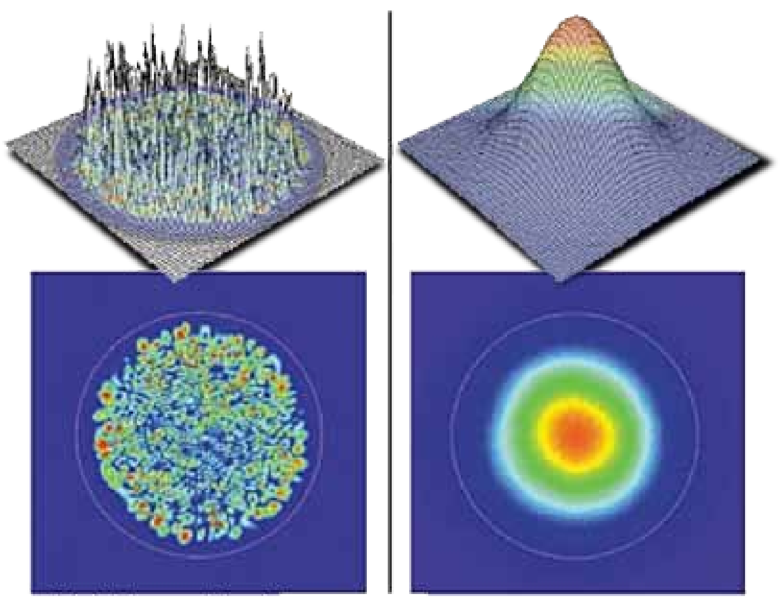

Experimental comparison of the output beam of a 25-μm diameter fiber that is not mode-filtered, versus one that is mode-filtered. The mode-filtered beam is diffraction-limited, stable, and of high quality, while the unfiltered beam is multimode and highly divergent.

remarkably, the fiber laser still achieves an approximately 1 megawatt output with 39% efficiency; that is, of the electrical energy input to drive the generation of light, nearly 40% is emitted as photons, light energy, the remainder as waste heat. Comparing this figure with the approximately 5% (or lower) efficiency of most standard lasers focuses the discussion. The latter instruments must be equipped with elaborate water-based cooling systems to remove the large quantity of waste heat. While cooling is required by the high-power fiber laser, the size and elaborateness of the cooling apparatus is greatly reduced, retaining the unit's overall portability.

As successor to the foundation laid by the team under Dr. Kliner, Jeff Koplow leads the research effort in this area, which includes heat distribution, the extraction of heat from the instrument, and heat dissipation, the transfer of that heat to surrounding air. State of the art — blowing air over the fins — still seems rather crude in the overall scheme of things, and Sandia researchers are examining alternative solutions. This issue is critical because the overall size of the instrument will, to a significant extent, be affected by this cooling apparatus.

Koplow is unreserved in crediting LDRD for its Grand Challenge support. "LDRD allowed us to take the time to think coherently, to take the longer view; not just to do research in short bursts," he asserts. He recalls that, initially, many in the fiber laser research community were somewhat incredulous at the notion that coiling the fiber could be effective at scaling-up power, given that dogma asserted that there was a fundamental barrier in this respect. As it turned out, this dogmatic view, although qualitatively correct, was quantitatively in error because most of this conventional wisdom was derived from experience with very long optical fibers (on the order of tens of kilometers) in the telecommunications industry. Instead, working with much shorter fibers, perhaps 10 meters in length, the Sandia team was able to demonstrate the idea's viability during the period 1997 through 2000. The demonstration, published in the scientific literature in 2000, exhibited its influence in the form of widescale adoption of the technique in the research community. Hence, even before the LDRD Grand Challenge funding began, the idea's viability had been widely accepted.

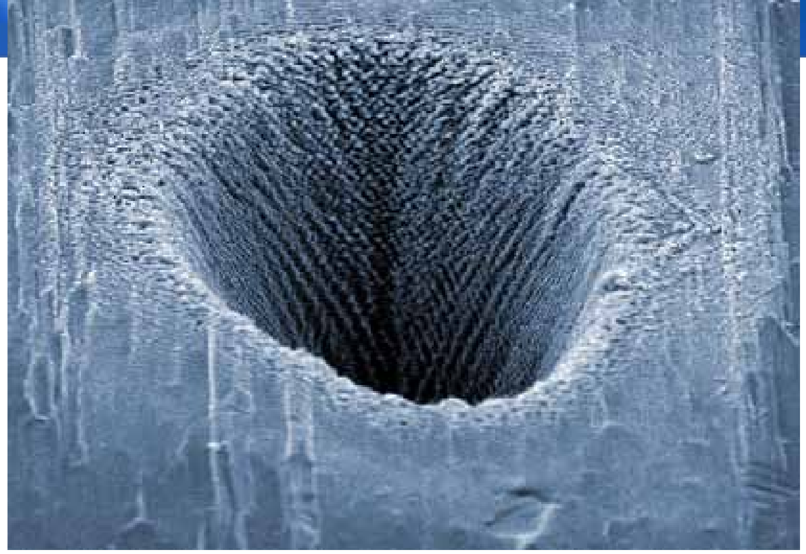

A laser-drilled hole in solid metal, illustrating the type of application requiring the higher power achieved by mode-filtering fiber lasers.

Commercialization and Spin-offs

But what the Grand Challenge funding subsequently permitted was the careful development of all aspects of the discovery, such that a truly ground-breaking engineering solution could be gradually synthesized, by a team, which Koplow characterizes as "an oasis of sanity." "LDRD gave us license to think about fiber laser technology in its entirety," Koplow makes clear. Both the Liekki Corporation of Lohja, Finland and the Nufern Corporation of Connecticut were early believers in the assertion that at least the issue of higher power in a single-mode fiber had been resolved, that in areas such as materials processing — where lasers had failed to fully deliver, because of size, weight, inefficiency, and reliability issues — fiber lasers were now poised to change that outcome. Liekki and Nufern became the first two licensees and co-applicants with Sandia and the Naval Research Laboratory on the R&D100 application. Ultimately, at least three other companies secured licenses, as the technology demonstrated its rational design and global optimization of a whole spate of advances gleaned from a decade of research.

A WFO project for General Motors spotlights the size issue. In a bold step designed to investigate what could be a game-changing technology for internal combustion engines, the project is investigating laser-spark ignition, that is, the substitution of a pulsed laser beam for the traditional electronic spark plug in the ignition of the fuel-air mixture in an engine's cylinders. Such engines would run at a higher compression ratio and a leaner fuel-to-air-mixture, the result being more-efficient use of fuel; generating power with a potentially large increase in mileage. Although the idea of laser ignition of gasoline has been circulating in the combustion community for some time, the pragmatic barrier of laser-system size had essentially prevented a serious attempt at actualization. The small size with adequate power of the mode-filtered fiber laser system has moved this possibility a step closer. This outcome illustrates how significant engineering advances can re-energize R&D communities to pursue potentially revolutionary concepts that had slipped into a state of latency for lack of the appropriate real-world technology.

Numerous patent applications, covering different aspects of the design are, of course, in process, and this technique is clearly a critical piece of intellectual property for Sandia. Hence, this technology remains the center of significant attention. The Grand Challenge had also included modeling of fiber laser performance as an important component of the initiative, and this modeling effort continues.

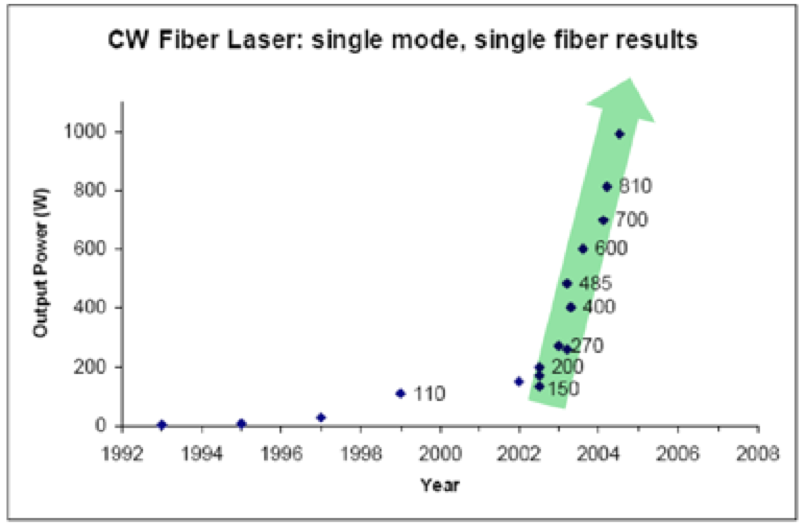

Output power achieved in fiber lasers over the time period 1992 to 2004. Note the very gradual increase in power until the application of mode-filtering, which quickly raised output power to the megawatt level (courtesy Nufern Corporation).

> *". . . of the electrical energy input to drive the generation of light, nearly 40% is emitted as photons."*

Considering the system's portability and the conviction that even higher power output can be achieved, the military applications of the laser are now being examined, as well. In addition to DARPA, several other potential funding sources will be solicited for the purpose of investigating power increases, while retaining single-mode functionality. And with increased power and increased physics understanding comes increased applicability. Highly sensitive bioaerosol detection is still a key application, and free-air communication with portable laser systems is now an even more viable notion given the ability to transmit signals over greater distances.

Globally Relevant Offshoots

Even beyond corporate and military applications is the notion that by allowing the team to step back and examine the problem both bottom-up, from physics first-principles, and top-down, from the global perspective of considering all possible aspects of the design, the LDRD investment set into motion a powerful force of creativity that continues to this day. Once allowed to operate in the realm of science and technology, this creative force cannot help but see challenges in need of attention, challenges that often significantly transcend the scope of the initial problem. In this instance, the issue of thermal management presented itself as a size management consideration.

But as one pondered potential paths to solutions, the creative intellect would also be struck by just how many high-technology arenas were being forced to butt heads with this very issue. Most prominent is the growing concern over the so-called "thermal barrier" to Moore's law in the microchip industry (see the article on massively parallel computing, pages 6–17 of this brochure). As nanotechnology advances allow microchip manufacturers to imprint even more transistors and other components onto their silicon semiconductor chips, that higher density of components transforming electrical energy may mean greater processing power, but it also means a greater thermal load problem. For none of those energy transformations is 100% efficient, and hence, heat is a byproduct of chip operation. Without efficient removal of this heat, and the consequent rise of operating temperatures into realms that affect operational efficiency, the ability to place an even greater density of processing elements onto a silicon wafer becomes irrelevant, by comparison to the ability to remove heat from that chip.

In this fashion, the thermal management issue in fiber laser operation globalizes to several other research areas. Consequently, Koplow has proposed an ingenious solution to the issue of so-called CPU coolers, the small fan units that cool most computers. The proposal, generally seeks to both markedly increase the efficiency with which electricity is used to rotate fans that dissipate heat in traditional CPU coolers, and also to reduce the thermal resistance of the so-called boundary layer, which normally prevents efficient transfer of heat away from the source (be it computer chip, fiber laser or applications in the energy-generation arena). A multimillion-dollar DARPA grant is in the offing.

"This would have never happened in a normal R&D environment," asserts Koplow, that is in an environment that did not include a discretionary R&D fund like LDRD, in which scientists often accused of "playing in sandboxes" are actually provided enough intellectual freedom to think so globally and with such first-principles clarity, that the scientific teams pursue solutions that shake the foundations of traditional thinking, breaking down barriers that sometimes only exist because a cadre of frustrated investigators have agreed to erect them. In this sense, the story of the Mode-filtered Fiber Amplifier team transcends awards, patents, and funding success, going straight to the heart of the type of barrier-breaching persistence and confidence in rationality that engender significant scientific advances in the face of seemingly intractable challenges. ♦

> *". . . the most enabling technology (of 2007) . . . the single advance most likely to impact changes in its field . . ."*

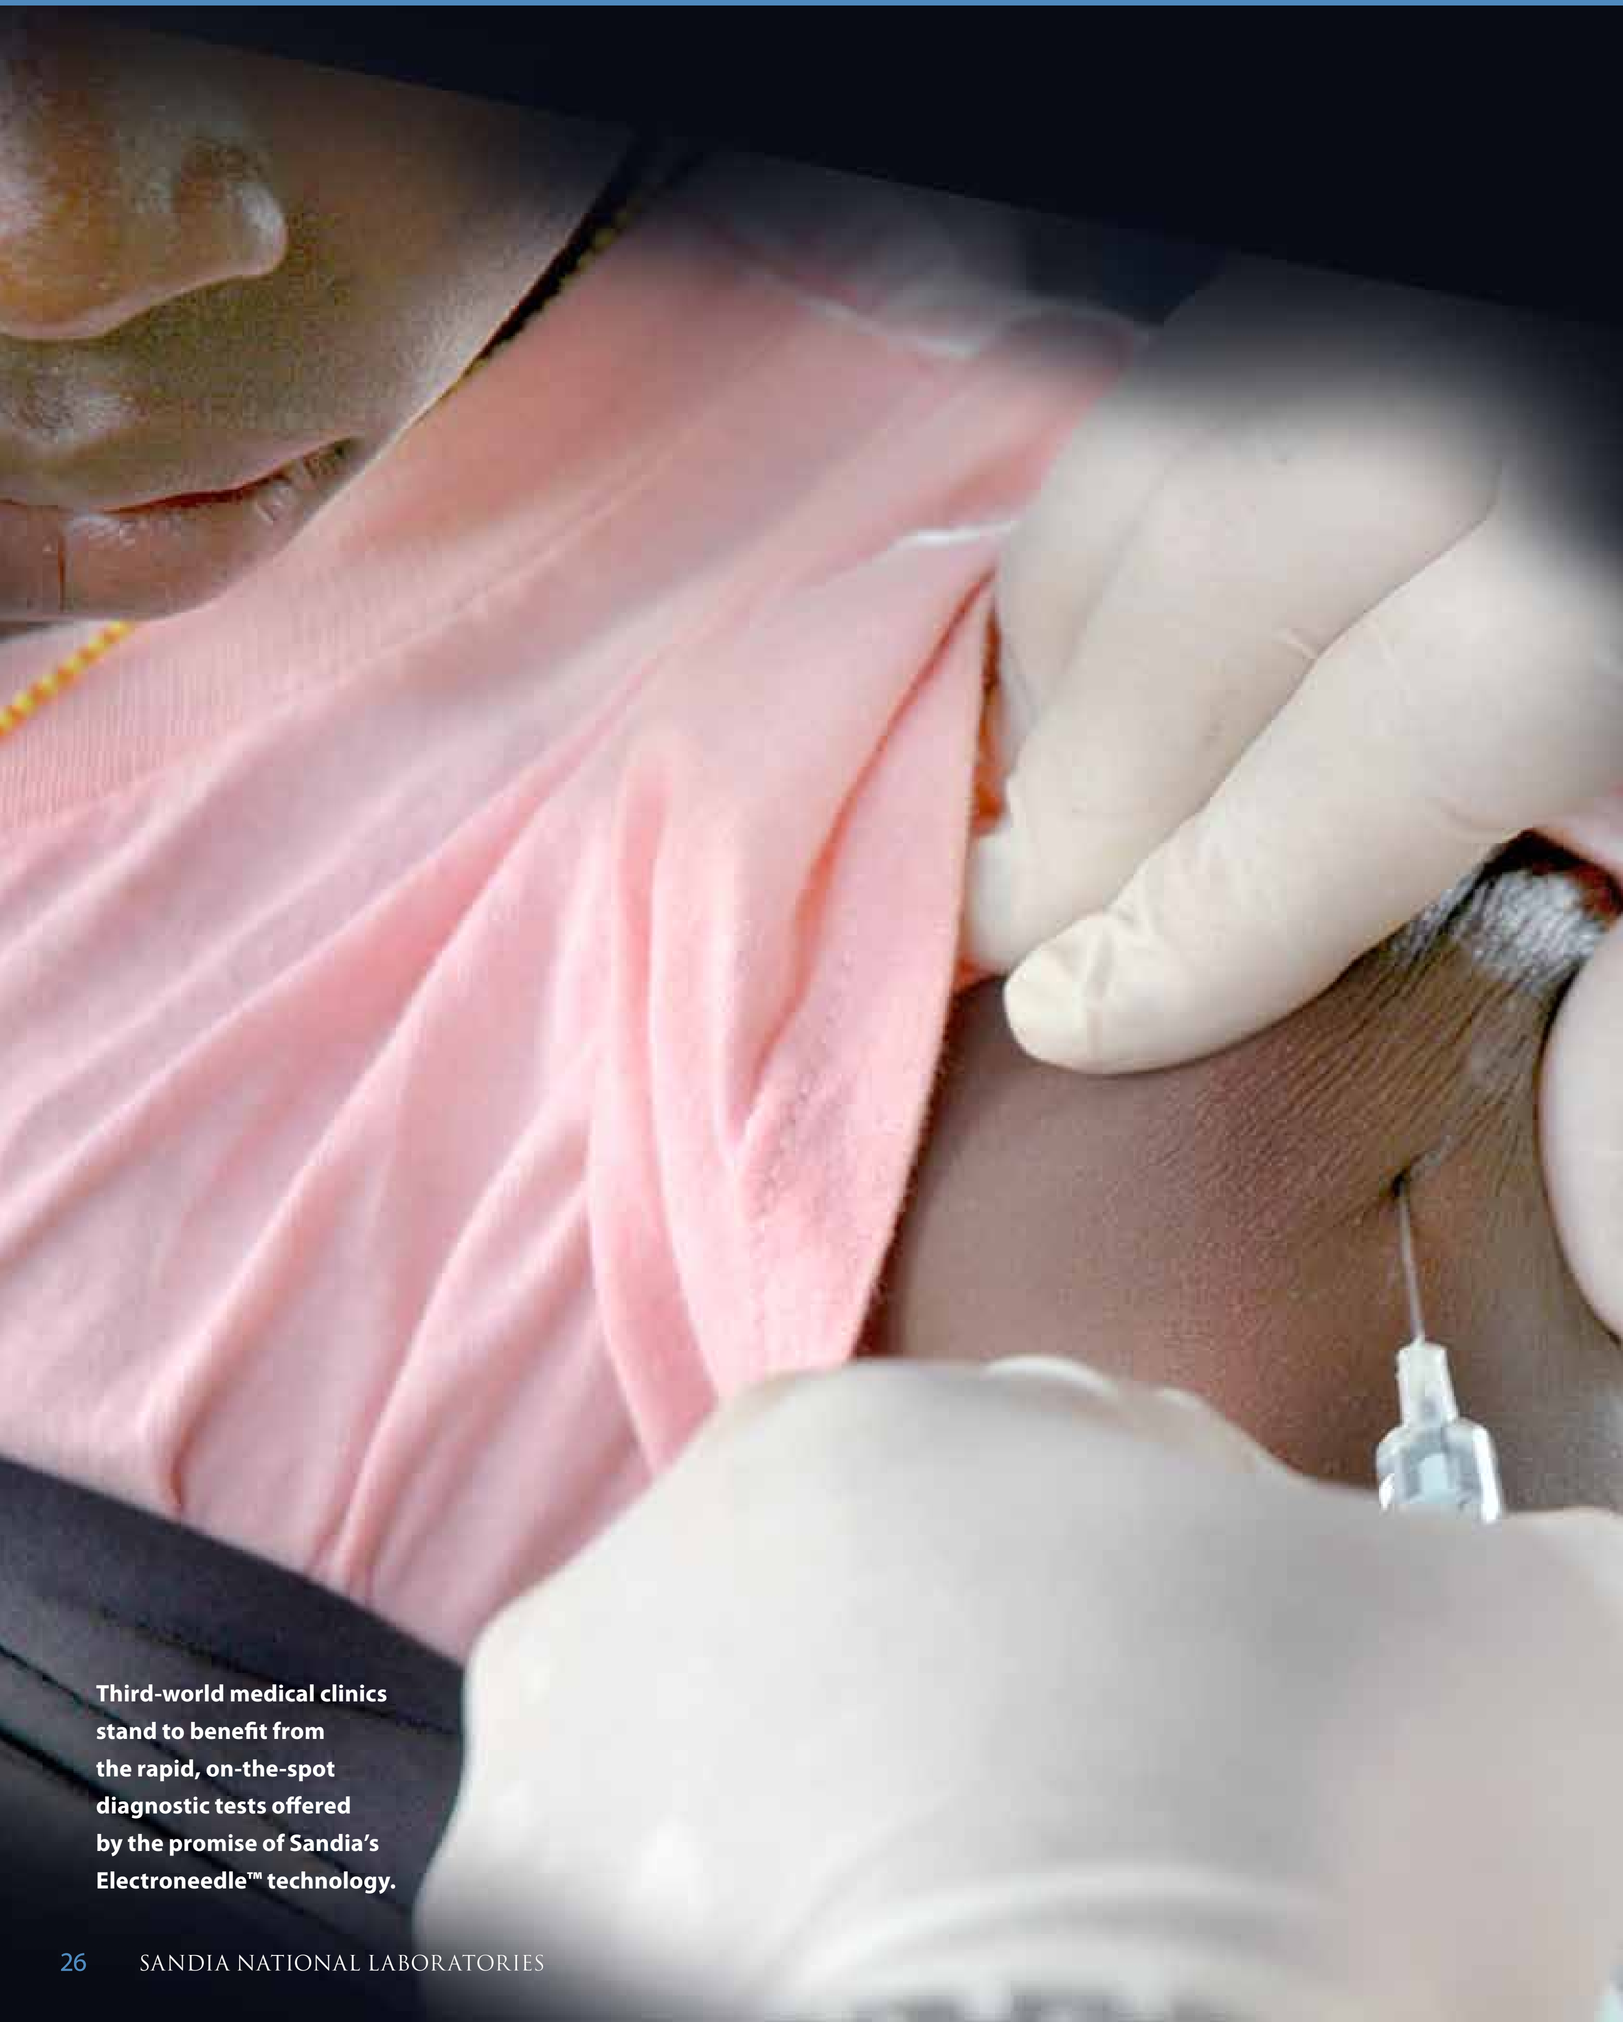

Third-world medical clinics stand to benefit from the rapid, on-the-spot diagnostic tests offered by the promise of Sandia's Electroneedle™ technology.

Anywhere Instant Diagnosis

ElectroNeedle™

A first-responder at an accident scene can treat a patient-in-crisis much more effectively if s(he) can gain an immediate reading on a particular aspect of blood chemistry; a physician treks to a clinic in the deepest recesses of the underdeveloped world — no hospitals, perhaps even no electricity. An environmental manager makes her way to a remote intermittent stream in the Southwest. The first-responder or physician, equipped with a small, portable instrument, carries the promise of on-the-spot diagnosis; the environmental manager, with a similar instrument in a backpack, is equipped with the possibility of an immediate accurate reading to detect a pollutant in a pristine ecosystem.

Building on a foundation of Sandia LDRD work in microsystems engineering, the technology known as ElectroNeedle™ offers the potential for point-of-care diagnosis of abnormalities that manifest as changes in blood and body fluid chemistry. For example, this capability could immediately alert a physician to changes such as the presence of a particular viral or bacterial infection, overcoming the delay inherent in current clinical laboratory test-and report procedures. Moreover, it could potentially bring such diagnostics into regions of the world where such laboratory services are either sparse or entirely unavailable. The equivalent laboratory-free, on-the-spot assessment of environmental changes applies to the environmental-monitoring scenario.

A minuscule square chip a few millimeters on edge, ElectroNeedle™ is designed to be placed on the skin's surface. The device interacts with the patient's body fluids by way of an array of patented microneedles, which depending on the application, are inserted shallowly into the superficial tissue fluid surrounding skin cells in the skin's outer

A square chip, about the size of a dime, Electroneedle™ is as inconspicuous as it is clever.

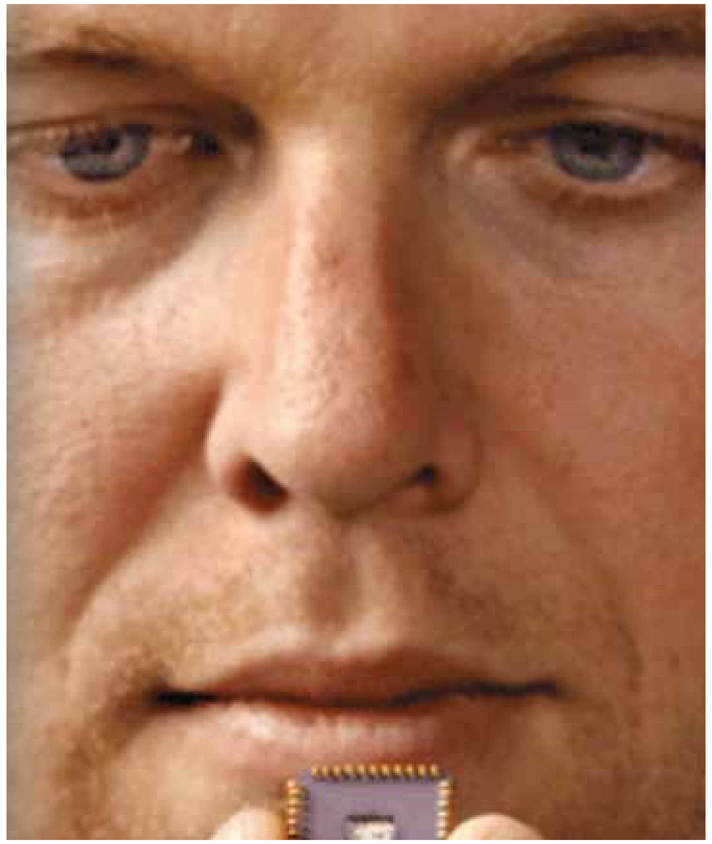

“. . . offers the potential for point-of-care diagnosis . . . overcoming the delay inherent in current clinical laboratory test-and-report procedures.”

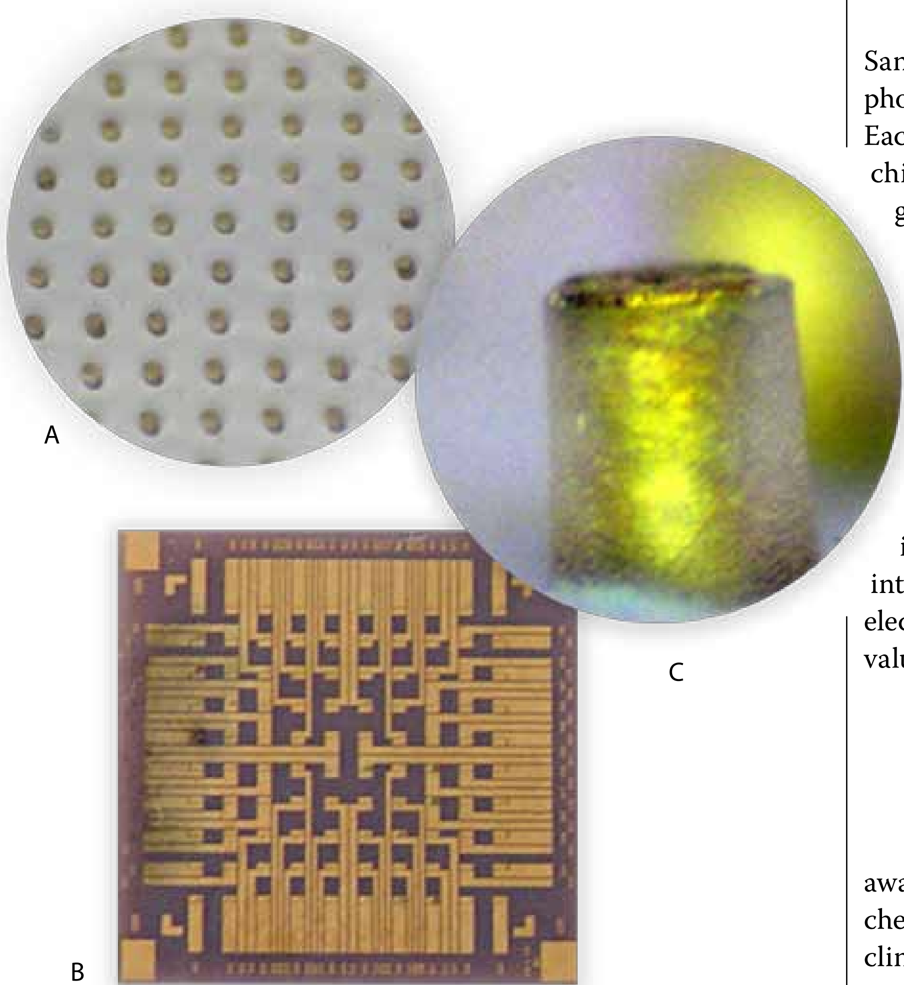

A) A 1 cm², 10 x 10 array of individually addressable electroneedles; B) The backside electrical contacts; and C) A single ElectroNeedle™ showing its protective sheath, leaving the metal electrode exposed only at the very tip.

layer (epidermis) or more deeply, into the microscopic capillaries in the deeper layer (dermis) of the skin. In the case of smaller molecules, such as glucose, their concentration (abundance) in the tissue fluid accurately reflects their blood concentration, since there is free movement (diffusion) of such substances between blood and tissue fluid; for certain larger molecules, for example, many proteins, their concentration must be measured directly from blood. Both measurements can be readily made from the skin, the human body's largest organ.

Underpinning the technology are Sandia microfabrication techniques, such as photolithography, etching and film deposition. Each of the microneedles in an array on the chip is composed of a commercially available glass known as Foturan® that is etched to create a hollow center subsequently filled with metal that serves as a sensing electrode, contacting tissue fluid at one end, within the skin, and making an electrical connection to sensor circuitry at the other end. This circuitry, in turn, connects to a small potentiostat that controls and measures changes in electrical currents and voltages and interfaces with a PDA, which converts these electrical measurements into body-chemistry values, such as blood-sugar concentration.

NO MIXING, NO MESS

The key to this technology and its R&D100 award-winning design is its reagentless chemistry. That is, unlike a traditional clinical laboratory, which adds measured quantities of chemical reagents to a patient's blood or urine sample to test for a specific body biochemical — a nutrient (e.g., glucose or cholesterol) or hormone (e.g., thyroid hormone) — ElectroNeedle™ relies on chemical modification of the metal tips of its needles. This so-called “functionalization” entails

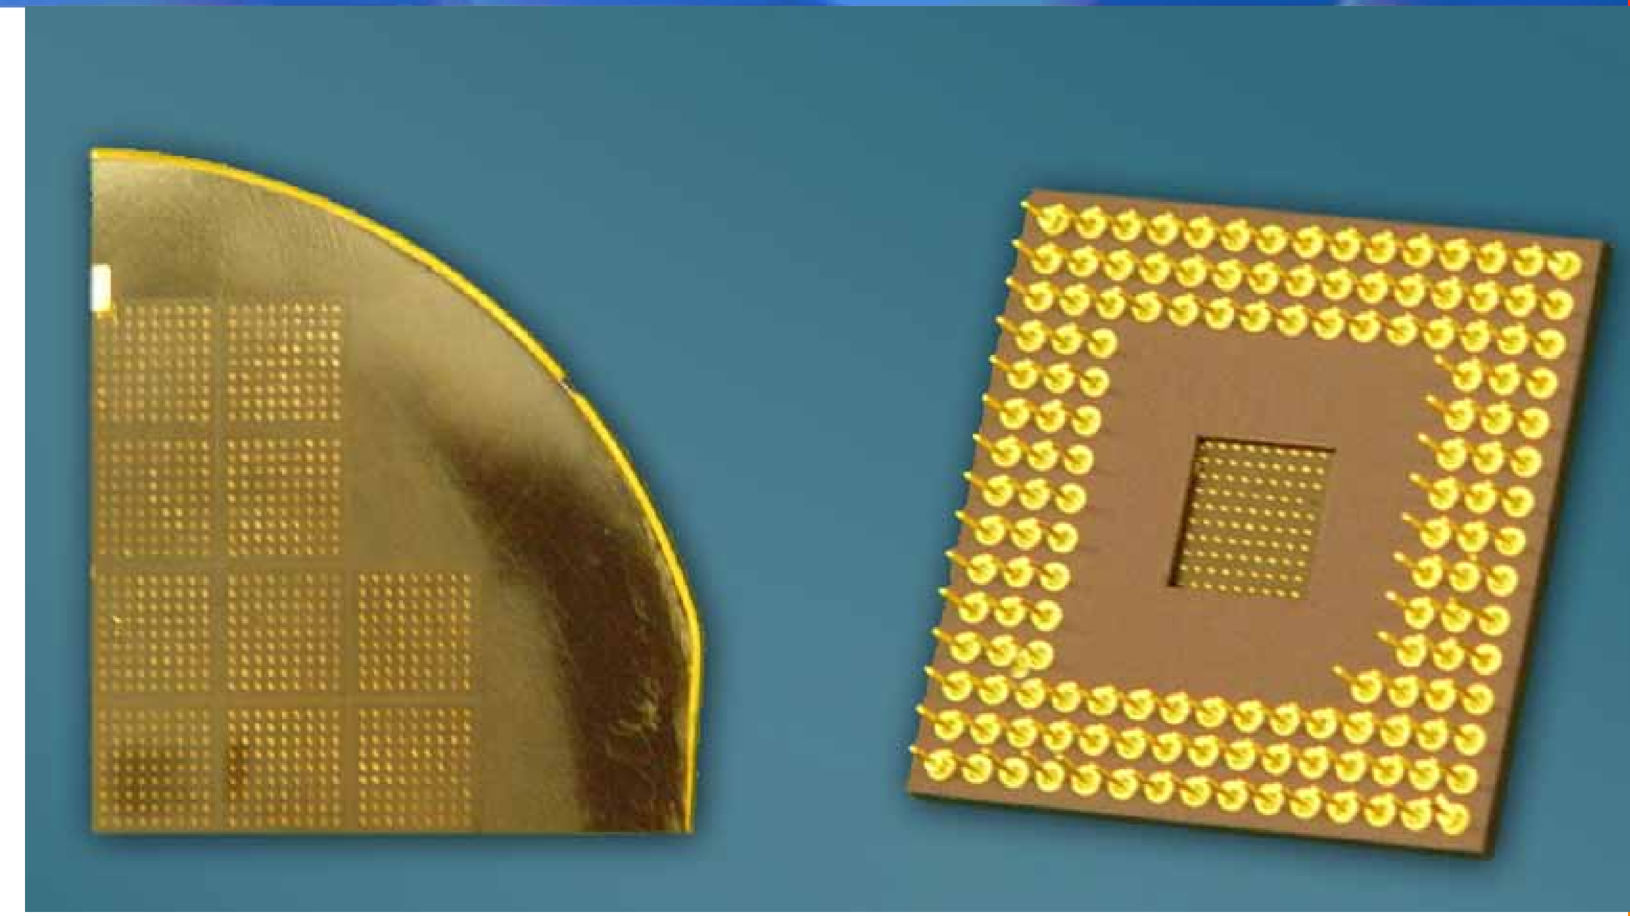

Microfabricated ElectroNeedle™ array. Left: Ten sensor arrays; Right: single array ready for use.

coating the tips of each needle with a particular reagent, one specific to detection and measurement of the concentration of a given body biochemical. Ideally, within an array of perhaps a hundred needles, one needle would be pre-functionalized to measure glucose, another to measure thyroid hormone, another for cholesterol, and so on for each needle in an array on the chip. Hence, under the best of circumstances, an array of 100 needles could assay for 100 different metabolites, in reagentless fashion.

Currently, some of these assays rely on the specificity of antibodies for identifying substances, that is, they are immunoassays, and as is frequently the case with such assays, they identify metabolites using fluorescence techniques. Ultimately, however, the goal is to adopt or develop methodologies that are strictly electrochemical in nature, that is, in which the reaction of a physiologically relevant substance (e.g., thyroid hormone) or infecting parasite (e.g., Dengue virus) with the reagent coated on a needle will be detected as a change in electrical voltage or current by the potentiostat, and subsequently, converted by the software into a measurement of that target substance's concentrations in tissue fluid and/or blood. The accomplishment of this goal would mean that this completely reagentless technology could bring a battery-powered laboratory to the patient, regardless of where that patient happened to be located.

Application of the technology in the area of environmental monitoring is a logical offshoot of its reagentless configuration. For example, currently, an environmental researcher, manager, or field technician must either gather samples and return them to a testing laboratory or carry reagents into the field. With a reagentless technology, measurements can be made directly and with an immediacy that allows decision-making regarding other relevant measurements, which it might behoove a researcher to make immediately upon discovery of a change in environmental conditions; important data are less likely to be missed.

> *“. . . the value of LDRD’s support of S&T endeavors that test and pursue novel directions in the quest for enhanced national security.”*

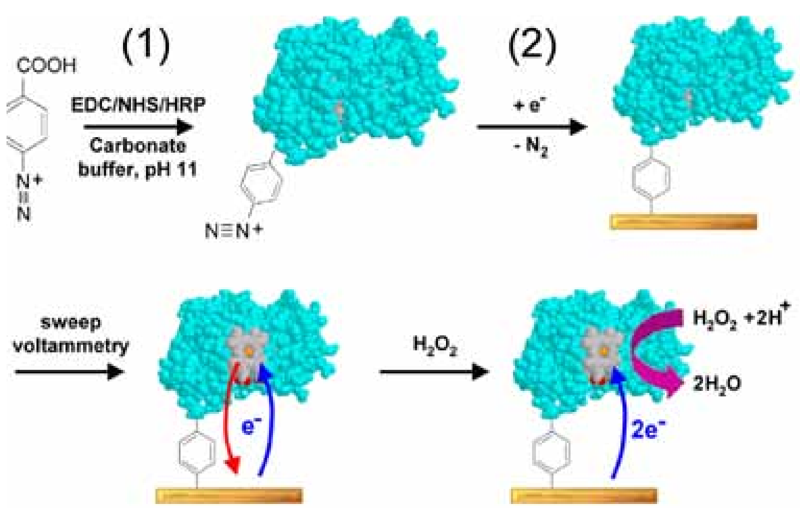

One methodology for functionalizing an ElectroNeedle™ tip by attaching an enzyme to it, and the enzyme’s subsequent reaction with its substrate, hydrogen peroxide (H_2O_2), resulting in an electrical current.

A NEEDLE FINDS MULTIPLE FUNCTIONS

The technology’s LDRD connection is an interesting combination of both precursor and follow-on LDRD projects, as well as non-LDRD projects. At its inception, the idea for the critical hardware component of ElectroNeedle™ was actually a spinoff of an LDRD project whose investigators were interested in the possibility of a fuel cell that would be powered by glucose. Although it is best known as blood sugar, glucose is a universal biological energy-currency carbohydrate. Living cells are adapted to “burn” sugars to liberate energy that they use for myriad cellular processes, and glucose is the evolutionary sugar-of-choice. Hence, the idea of a glucose-powered fuel cell is reasonable and potentially groundbreaking. As part of the fuel cell development, this LDRD team designed arrays of hollow glass microneedles to extract glucose from plant tissue and transfer the glucose to the fuel cell. While working on this project, the team recognized that these microneedles provided a unique capability to interface with a variety of cells and tissues. Could the needles, for example, be used to extract human tissue fluids for analysis? But, more creatively, since extracting fluids for downstream analysis requires complicated fluid manipulations, could such needles somehow be adapted to a methodology for in situ detection, that is, detection on the spot, without any extraction at all?

Discussions among the bio fuel cell LDRD project staff engendered the notion of filling the hollow glass needles with an electrically conducting metal, which they dubbed ElectroNeedle™, thereby opening up the possibility of performing electrochemical assays, that is, assays that yielded to the possibility of measuring a substance by way of a change in the electrical potential (voltage) or current. Typically, the team knew, this could be accomplished if a catalyst could be found that acted on the substance to be measured, promoting its transformation to a related substance with an altered electrical potential.

In the case of glucose, a simple technique was already well-established. The enzyme (biological protein catalyst) glucose oxidase (GOx) catalyzed (sped up) an oxidation–reduction reaction that transferred electrons from the glucose. This transfer of electrons, an electrical current, could clearly be detected and measured electrochemically; and as more and more reactions occurred, more current would flow, thereby detecting not only the presence of glucose, but also its concentration, a critical measurement for the control and treatment of diabetes. In the lab, using samples of pig skin — procured at a local butcher shop — as a model for human skin, the team was able to demonstrate glucose detection using the microneedle electrodes.

What this particular sequence of events illustrates is the genesis of scientific and technical possibility that frequently occurs when scientists are placed in creative groupings, with the flexibility to pursue paths that are even just a tad removed from the constraints of a short-term definitive project outcome, yet where an overarching mission of contribution to national health, safety, and security is a continually provocative driver. With Sandia’s national security mission as a constructive constraint and its inherently interdisciplinary team organization, research in one domain can often seed a new direction — a backdrop against which LDRD funding can sometimes encourage creative solutions of types that might not otherwise have been considered.

The team realized that there were two significant problems that had to be addressed. One was a method to functionalize each individual microneedle electrode so that an array of differently functionalized needles could perform measurements of multiple different aspects of body chemistry. Because of the small size of each needle and the large number of needles in each array — up to one hundred in the prototype — a manual approach was out of the question. The second problem was a shortage of electrochemically active assays, analogous to glucose oxidase, that could be bound to the ElectroNeedles™ to detect other medically (or environmentally) relevant substances. The solution to the first problem was found in the results of an earlier LDRD project that investigated the surface chemistry of the diazonium molecule. One end of this versatile molecule could be attached to a variety of other molecules — organic, inorganic, or biochemical — while the other end could be bonded to a metal surface by applying a voltage to that surface (without the applied voltage, no bonding occurs). This technique provided an automated way to functionalize the ElectroNeedles™.

PDA-controlled potentiostat for use with ElectroNeedle™ arrays.

To address the second problem — a need for more electrochemical assays — the team turned to nanomaterials research done at Sandia for a national security sponsor (defense). Metal nanoparticles such as platinum are known to exhibit catalytic properties. Sandia researchers showed that these nanoparticles could be attached to metal electrodes using the same diazonium chemistry employed to functionalize the ElectroNeedles™. To this needle-attached nanoparticle, the next step was to attach a biological molecule that — by virtue of its shape and charge — would specifically recognize a

> “. . . under the best of circumstances, an array of 100 needles could assay for 100 different metabolites, in reagentless fashion.”

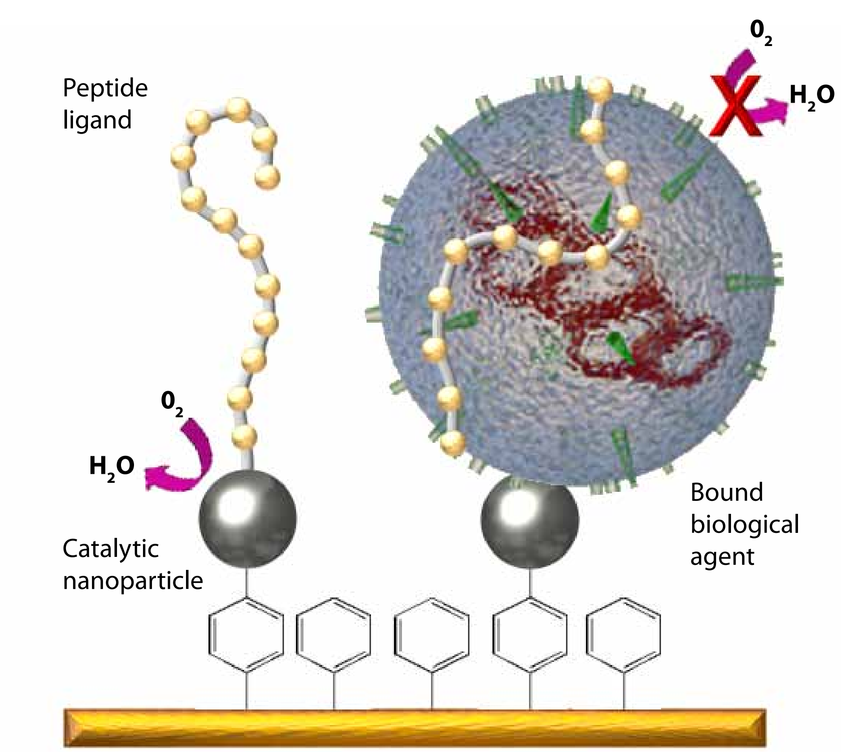

Drawing of the combination of biomolecule (a peptide in this instance) and catalytic nanoparticle (left); the former enables specific molecular recognition of a virus (right), which in this case, shuts off the nanoparticle's catalysis of the reduction of molecular oxygen to water (a change in electrical current).

biochemical substance of interest (such as thyroid hormone or the surface molecules of an infecting virus.) Antibodies and other proteins and single-stranded DNA exemplify such specificity-exhibiting biomolecules. This “molecular package” of needle-bound nanoparticle and specific biomolecule provides catalytic activity and molecular (shape-charge) recognition on the ElectroNeedle™, much like a naturally-occurring enzymes such as GOx. This development is the subject of ongoing research. According to ElectroNeedle™ Project Manager, Steve Casalnuovo, reagentless assay development continues to be the focus of electrochemical biodetection efforts because of its impact on Sandia's national security programs and the potential for improved healthcare diagnostics.

COLLABORATIONS

The evaluation and refinement of the ElectroNeedle™ technology's performance has been supported by a recent LDRD-funded effort to test the ElectroNeedle™ sensor array against several metrics defined by the FDA (sensitivity, specificity, linearity, limits of detection and quantitation, repeatability, ruggedness, and recovery), and in parallel, to evaluate small-scale manufacturing strategies in Sandia's microfabrication (MESA) facility. Using p-cresol, an environmental toxin with known kidney toxicity, as the target substance for detection, the Sandia team is collaborating with kidney-disease researcher, Dr. Dominic Raj at the University of New Mexico School of Medicine, to validate ElectroNeedle™ performance in situations that more-closely reflect those that would ultimately be encountered with patients.

In affirmation of the national security value of this technology, Lockheed Martin is likewise providing funding directed at the development of the technology, specifically for environmental monitoring. Hence, development on the medical and environmental fronts is proceeding simultaneously. And beyond the corporate-, LDRD-, and defense-funded studies, the ElectroNeedle™ story also includes forays into the business environment of New Mexico. Two biotechnology startup companies, New Mexico Biotech, Inc., and Life Bioscience, Inc., have been formed to commercialize ElectroNeedle™ technology, and as the various R&D pathways move the technology forward, both companies are exploring market for a technology with enormous global potential.

From a clever idea that advanced both scientific understanding and technology, through reconceptualization, to an R&D100 award, to external sponsorship supplementing and validating LDRD funding, the ElectroNeedle™ story illustrates the value of LDRD's support of S&T endeavors that test and pursue novel directions in the quest for enhanced national security. It is a story that validates the value of basic-science research as an endeavor in which, at times, the ultimate applications of a body of research cannot be anticipated, but which, when guided by the overarching national-security imperative, will often yield some very promising outcomes. ♦

ElectroNeedle™ arrays are fabricated in Sandia's MESA Microsystems facility.

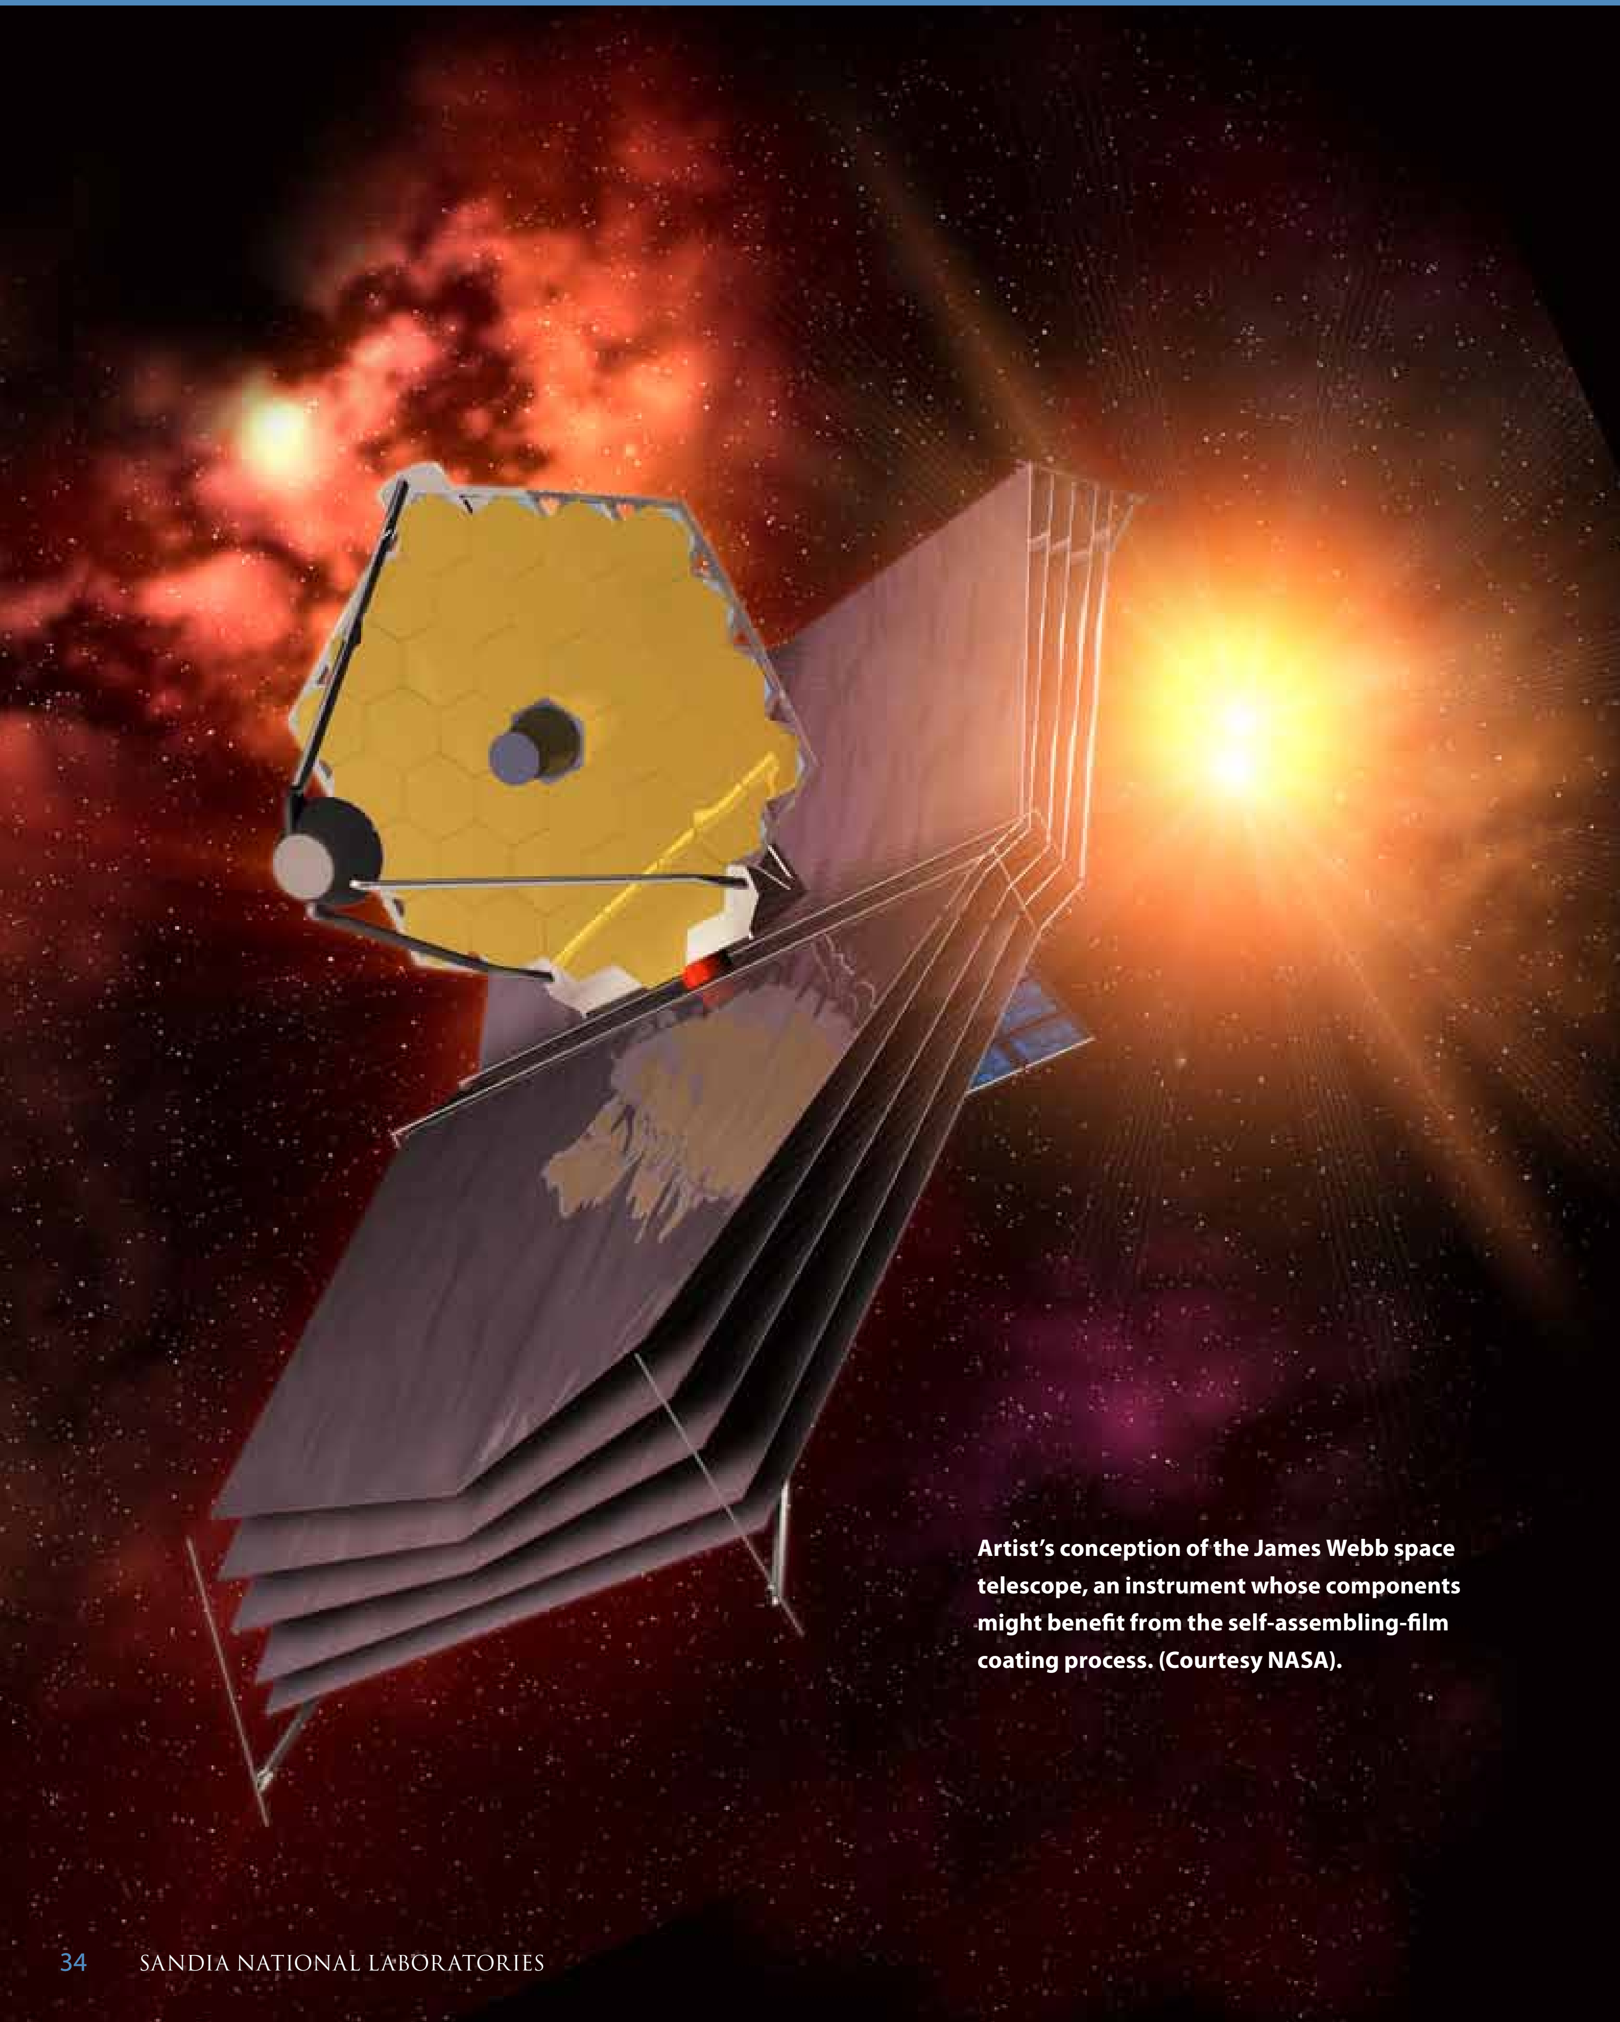

Artist's conception of the James Webb space telescope, an instrument whose components might benefit from the self-assembling-film coating process. (Courtesy NASA).

Gentle Smart-coatings

Self-Assembling Thin Films

From consumer electronics and photographic lenses to sensors on aircraft and lightweight plastic goggles for troops in the field: optical and electronic thin-film coatings have become an everyday fact of civilian and military life. Since it is quite likely that the applications for such coatings will continue to expand, it is also probable that both optics and electronics industries will greatly benefit from a revolutionary Sandia process for depositing thin-film coatings on surfaces.

Until this development, thin-film coatings have most often been created by techniques such as metal organic chemical vapor deposition (MOCVD). Although effective, these techniques are expensive, requiring elaborate and costly technology, and they are also demanding, requiring conditions such as high temperature and high vacuum. Not only do such conditions require elaborate equipment, but they also have the potential to damage certain materials — such as the aforementioned plastic goggles, for example. By contrast, Sandia's R&D100 award-winning coating process utilizes clever solution chemistry to deposit a thin film with desired nanodimensions (arrangement and spacing of individual molecular components), and without either a risk of environmental toxicity or the stringent equipment requirements of currently employed processes. Moreover, the chemistry involved has an amazing breadth of potential applications beyond what might normally be thought of in connection with such coatings — for example, in ultrahigh sensitivity chemical and biological sensors.

Multiple Paths from Fundamental Chemistry

The film-deposition process hearkens back to underpinnings of clever solution chemistry that LDRD played an important, if not a crucial role in funding. This larger story offers more evidence for the notion that LDRD provides the fertile soil in which ingenuity can test its DNA, find out if its genes are adequate to the task of carrying forward more-completely developed science that can then find its way into technologies using funding sources better adapted to technological applications. In this case, the "better mousetrap" scenario also applies, with Sandia chemists convinced that they could improve on existing work and develop a process that would lend itself to the resolution of fundamental dilemmas facing materials science in the area of semiconductor growth. That this research path also led to numerous side trails with quite interesting and fruitful destinations was both serendipitous and somewhat anticipated.

Several lines of LDRD-funded research, over the years 2003–2006, developed the chemistry eventually leading to the 2007 R&D100 award, whose specific proposal was merely one of several possible applications of this ingenious chemistry. The first of these developmental lines of research derived from an LDRD-funded project seeking to improve the methods by which Sandia grows gallium nitride (GaN) and other semiconductor materials crucial to many aspects of electronics (e.g., microprocessor chips) and optoelectronics (e.g., solid state lighting). Currently, semiconductor crystal growth takes place in a several-step process upon a sapphire or silicon carbide substrate, and because of mismatch

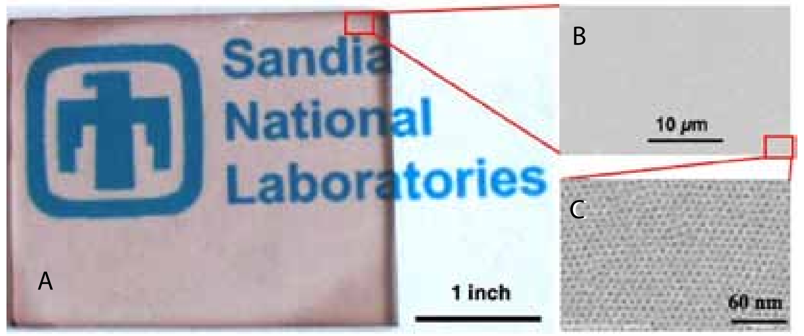

Views of a gold nanoparticle thin film at three different magnifications, showing its transparency (a), and at highest magnification, its ordered crystalline structure (c).

> *“. . . a chemist can tailor films with a range of optical, electronic, and structural characteristics.”*

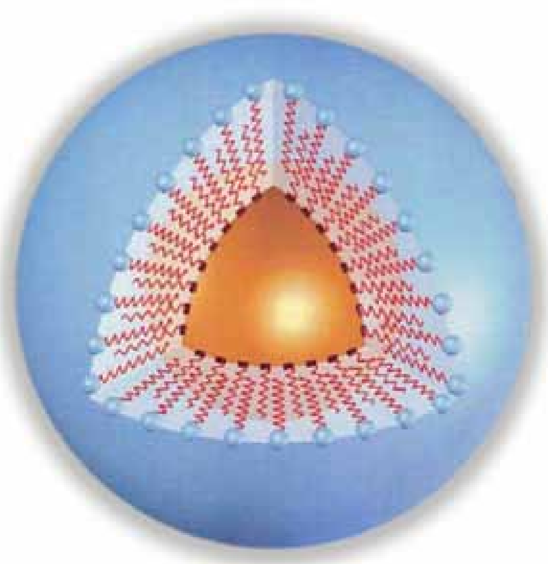

Conceptual drawing of gold nanoparticle structure, the gold atom, at center, surrounded by a sphere of amphiphilic molecules, their hydrophilic head groups (blue) facing outward toward surrounding aqueous solvent, their hydrophobic tails (red “zig-zags”) facing inward, toward the central gold atom. (Courtesy American Chemical Society).

issues in the crystal structures of the sapphire and the growing semiconductor crystals (e.g., GaN), a significant number of defects arise in the semiconductor. While other methodologies have been developed to address this problem at the microscale (see, for example, page 42 of this publication), this LDRD research proposed to address it at a more-precise nanoscale level of resolution by developing films with pores of nanoscale dimensions that would ultimately allow the growth of reduced-defect semiconductors. Honored by a 2007 LDRD Award for Excellence, this research not only demonstrated that the method had the potential to produce reduced-defect semiconductor growth, but it also gave the investigators a “handle” on the chemistry involved in the growth of self-assembled nanoscale films of different thicknesses and properties. In the end, the “chemistry lesson” turned out to be as valuable as the advance toward improved semiconductors.

Nanocrystal Micelles

The second, and parallel research initiative centers around a fundamental unit known as a nanocrystal-micelle. Essentially, prior work had shown the value of encapsulating both precious-metal (i.e., gold, silver, and platinum) and semiconductor (e.g., cadmium selenide) nanocrystals in fat-soluble, “water-shunning” (hydrophobic) organic molecules, to render them unreactive (passivate them), in order to retain their unique optical qualities. But in order to make these optical qualities useful in studying living cells, for example — that is, in order to render nanocrystals useful as bio-labels —the nanocrystal would have to be rendered water-soluble, since almost all cellular processes occur in aqueous (water-based) solution inside and outside the cell’s many nanoscale compartments. So Hongyou Fan, Jeff Brinker, Bruce Burckel, and their co-workers accomplished a process for forming nanocrystal micelles, the result essentially

Spray-coating thin films onto a diversity of surfaces is a relatively easy process.

encapsulating the metal or semiconductor nanocrystal into an “amphiphilic” particle with a water-shunning (hydrophobic) interior and a water-seeking (hydrophilic) exterior, thereby effectively dissolving these particles — “nanocrystal micelles” — in aqueous solutions such as the interior of living cells. As a crude analogy, imagine a tiny hollow sphere, its exterior wall composed of table salt, but its hollow center filled with cooking oil that is confined within the sphere. Immersed in water, a small enough sphere might technically be dissolved, in that water molecules would chemically interact with the external salt, as they do in a salt-water solution, (forming what are known as ion-dipole bonds). But the oil in the interior of the sphere would not really be dissolved. Such is the paradoxical nature of micelles in water, with the exterior of the micelle “dissolving,” while the interior remains undissolved. (Problematically, in the analogy, as the salt coating the sphere dissolved, the surface would disintegrate, and so the analogy quickly, and quite literally, breaks down.)

Most critical to their micellar encapsulation was the fact that the nanocrystals retained their optical properties in this chemical configuration, including the ability to display multiple colors simultaneously. Additionally, the nanocrystal micelles remained dispersed in aqueous solution, that is they did not aggregate (clump together). These were important qualities in ensuring their ultimate use as biosensors. By employing phospholipids, biology’s own amphiphilic (part water-shunning, part water-interacting) cell membrane molecules as the structural capsule of the micelle, and chemically modifying the outer, water-soluble face, these structures have the potential for use as a diverse array of bio-sensors. As an example, by chemically attaching an antibody molecule to the outside of the micelle, with the antibody specific to the identification of a particular virus, the presence of that virus in biological fluids might be identified by an optical signal coming from the nanocrystal within the micelle.

With the nanocrystal micelles as fundamental molecular entities, Sandia’s patent application for the thin-film coating process is predicated on its unique self-assembly characteristics that allow control of the resulting films at a nanometer scale. This astounding ability to control chemistry

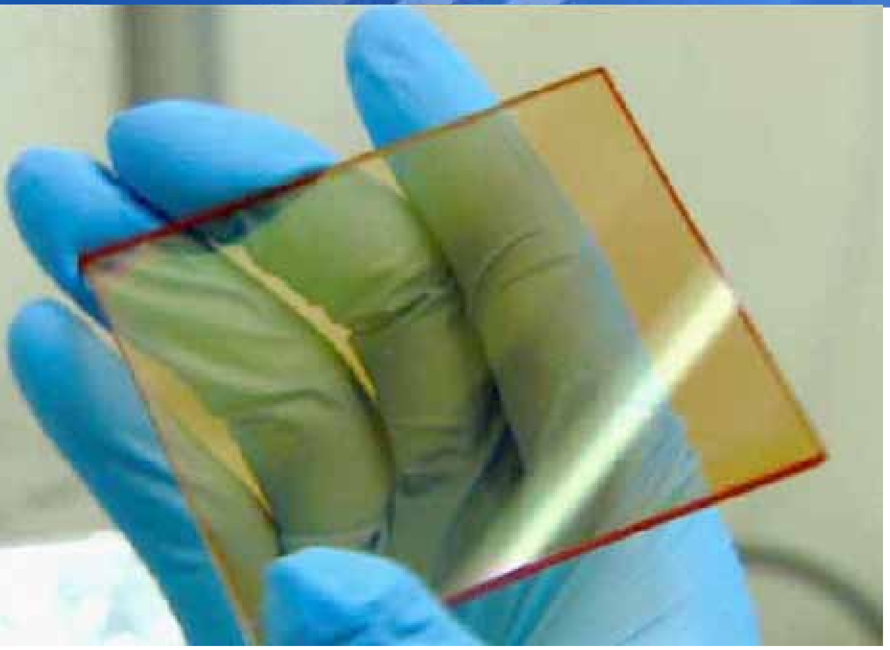

A gold nanoparticle film is perfectly transparent after drying onto the surface of this glass plate.

The LANTIRN sensor assembly on the underside of this F-16 aircraft illustrates a military application for the smart-coating process.

at such a fine level of resolution (crystal structure, spacing of the nanocrystals in the film, etc.) received world-wide attention and acclaim, the chemistry and biological science communities immediately aware of the implications. Several invited cover stories in major technical publications (Chemistry of Materials, Advanced Functional Materials and others) have publicized the discovery, propelling it into international prominence.

Such fine control of crystal structure simply cannot be obtained by current chemical or physical deposition processes. Moreover, because the collective physical properties of such crystals derive from coupled interactions among the nanoparticles, the implications for this control of crystal structure open the possibilities for new, potentially as-yet unexplored physical properties. Not surprisingly, the biosensor work is proceeding, funded both by LDRD and by a Sandia/Lockheed Martin Shared Vision Grant. But, according to Fan, the LDRD "investment in fundamental science" paid off in another way as well, in the technology that garnered the R&D100 award, which drew both from the experience of creating self-assembled templated films for improved semiconductor growth, as well as from the research employing nanocrystal micelles as chemical building blocks and potential sensors in aqueous systems. The two lines of LDRD research coalesced during the development of this coating process.

In the semiconductor industry, thin film coatings are usually generated by metal organic chemical vapor deposition (MOCVD), a process occurring at high temperature and under high pressure, conditions that limit the application of semiconductor films to materials that are resistant to deformation or damage under such extreme conditions. By contrast, the R&D100 process describes a methodology for coating surfaces with either semiconductor or optical films under much more gentle ambient conditions by means of dip, spray or spin coating. Essentially, what this requires is further chemical modification of the water-soluble nanocrystal micelles to incorporate chemical substances into their structure that desirably control their self assembly into highly ordered two and three-dimensional sheets (films) of different thicknesses.

Best illustrating such an incorporated substance is silica (silicon dioxide), the primary components of most glasses and quartz. Added to gold nanocrystal micelles in the form, tetraethyl orthosilicate (TEOS), silicic acid binds to the hydrophilic exterior surface of the micelles, at the water interface. A solution of these micelles readily coats surfaces by means of dipping or spin coating. Subsequent evaporation of the solvent water leaves an orderly array of gold nanocrystals embedded within a matrix of silica — a thin film of quite orderly structure and known optical properties. In this case, the film possesses the physical properties of silicate glasses (such as strength). But to create, for example, a more-flexible, plastic-like film, chemists can employ other polymeric substances with the desired characteristics, in substitution for the silica. In all cases, the nanocrystal micelle is the essential building block, and by varying the nanoparticle (gold or silver for optical films; semiconductors such as CdSe for electronics films), the structural component (silica, titania, various organic polymers), and controlling the concentration in solution and the coating speed and evaporation conditions, a chemist can tailor films with a range of optical, electronic, and structural characteristics.

The ingenuity of this material process is evident from the fact that films form at ambient conditions, without a requirement for high temperature or high vacuum. This allows the coating of, for example, softer plastics whose molecular structure and shape might deform under the extreme conditions required for MOCVD. It also allows coating of hidden surfaces in irregularly-shaped parts that might be inaccessible to other methods of film deposition. Furthermore, the coating process is not limited by the size of parts to be coated, as is the case with vapor deposition processes that must occur within a vacuum chamber.

Gentle Conditions

With control of film structure and thickness comes the possibility of producing high-performance optical coatings, such as those for lenses and mirrors, and critical sensor widows on high-performance military aircraft, the reflective properties of the coatings with respect to specific wavelengths of light precisely controllable by depositing multiple film layers with different optical properties. Should this technology be fully realized, it may ultimately be possible to replace damage to these coatings (suffered during routine operations or through rain erosion) by simply applying a new "touch-up layer of coating, similarly to an automotive touch-up paint.

> *". . . the readiness and willingness of the LDRD program to invest in basic chemical science have . . . opened up a universe of possible applications."*

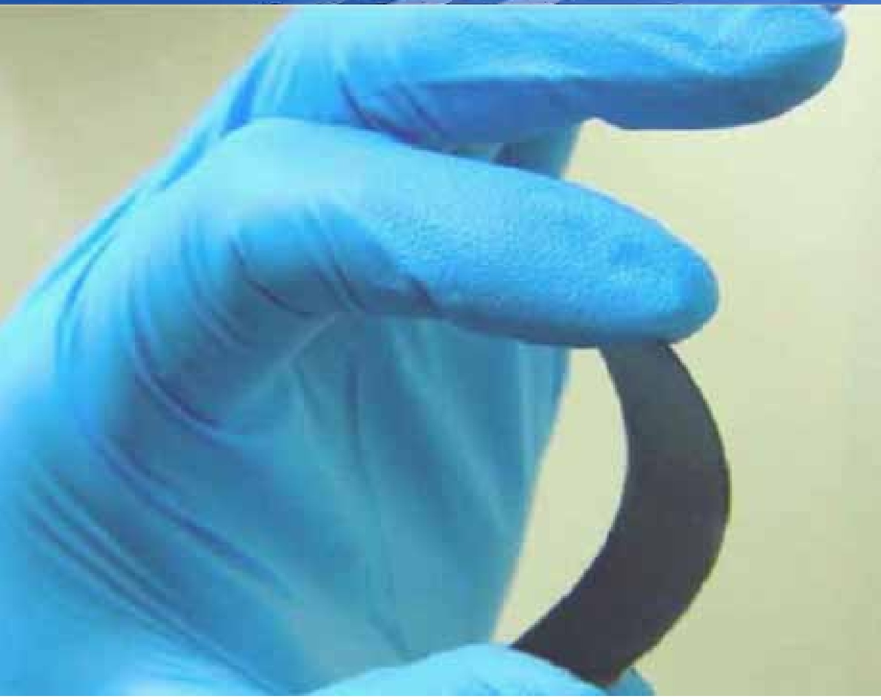

Incorporating polymers rather than silicates as the structural backbone of a thin film confers the property of flexibility.

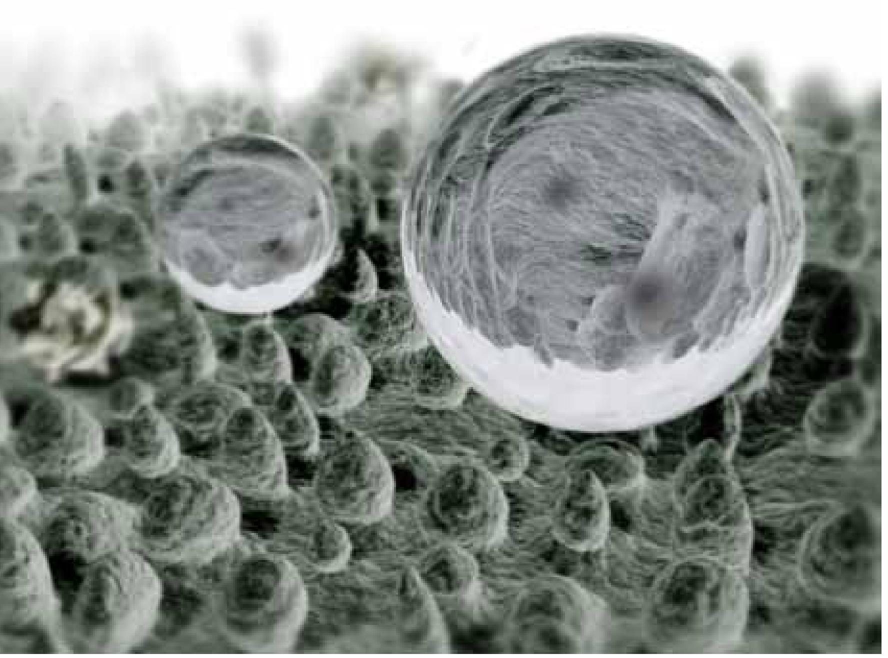

Thin films with superhydrophobic properties have recently been made at Sandia, which reject water similarly to the surface of this lotus leaf.

This is a far cry from current operations, where entire windows must be removed, recoated and replaced. Additionally, this method has the further advantages of environmental friendliness, by comparison to, for example, MOCVD, whose gas precursors can be highly toxic. Support for ongoing development of optical coatings derives from Lockheed Martin, with whom licensing negotiations are in progress. Other applications of the coating technology include irregularly shaped parts in MEMS (microelectromechanical systems), in coatings for high-density flash memory devices, and in interconnects for minuscule wiring networks for the electronics industry.
The breadth of possibilities for the basic structural unit, the nanocrystal micelle, is testified to by ongoing work, funded both by LDRD, and by Sandia Basic Energy sciences (BES), with other external support pending from DARPA. First, work to use templated films for reduced-defect semiconductor crystal growth is ongoing. Second, research into both chemical detectors and bioagent detectors likewise continues, for example by employing antibodies and other proteins attached to the outside (aqueous) surface of the micelle to function as specific molecular binding agents for identification of specific biotoxins or organisms. By using surface enhanced Raman scattering–based detection methods, scientists may be able to leverage such Raman fingerprint identification of molecular components to field detectors with potentially single-molecule sensitivity. That is, unlike traditional detectors that usually require a threshold concentration in the environment before reliable detection, there is no threshold for nanocrystal micelle detectors; just a single molecule would be reliably detected. The likelihood that threat agents would escape detection is thereby greatly diminished.

And in what, for biologists, is an even more-fascinating application, nanocrystal micelles are now being developed for in vivo use in neuroscience in another LDRD-funded project. This project seeks to bridge the type of recordings that are available from the brains of living animals in order to better understand the relationship between specific aspects of brain activity and behavior. For example, electroencephalography (EEG) and functional MRI (fMRI) allow recordings from large areas of the brain during activity, and at the other extreme, microelectrode recordings can inform biologists about the electrical activity of single nerve cells. But tools are lacking for the intermediate structural level, that is, recording the activity of networks of interconnected brain cells in a living, behaving animal. This project proposes to bridge that gap through the use of biocompatible nanocrystals that would act as detectors of various forms of neurological activity at the cellular and subcellular levels, essentially allowing visualization of the activity of populations of brain cells as an animal behaves in its environment. Accomplishing this goal could represent a breakthrough in understanding the connection between brain processes and behavior, which, with few exceptions, is poorly appreciated.

Repeat Performance

Not so remarkably, another branch of this LDRD-seeded Sandia chemistry staged an immediate repeat performance in the person of Brinker's Superhydrophobic Coating, a 2008 R&D100 winner. Again underpinned by a silicon dioxide–based nanostructure that completely self-assembles after solvent evaporation, these films repel water so strongly that the beading of water droplets occurs at remarkably large angles to the film's surface, an indicator of the strength of "repulsion" of the surface for water and therefore of the tendency of the water molecules within the droplet to "want" to cling together away from the coating. Inspired by biological structures in plants and insects that attain this outcome through physical patterning and the ingenious biochemistry of Nature, the Sandia films likewise attain their properties through nanoscale patterning and chemistry, and can be applied in similar fashion as the nanocrystal films — by spin-coating, spraying, or dipping. The films are transparent, and therefore, have an obvious application in the protection of objects of antiquity and other precious art that can be (and have been) damaged by such phenomena as acid-rain corrosion. Facilitating flow of aqueous fluids in narrow microfluidics channels and within pipe networks, decreasing the pump-energy required to propel such fluids is another important energy-conservation application.

In addition, UV light can be used to pattern hydrophilic zones into the overall superhydrophobic film, in situations in which it is desirable to allow water condensation, and subsequent runoff, such as its application in water-harvesting in arid environments. Overall, the winning technology is yet another exhibition of the myriad possible offshoots that are sprouting from the solid foundation in chemical nanoscience established through the auspices of Sandia's LDRD program.

As Fan himself emphasizes, the readiness and willingness of the LDRD program to invest in basic chemical science have — because of the leverage conferred by the strength of that understanding — opened up a universe of possible applications. This underscores the value of LDRD research into fundamental principles, the outcome of the understanding thus gained inscribing a considerable listing of applications that will likely continue to grow. Moreover, in a field as young as nanomaterials chemistry, there will undoubtedly be new lines of basic research fueled by these discoveries, some of which will likely point to applications as yet unimplemented and possibly yet to be conceptualized. ♦

A near-ultraviolet light-emitting diode (LED). Reduced-defect semiconductor materials would improve the performance of these devices.

Reduced-defect Materials

Cantilever Epitaxy

High-quality semiconductor materials form the basis for numerous current and future advances in energy-efficient technologies that will likely be crucial to US and planetary solutions to both diminishing traditional energy sources and to human-induced climate impacts. For example, so-called III-V materials, such as gallium nitride (GaN) — chemical compounds deriving from combinations of elements from the III and V columns of the periodic table of the elements — have very desirable properties in applications such as solid-state lighting. The light-emitting diodes (LEDs) with which we have become so familiar on cell phones and digital music players form but a tiny slice of the more extensive marketplace of applications potentially available should some of the fundamental materials science issues in the growth of semiconductor materials like gallium nitride be resolved. Alloyed with other metals such as aluminum or indium, GaN is useful in the formation of ultraviolet (UV) and blue LEDs, and potentially for other colors that have been difficult to optimize. The key nature of this materials science problem was illustrated by the conferring of a 2004 R&D100 Award on a Sandia-derived process for the growth of higher-quality GaN and other semiconductor materials.

Dubbed "cantilever epitaxy," the process was designed to compensate for the fact that many materials must be grown on an underlayer, or "substrate" whose chemical crystal structure — arrangement of atoms in space — is somewhat spatially different than that of the material being grown atop it. This so-called "lattice constant mismatch" results in defects in the newly growing crystal material and is best illustrated by the growth of gallium nitride on a sapphire (a type of aluminum oxide [Al_2O_3]) substrate. The end result of this situation (known as "heteroepitaxial growth") is that any GaN crystallized in contact with the sapphire crystal suffers from irregularities in its structure that are termed "threading dislocations," because they can thread or propagate from one layer of growing crystal to another.

The density of these irregularities in the GaN crystal structure can be quite high. Theoretically, these defects should influence the characteristics of this semiconductor material in serving as a suitable candidate for use in solid state lighting (SSL). For example, two key parameters are a material's bandgap, which determines its possible spectrum of light emission under the appropriate circumstances, and the material's efficiency in the process by which electrical energy is transformed to light energy. Since a material with more defects is expected to perform less suitably in both respects than one with fewer defects, materials scientists have been searching for new, reduced-defect methods of growing gallium nitride and other semiconductors. Cantilever epitaxy (CE) was developed as just such a methodology.

Growing Crystals in Thin Air

The general idea is eminently sensible: since threading dislocations occur when gallium nitride crystals grow in contact with the sapphire underlayers, why not reduce the areas of contact, in order to diminish the frequency of

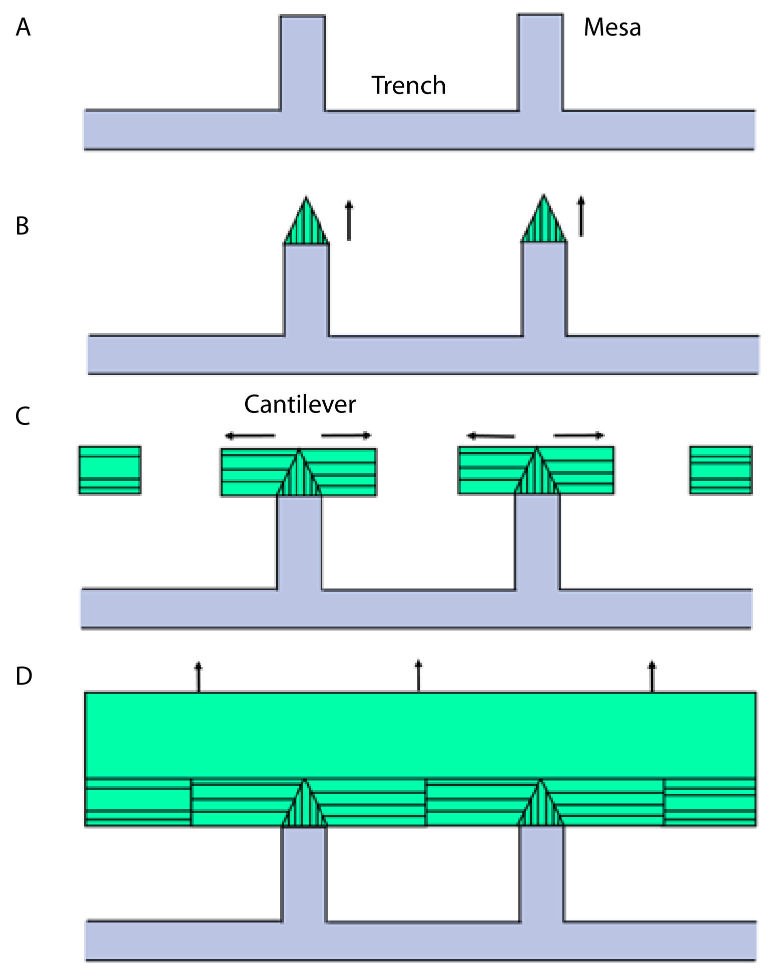

Schematic drawing of cantilever epitaxy, illustrating crystal growth initiating on mesas, then continuing as cantilevers over trenches, with defects (dislocations) turned horizontally, so that they do not "thread" into upper layers of the crystal.

> *". . . applicable to the growth of many other materials in situations where there is a mismatch in crystal lattice parameters."*

those defects, and also devise a way to prevent the contact-defects that must be tolerated from propagating into upper regions of the crystal? Several schemes had already been proposed to do just that including the most popular at the time, ELO (epitaxial lateral overgrowth). The use of ELO had improved several measures of performance in optoelectronic crystals grown by the method — such as those of GaN and indium gallium nitride (InGaN).

LDRD had been funding several Sandia teams studying GaN and related semiconductors (such as aluminum gallium nitride, AlGaN) for at least a decade, and the solid state lighting Grand Challenge from FY2001 through FY2003 had several different foci, among them a concentrated effort to more fully understand the fundamental materials science of semiconductors. In this atmosphere, a UNM student, Christine Mitchell, working with Sandia scientists, Carol Ashby and David Follstaedt, undertook a project to better appreciate the theoretical underpinnings of ELO, specifically as applied to the growth conditions that Sandia scientists, Andy Allerman and Jung Han (now a Yale faculty member) encountered in their work. The difficulty and hence research-limiting step of actual crystal growth were significant in slowing the rate at which such investigations could proceed.

It was during an ELO-related conversation between Mitchell and Ashby that the "better-way" notion about cantilever epitaxy was hatched. Essentially, ELO required two growth steps, necessitating removal of the growing crystal from the growth chamber, then its subsequent reintroduction for a second round of growth. Suppose that growth of GaN with diminished substrate contact could be more-simply performed in only one growth step, thereby rendering the process quicker, more efficient, and less expensive?

Beginning with a sapphire substrate into which trenches (grooves) have been etched, growth of gallium nitride is initiated atop the remaining posts, (dubbed "mesas") between the trenches. At the temperature used for crystal growth, the GaN grows vertically atop the mesas, forming pyramids that reflect the material's crystal symmetry. These pyramids allow growth to be turned horizontally, so that much of the GaN subsequently grows over open space, as cantilevers not in contact with the sapphire, and hence, of reduced-defect quality. The defects already present propagate (thread) horizontally, only within that initial growth layer. Hence, the uppermost growth layer of GaN contains a significantly reduced number of defects propagating from the initial, deepest region of growth. The cantilevers grow toward each other over the trench and eventually meet to form a continuous crystal layer. (See drawing, p. 44).

Broad Applicability

CE reduces the density of defects by about two orders of magnitude, from about 10^9 or 10^{10} to about 2×10^7 defects per square centimeter of GaN crystal, and more importantly, only a single growth step occurs, the changes in growth direction controlled by employing different temperatures within the growth reactor; this reduces both time consumption and cost by comparison to ELO. The process obtained a patent in 2003 and was promptly honored with the R&D100 award in 2004. The patent was awarded for the process itself, which far from being limited to GaN growth, is applicable to the growth of many other materials in situations where there is a mismatch in crystal lattice parameters between the growth of a material and the substrate upon which growth is occurring.

Although LDRD funded much of the preliminary work, in the characterization of ELO and some of the growth conditions for CE, work that led to a Master of Science thesis for Christine Mitchell, funding also derived from the National Energy Technology laboratory of DOE, as part of a joint project between Sandia and the Silicon Valley, CA lighting company, Lumileds (which later became part of Philips Inc.). Sandia provided significant quantities of reduced-defect materials to Lumileds as part of the company's developmental efforts in the SSL area. The Sandia developers point out that it was unlikely that Lumileds would have participated in funding this joint project without the proof-of-principle work that LDRD had initially funded. This pattern of proof-of-principle funding by LDRD, followed by external sponsorship for further development of a process or product is one that frequently transpires in the LDRD arena.

> *". . . the readiness and willingness of the LDRD program to invest in basic chemical science have . . . opened up a universe of possible applications."*

Ultimate Payoffs

Problematically, ELO was an entrenched process in the marketplace, and so CE did not make significant inroads as a commercialized process. Additionally, at the time of the process' development, silicon carbide

> ". . . it was unlikely that Lumileds would have participated in funding this joint project without the proof-of-principle work that LDRD had initially funded."

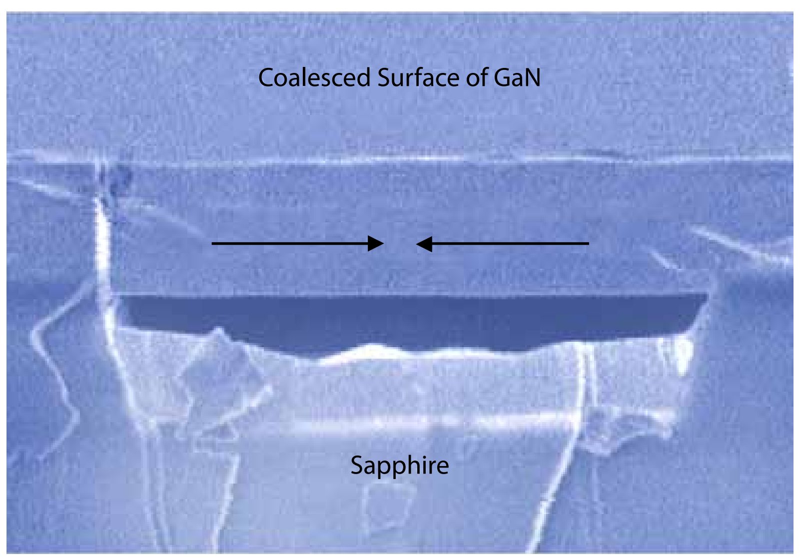

Electron micrograph of cantilevers growing together over a trench.

(SiC), a somewhat better substrate upon which to grow GaN, because of its thermal properties, was just beginning to come into use and was prohibitively expensive. Thermal management, i.e., heat dissipation, is an important consideration because the bandgap properties of materials can be altered by excessive heat. Had SiC been more readily available, it is possible that CE, applied to SiC as a substrate for GaN growth might have become a more-significant player in the marketplace. Additionally, although CE is economically viable as a process for the creation of diode lasers, predictions are that more-direct methods for the creation of wafers of GaN and other III-V materials may ultimately replace all heteroepitaxial growth processes. At this point in time, there are, as yet, significant uncertainties regarding both the exact materials that will be appropriate for various SSL applications and also the processes that will ultimately create those materials. Consequently, the availability of a process such as CE is potentially a boon for further materials development. And, of course, LDRD continues to fund relevant research in semiconductor materials science, one goal to sustain Sandia's leadership in the solid state lighting arena, given its position as lead laboratory of the National Center for Solid State Lighting R&D. The importance of this US initiative is underscored by the observation that eventual worldwide replacement of incandescent and fluorescent lighting by SSL of 50% efficiency (70% considered an ultimately achievable goal) would have an enormous impact on electrical energy consumption and hence, potentially, on global greenhouse-gas emissions. ♦

Writer
Vin LoPresti
Sandia Staffing Alliance, LLC

Graphic Design
Chris Brigman
Global Analytic IT Services

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

SAND No. 2008-7400P