


DARPA

DEFENSE ADVANCED RESEARCH PROJECTS AGENCY

1958-2018



60
YEARS
1958 — 2018



60
YEARS
1958 — 2018

DEFENSE ADVANCED RESEARCH PROJECTS AGENCY

Cover photo: Shown with artistic flair is an injector plate of an F-1 engine, the workhorse engine that would, among other lifts, power the Saturn V rockets of the Apollo missions. From the injector plate's holes, liquid fuel and liquid oxygen would spray out, much like water from the head of a garden hose, but under enormous pressure. This particular F-1 engine is on display in Huntsville, Alabama, at the U.S. Space and Rocket Center. In August 1958, just six months after ARPA was established (later renamed as DARPA), the agency authorized a proposal by Wernher von Braun and his research team at the Army Ballistic Missile Agency in Huntsville to design and build a large, heavy-lift rocket vehicle. To quickly and cheaply achieve massive thrust in the first stage, ARPA suggested a design featuring a cluster of available rocket engines, the powerful F-1 developed by Rocketdyne in the mid-1950s. Also expediting the successful development of the Saturn booster was the reliance in the upper stages on liquid hydrogen technology developed earlier for the ARPA-supported CENTAUR vehicle. As DARPA enters its seventh decade, the agency remains at the forefront of rocket design, now focusing on the challenges of boosting assets to orbit quickly, with little notice, and at low cost. Photo by Lee Hutchinson



DEFENSE ADVANCED RESEARCH PROJECTS AGENCY
675 NORTH RANDOLPH STREET
ARLINGTON, VA 22203-2114

September 5, 2018

On February 7, 1958, the day the Advanced Research Projects Agency was officially established, humanity had just entered the Space Age. In all, the United States and our Cold War rival, the USSR, had managed to place three satellites into orbit by that date, and one of those already had returned to Earth. Today, thousands of sophisticated satellites are overhead, providing us with critical intelligence, meteorological data, navigational tools, global communications, and many other technological wonders that have become necessities. Hints of the coming microelectronics age, in the form of the world's first integrated circuits (ICs), also rolled out during ARPA's maiden year. Sixty years later, multibillion-transistor ICs are commonplace, and the artificial intelligence they are making possible appears slated to open our roads to millions of driverless cars; to identify hidden threats amidst oceans of data; and to provide machine-assistance for diagnosing disease, making scientific discoveries, and designing new weapon systems.

We are honored to know that the agency we lead has been one of the most consequential drivers of these and other technological evolutions and revolutions over the past 60 years. The seeds of many technologies in our lives today were sown by DARPA (the acronym acquired a D for "Defense" in 1972), and the same is happening now for the technologies of the future.

It is our pleasure to share this publication with you to mark DARPA's first 60 years. On these pages, you will encounter an eclectic mix of articles that chronicle the agency's past and present, and that project some of the exciting and potentially world-changing directions in which DARPA is heading. Written by both insiders and outsiders, the articles collectively reveal many angles on the institutional innovation that DARPA has brought to national security and beyond.

DARPA's storied record of accomplishment is a synergy of people, partnerships, and projects. It is our uncommonly creative and capable program managers who bring to DARPA unique expertise, knowledge, and experience in the technological territories that matter for national security. It is through partnerships with academia, industry, and warfighters within the Department of Defense that DARPA embraces the opportunity to change the world. And it is by way of the Agency's projects (the P in DARPA) that engineering imagination and drive become hold-in-your-hand technological capability for the country.

You will read about these critical components of DARPA in these articles. But the book on DARPA is open-ended. Those chapters not yet written—and that will show up in future anniversary publications and could well include readers like you as major characters—will be informed by the constant kaleidoscope of people, partnerships, and projects that for the past 60 years has been the right stuff of DARPA.

Steven H. Walker
Director

Peter Highnam
Deputy Director



Congressional Record

PROCEEDINGS AND DEBATES OF THE 115th CONGRESS, SECOND SESSION

Vol. 164

WASHINGTON, WEDNESDAY, FEBRUARY 7, 2018

No. 24

Senate

DARPA'S 60TH ANNIVERSARY

Mr. REED. Mr. President, today I pay tribute to DARPA, the Defense Advanced Research Projects Agency, on the 60th anniversary of its inception. After the Soviet Union launched Sputnik I, President Eisenhower determined that the United States would never again be caught off guard by technological surprise. DARPA was established to anticipate new technological capabilities and pursue strategic technological surprise for our military forces.

DARPA works collaboratively with academic institutions, corporate and government R&D labs, and small business enterprise. While the primary focus is to discover fundamental new concepts that lead to breakthrough technologies for national security, many of DARPA's advances also benefit greater society. Some well-known examples include precision-guided weapons systems with miniaturized GPS components also found on many consumer products; the internet, used initially to link DARPA with performer partners, now widely used in commerce and every aspect of our lives; advanced antenna systems enabling more efficient warfighter communications and satellite signal reception for consumers; new breakthroughs in robotic technology for national security applications and the development of advanced prosthetic arms for wounded warriors and civilians alike. The list goes on.

By not accepting the parameters of what is widely accepted as the known possible, DARPA has proven that amazing achievements can be had by stretching to reach for what was once deemed impossible. In the realm of national defense, DARPA has pursued new systems, including unmanned aerial and underwater vehicles, hypersonic flight research, and new frontiers in biomedical research. From the giant engines of the

Saturn V rocket that took Americans to the Moon to the smallest microelectronics that populate our smartphones, DARPA has been ahead of the cutting edge of technological innovation.

By focusing its efforts at the boundaries of fundamental research in physics, chemistry, biology, mathematics, materials science, electronics, and engineering, DARPA has helped create new communities of scientists and engineers, both inside and beyond the traditional defense community. Along the way, new businesses and sometimes entire industries have sprung from DARPA-funded research, reflecting the Agency's commitment to pursue its ideas all the way from initial concept to demonstration of practical feasibility through prototype development.

DARPA programs are led by program managers who come from universities, industry, national laboratories, and other parts of government for limited postings that typically last 3 to 5 years—a time limit that helps drive the Agency's signature sense of urgency. Recognizing that some revolutionary goals inevitably prove unachievable, DARPA carefully manages risk by establishing appropriate milestone procedures and redirecting or discontinuing programs when further advancement stalls.

I congratulate DARPA for its many achievements over the past 60 years. The true assets that enable this kind of achievement are the men and women who work to make the visions of tomorrow become today's reality.

As DARPA moves into the future, I encourage my colleagues to join with me in recognizing this milestone and supporting DARPA so that it can continue to keep our warfighters and citizens at the leading edge of technology and out of harm's way.

DARPA

DEFENSE ADVANCED RESEARCH PROJECTS AGENCY

1958-2018

CONTENTS

8

DARPA DIRECTORS

10

REFLECTIONS OF FORMER DIRECTORS

12

DARPA

The innovation icon at 60

By Richard H. Van Atta

22

DARPA TILES TOGETHER A VISION OF MOSAIC WARFARE

Banking on cost-effective complexity to overwhelm adversaries

By Stew Magnuson

26

FURTHERMOORE

The intertwined history of DARPA and Moore's Law

By William Chappell

32

THE STUFF OF NATIONAL SECURITY

DARPA and the invention of new materials

By Ivan Amato and Steve Wax

38

GIANT STEPS

DARPA's X-planes and the quest to redefine the boundaries of flight

By Jan Tegler

46

DARPA AT THE TACTICAL EDGE

TIGR & RAA: Information from the bottom up

By Eric Tegler

52

SEMICONDUCTOR SAFARI

Exotic materials beyond silicon

By Dan Green

56

JUMP-STARTING INNOVATION

DARPA's Grand Challenges

By Ivan Amato

62

DARPA'S QUEST FOR A BENEFICENT CYBER FUTURE

The innovations it will take to assure our cyber security

By Brian M. Pierce

68

ENTERPRISE DISRUPTION

Shaking off entrenched ways is hard, but it is the gateway to better technology frameworks for all of the military domains

By Fred Kennedy

74

SECURITY AND SURPRISE AT BIOLOGICAL SCALES

The evolution of defense technology

By Justin Sanchez and Jacob Jordan

80

AUTONOMOUS TECHNOLOGY COMETH

The advent of machines that act and decide will benefit from the lessons humanity has learned from itself

By J.R. Wilson

86

FIGHTING IN MEGACITIES

DARPA takes on the challenge of warfare in expanding urban settings

By Stew Magnuson

90

TAKING NEUROTECHNOLOGY INTO NEW TERRITORY

With new interface technology, the human-machine relationship is deepening

By Justin Sanchez and Robbin Miranda

DARPA

DEFENSE ADVANCED RESEARCH PROJECTS AGENCY

1958-2018

CONTENTS

96

A PROTEAN TECHNOLOGY

DARPA and the exploration of artificial intelligence

By John Everett

103

THE NETWORK OF OUR TIMES

A DARPA perspective on the development of the internet

By Henry Kenyon

110

SPINTRONICS

A DARPA-spurred spin on fundamental electron physics just keeps on giving

By Ivan Amato

116

TECHNOLOGY THAT MATTERS

Ushering game-changing technology into the real world

By Craig Collins

122

DARPA'S 60-YEAR SPACE ADVENTURE

A maker of the Space Age continues its expansive take on space technology

By Edward Goldstein

130

ARTIFICIAL INTELLIGENCE TO ACCELERATE SCIENCE

How in-silico minds will open new pathways to discovery

By Paul Cohen

134

UNIVERSITIES IN SERVICE TO NATIONAL SECURITY

DARPA's call to academia

By William Regli

DARPA

DEFENSE ADVANCED RESEARCH PROJECTS AGENCY

1958-2018

Published by Faircount Media Group
4915 West Cypress Street
Tampa, FL 33607
Tel: 813.639.1900
www.defensemedianetwork.com
www.faircount.com

EDITORIAL

Executive Editor: Ivan Amato
Editor in Chief: Chuck Oldham
Managing Editor: Ana E. Lopez
Editor: Rhonda Carpenter

Contributing Writers: Ivan Amato, William Chappell,
Paul Cohen, Craig Collins, John Everett, Edward Goldstein, Dan Green,
Jacob Jordan, Fred Kennedy, Henry Kenyon, Stew Magnuson,
Robbin Miranda, Brian M. Pierce, William Regli, Justin Sanchez,
Eric Tegler, Jan Tegler, Steve Wax, J.R. Wilson, Richard H. Van Atta

DESIGN AND PRODUCTION

Art Director: Robin K. McDowall
Designer: Rebecca Laborde

ADVERTISING

Advertising Sales Manager: Geoffrey Weiss
Representatives: Art Dubuc, Damion Harte,
Bryan St. Laurent, Geoffrey Weiss

OPERATIONS AND ADMINISTRATION

Chief Operating Officer: Lawrence Wayne Roberts
VP, Business Development: Robin Jobson
Business Development: Damion Harte
Business Analytics Manager: Colin Davidson
Accounting Manager: Joe Gonzalez

FAIRCOUNT MEDIA GROUP

Publisher: Ross Jobson



©Copyright 2018 United States Government, as represented by the Secretary of Defense. All rights reserved. Published by Faircount LLC, a private firm in no way connected with the Defense Advanced Research Projects Agency or the Department of Defense, under exclusive written agreement with the Defense Advanced Research Projects Agency. This commercial enterprise publication is an authorized publication for members of the U.S. military services and DoD civilian employees. Contents of the publication are not necessarily the official views of, or endorsed by, the U.S. Government, DoD, or the Defense Advanced Research Projects Agency. The editorial content of this publication is the responsibility of the DARPA Public Affairs Office. The appearance of advertising in this publication, including inserts or supplements, does not constitute endorsement by DoD, the Defense Advanced Research Projects Agency, or Faircount LLC of the products or services advertised. Everything advertised in this publication shall be made available for purchase, use, or patronage without regard to race, color, religion, sex, national origin, age, marital status, physical handicap, political affiliation, or any other non-merit factor of the purchaser, user or patron.

SIX DECADES OF DARPA DIRECTORS



Steven Walker
2017-Present



Arati Prabhakar
2012-2017



Regina E. Dugan
2009-2012



Anthony Tether
2001-2009



Fernando L. Fernandez
1998-2001



Verne L. "Larry" Lynn
1995-1998



Gary L. Denman
1992-1995



Victor H. Reis
1990-1992



Craig I. Fields
1989-1990



Ray S. Colladay
1988-1989



Robert C. Duncan
1985-1988



Robert S. Cooper
1981-1985



Robert R. Fossum
1977-1981



George H. Heilmeier
1975-1977



Stephen J. Lukasik
1970-1975



Eberhardt Rechten
1967-1970



Charles M. Herzfeld
1965-1967



Robert L. Sproull
1963-1965




Jack P. Ruina
1961-1963



Austin W. Betts
1960-1961



Roy W. Johnson
1958-1959



REFLECTIONS from former directors

“DARPA is a place of doing. It is a place where vision is paired with execution, where doing is as powerful a force as thinking, and where it is commonplace for something once deemed impossible to become improbable and then eventually inevitable.”

— Dr. Regina Dugan

“The thing that DARPA does, and DARPA does well, is show that something can be done... And once you show somebody it can be done, it’s amazing how many people suddenly are brave enough to go and try it.”

— Dr. Tony Tether

“DARPA is as true to its founding mission today as it was in 1958. The DARPA brand is an icon for technological innovation, entrepreneurial spirit, and creative solutions to national defense needs.”

— Dr. Raymond Colladay

Early in the summer of this year, DARPA's Experimental Spaceplane advanced toward its goal for rapid turnaround and on-demand launch with the completion of a rigorous series of 10 engine firings in 10 days for the spaceplane's main RS-25 engine.

“DARPA is an expression of the quintessence of America. A place where individuals reach beyond their grasp, and where the only failure is small dreams. A home for those who yearn to expand the frontier of the possible. An enterprise imbued with an abiding belief that tomorrow can be better than today, if we make it so.”

— Dr. Arati Prabhakar

“There has never been a better mechanism for accelerating the development of technology than DARPA. Why? A culture of acceptance of risk and failure in pursuit of great achievement. DARPA technology staff are irrepressible and compulsive – ‘the best and the brightest’ – and DARPA management staff are intent on removing barriers, not erecting barriers to getting things done ... DARPA accelerates time, DARPA brings the future faster. It makes things happen that would otherwise take years, a decade, perhaps two decades.”

— Dr. Craig Fields

“DARPA was unique in its exposure to hard, important national technical problems, in its freedom from military intervention, and most importantly in its exceptional and dedicated technical and support staff. We were not bound by short term requirements but were instead allowed to work on a number of ‘what if’ problems, some of which materialized in the future when we went to war against terrorists all over the world.”

— Dr. Fernando “Frank” Fernandez

“ARPA played roles in numerous conflicts. We worked with the Air Force on strategic offense and with the Army in strategic defense. We added programs in naval technology to ensure the survivability of our sea-based forces and devoted attention to working with Combatant Commands and the Intelligence Community.”

— Dr. Stephen Lukasik

“The people that I managed and worked with at DARPA were some of the brightest and most motivated people I ever had the pleasure to work with. The interactions with program managers were extremely lively as most had great ideas and most thought I had a bucket of money under my desk. Some had outrageous ideas but were passionate about them. DARPA, in my opinion, is a national treasure and is absolutely essential to maintaining our technology edge in military capabilities.”

— Dr. Gary Denman

DARPA

The innovation icon at 60

By Richard H. Van Atta¹

The Defense Advanced Research Projects Agency (DARPA) is an “innovation icon.” It has created many breakthrough technologies that have revolutionized defense, including stealth, unmanned aerial systems, and precision-guided munitions. Many DARPA-developed technologies have had sweeping societal and economic impacts, including portable GPS receivers, new types of computer chips, voice-recognition software, interactive and personal computers, and, most famously, the ARPANET and its successor, the internet. Current DARPA research also may have dramatic future impacts, including self-driving vehicles, robots and exoskeletons, and cognitive computing – computers that emulate brain-like processing. In addition, DARPA has fostered the creation of important technical communities, including the fields of materials science, computer networking, and engineering biology.

How and why has DARPA succeeded? Which features of its operation and environment contribute to this success?² How can it succeed today and in the future in a vastly changed technological and security environment?

This introduction focuses on DARPA's key features – its mission, organization, linkages to other organizations, and “political design” – and how these have contributed to its success. It will review the evolving foci of DARPA's research and conclude with some lessons for the agency as it strives to continue to deliver transformative technologies in the coming years.

Over its 60 years, DARPA has evolved and, on a couple of occasions, come close to being dissolved. Its organizational structure, as well as some important operational mechanisms, have changed. There is no singular accurate depiction of DARPA because it has adapted based on how the world has changed – especially the national security environment. At a given point in time, one can discern several DARPAs as different parts of the organization have embraced different technologies and development approaches. DARPA's history has been perturbed by political dynamics as well as the dynamics of the technologies it has pursued. Perhaps the most important hallmark of DARPA has been its adaptability and flexibility to respond, sometimes extremely rapidly, to changing circumstances.

DARPA's Origins – 1958-1970


In October 1957, the Soviet Union launched the first artificial satellite, Sputnik I, an event that shocked the United States. Many Americans worried that the country was losing technological leadership to its Cold War adversary. After the launch of the first Sputnik, President Dwight D. Eisenhower followed the advice of Secretary of Defense Neil McElroy and leading scientists to propose the creation of the Advanced Research Projects Agency (ARPA).³

Initially the agency focused on three assignments from the president: space, missile defense, and the detection of nuclear weapons tests. However, Eisenhower emphasized that space was to be the realm of a civilian agency. Toward that end, later in 1958, Congress and the president created the National Aeronautics and Space Administration (NASA), a civilian agency responsible for the country's principal space programs. As such, NASA absorbed



On Oct. 4, 1957, Sputnik 1 was successfully launched and entered Earth's orbit. Thus began the Space Age. The successful launch shocked the world, giving the former Soviet Union the distinction of putting the first human-made object into space. The surprise advance by the United States' Cold War nemesis helped catalyze the establishment of ARPA in February 1958.

NASA PHOTO/ASIF A. SIDDIQI



much of DARPA's space program. The two other presidential assignments for ARPA, missile defense and nuclear test detection, continued as the dominant foci for about 15 years.

Soon after its founding, ARPA took on Project AGILE – a decade-long, highly classified program supporting U.S. efforts in Vietnam – proposed by its Deputy Director, William Godel. In retrospect, much of AGILE was naive, poorly managed, and rife with amateurism. AGILE was run with little oversight; it was hardly scientific and focused on near-term solutions. Thus, it serves now as a lesson of what DARPA should not be.

With the transfer of the space program to NASA, ARPA spent the 1960s focused on missile defense, nuclear test detection, and AGILE. In the early 1960s, however, ARPA began to pursue a set of smaller, technically focused programs to “prevent technological surprise.” Initial programs were in materials science, information technology, and behavioral science. Arguably, ARPA “invented” these as areas of technological pursuit. For example, in 1961 ARPA Director Jack Ruina hired J.C.R. Licklider as the first director of the Information Processing Techniques Office (IPTO), which played a vital role in the creation of personal computing and the ARPANET – the basis for the future internet.

DARPA's Resuscitation in the 1970s

While there were successes in these early pursuits, by the mid-1960s DARPA was a troubled agency – a victim of the Vietnam malaise, resource cutbacks that affected all of the Department of Defense (DOD), and the fact that its post-space program thrusts – missile defense (Defender) and nuclear test detection (Vela), as well as AGILE – had essentially run their respective courses. Indeed, in 1965, Deputy Secretary of Defense Cyrus Vance “came to advocate abolishing the agency.”⁴ However, DARPA evolved both organizationally and programmatically from this crisis under Director of Defense Research and Engineering (DDR&E) John Foster. By the mid-1970s, DARPA had jettisoned AGILE and transitioned Defender to the Army. DARPA sought new directions first under Director Eberhart Rechtin, who created a Strategic Technologies Office, and his successor Steven Lukasik, who closed AGILE down for the most part, though salvaging portions of it within a new Tactical Technology Office. Thus, by the mid-1970s, DARPA leadership had substantially restructured the agency around technology offices rather than mission-focused assignments.

ABOVE: On Feb. 7, 1958, Defense Secretary Neil McElroy issued DOD Directive 5105.15 establishing the Advanced Research Projects Agency (ARPA), later renamed the Defense Advanced Research Projects Agency (DARPA). The agency's first three primary research thrusts focused on space technology, ballistic missile defense, and detection of nuclear weapons tests. LEFT: ARPA-funded research fueled the nation's early efforts in the Space Race, including development of the Saturn launch vehicle. Innovations in the Saturn I shown here led to the Saturn V that launched astronauts to the moon.





U.S. ARMY PHOTO

ABOVE: Pfc. Michael J. Mendoza fires his M16 rifle into a suspected Viet Cong occupied area on Sept. 8, 1967. Among the sub-projects of Project AGILE was development of the M16 rifle, which still serves the U.S. armed forces today. RIGHT: The ARPA Vela program developed sensors to detect nuclear explosions in space, the upper atmosphere, and underwater to support the 1963 Limited Nuclear Test Ban Treaty, under which the United States, Great Britain, and the Soviet Union banned atmospheric tests of nuclear weapons.

DARPA was rejuvenated when it took on the job of seeking technological alternatives to nuclear weapons to respond to the Soviet Union's force posture in Europe. This initiative was highly important to President Richard M. Nixon and his National Security Adviser Henry Kissinger, and continued under President Gerald R. Ford and his Secretary of Defense James Schlesinger. Under Lukasik, DARPA formulated a Long Range Research Planning Program in 1973 that guided DARPA research programs for the next 20 years. Among these were network-centric warfare, stealth technology, precision and autonomous systems, and theater missile defenses. In 1975, DARPA Director George Heilmeier — most famous for his “catechism” of questions he asked program managers to help determine the worthiness of the programs they proposed — scaled up major efforts in stealth aircraft, space-based lasers, space-based infrared technology, and artificial intelligence (AI).



DARPA IMAGE


DARPA in the 1980s – Transformative Technology Development and Transition

In the 1980s, DARPA developed and demonstrated new tactical capabilities through programs on stealth, standoff precision strike, and tactical surveillance via unmanned aerial vehicles (UAVs). Through persistent efforts, working with the DDR&E and the Secretary of Defense, DARPA worked to transition these capabilities to the military, creating what Under Secretary of Defense William Perry and Secretary of Defense Harold Brown (under President Jimmy Carter) would call the “Offset Strategy.” These DARPA programs are among its most important programs in terms of the agency’s impact on defense capabilities and are touted as DARPA’s impact in ushering in a “revolution in military affairs” – that offset Soviet military capabilities by transforming tactical warfare.⁵

Parallel to DARPA’s transformational military programs in the 1970s and 1980s were programs revolutionizing information technology, building on Licklider’s vision of “man-computer symbiosis.” DARPA’s research was foundational to computer science. ARPANET was one element of a much broader, increasingly coherent program based on the technological future that Licklider imagined. He and his IPTO colleagues conceived a multi-pronged development of the technologies underlying the transformation of information processing from clunky, room-filling, inaccessible mainframe machines to a ubiquitous network of interactive and personal computing capabilities.⁶ This transformation continues today in DARPA’s pursuit of artificial intelligence, cognitive (brain-like) computing, and robotics.

DARPA in the 1990s – The Cold War Ends

Early in the 1990s, DOD, and with it DARPA, adapted to the fact that the United States’ primary adversary, the USSR, had collapsed. With this momentous change in the world order, the focus of DARPA’s weapons research had disappeared. Moreover, the United States was in a budget crisis partly due to the vast defense spending of the 1980s. The Clinton administration employed the rubric “dual-use” – technologies that would have both defense and civilian economy payoffs – aiming to make the economy more competitive, while maintaining leadership in defense technologies. Under this approach, DOD sought to leverage the civilian sector to cut costs to develop new technologies. These dual-use programs were a major redirection of DARPA (renamed ARPA in 1993 to emphasize the dual-use focus), a shift that became highly contentious with elements in Congress. The Technology Reinvestment Program (TRP) was created to promote collaborations amongst defense technology developers, commercial firms, and universities. Despite the lack of a peer adversary, the Office of the Secretary of Defense (OSD) and DARPA worked with the White House to continue DARPA’s exploration and development of “breakthrough” technologies. Programs in unmanned systems and precision-strike capabilities continued with Predator and Global Hawk UAVs, and the agency fostered new technologies in information and electronics, advanced sensing, and battlefield surveillance. Programs in biotechnology were started. DARPA capped off the 1990s by taking on a program with the Army for a radical approach to ground combat using networked robotics – the Future Combat System – which proved



The log entry for the IMP (Interface Message Processor) at UCLA, recording the first ARPANET message. It was sent on Oct. 29, 1969, at 22:30 Pacific Time from Boelter Hall 3420.

unsuccessful. It was overly ambitious, rushed into acquisition by the Army, and eventually cancelled by the Secretary of Defense.

DARPA in the 2000s – War on Terror

At the outset of the 2000s, DARPA continued embracing major technology themes of precision strike, networking of operations, advanced sensors, and unmanned systems. However, within months of Anthony “Tony” Tether becoming director, the terrorist attacks of Sept. 11, 2001, occurred and DARPA became enmeshed in the “War on Terror.” The Total Information Awareness (TIA) program became the most notable DARPA response. This was a highly controversial program, as the use of information technologies to identify possible terrorists and terrorist attacks raised issues of privacy, and it was terminated amidst this controversy. DARPA also developed fast-response programs to support combat needs in Iraq and Afghanistan, including tactical robots and sensor systems. Secure, robust, self-healing sensor and communications networks also became a priority. Two new offices – the Information Awareness Office and the Information Exploitation Office – were created to address the terrorist threat. The agency pursued an array of robotics programs and related efforts in human-machine interaction, such as exoskeleton technology and advanced prosthetics. During this period, DARPA also emphasized programs in artificial intelligence, cognitive computing, quantum computing, and autonomous systems. The “DARPA Challenges” for self-driving cars in 2004, 2005, and 2007 were highly visible demonstrations of autonomous capabilities and were successful in creating interest and incentivizing teams of researchers. There was a renewed emphasis on “assured use of space” illustrated by the ambitious Orbital Express program for robotic on-orbit refueling and reconfiguring of satellites.

DARPA in the 2010s – Technology for Security in a Globalized World

In the current decade, DARPA has recognized that technology has become globalized. In the past, the United States and DOD led in



The Orbital Express program proved the concept of robotic satellite repair, refueling, and maintenance on orbit.

technology development; now global competitors pursue many of the very technologies that DARPA pioneered. With growing peer competition in the security arena, and terrorism an ongoing concern, DARPA's mission of avoiding technological surprise and creating technological surprise for adversaries is more daunting and pressing than ever.

One area of increased DARPA focus is the life sciences. While DARPA had been conducting biotechnology research for some time, in 2014 it created a new Biological Technologies Office with one of its thrusts centered on synthetic biology, a form of genetic engineering that opens new opportunities for such applications as sensing hazardous compounds and the efficient bio-production of novel coatings, fuels, and drugs. Another driver for biological technologies was DOD's own massive health care concerns as injured soldiers from two wars required new medical responses. DARPA's flexibility enabled it to pursue a "convergence" approach combining biology with engineering, physical, and computational sciences to seek new kinds of therapies, including advanced prosthetic technologies.

In the information domain, DARPA is focused on artificial intelligence, cyber threats, and approaches for advancing microelectronics, including efforts in quantum computing and neuro-synaptic processors based on how the brain processes information. With a foundation on previous research in aeronautics and propulsion, DARPA is embarking

on a major thrust in hypersonic systems. Meanwhile, growing cyber threats spurred several ambitious DARPA programs in cybersecurity.

This review shows that DARPA's technical and security foci have changed with the times, although its mission—ensuring technological leadership for the United States—has remained largely the same. To carry out this mission today, however, the agency must focus on creating and demonstrating breakthrough technologies for national security even as many highly capable players enter the technology competition and as technology disseminates globally more swiftly than ever.

DARPA's Organization and Approach

To achieve its mission, DARPA has been a highly adaptive, responsive organization. The hallmark of DARPA is agility. DARPA is organized into technology offices, the number and specific roles of which have changed over time. Today DARPA has six technical offices:

- Biological Technologies Office (BTO)
- Defense Sciences Office (DSO)
- Information Innovation Office (I2O)
- Microsystems Technology Office (MTO)
- Strategic Technology Office (STO)
- Tactical Technology Office (TTO)

DARPA is a funding agency; it has no laboratories or research staff of its own. At the heart of DARPA are approximately 100 program managers (PMs) and the office directors, deputy office directors, and agency director and deputy director who supervise them. These are all government employees,



Under a DARPA contract, the Rochester Institute of Technology developed the Blast Gauge, a small device worn by warfighters to measure blast exposure and cue medics for initial response. This phase of the project took just 11 months with a total development cost of approximately \$1 million.

|||||


most hired using special hiring authorities on a term basis — usually of three to five years. Importantly, none of these is permanent staff — all are, in essence, temporary.⁷ The agency's budget is approximately \$3 billion a year.

Program managers have great autonomy: they propose their programs, seek approval and funding from senior DARPA officials, write the funding solicitations, select the R&D performers, and supervise and assist these performers. A typical program will have specific

technical objectives, a budget of tens of millions of dollars, and will last three to five years. Often an individual program will fund multiple R&D projects run by different performers so as to compare and test different technical ideas.

Each DARPA technology office also can fund small “seed” programs, which provide a way for program managers to generate and test new ideas. In recent years, each office also has run an annual “open” competition in which applicants can propose work in areas of technology not covered in the office's current programs.

DARPA uses a “portfolio” approach: it funds a wide range of R&D programs and often funds multiple projects within a single program. Since the focus of DARPA is high-risk programs to achieve “high payoff” results, the agency and its program managers invest in a range of promising technologies to increase the chances of success while accepting the inherent risk that some research may not succeed.



How DARPA Transfers Its Technologies

As an advanced research organization, DARPA does not fund the construction of actual operational systems; it succeeds by developing and demonstrating “proof-of-concept” prototypes that others in defense and/or the commercial world can further develop, modify, and implement – a process that DARPA calls “technology transition.”⁸ DARPA succeeds when other organizations in government and the corporate world further develop, and then commercialize and buy the new potentially transformative technologies that it has fostered.⁹

DARPA also succeeds when its efforts lead to an understanding of whether something can be done, how well it can be done, or if it cannot be done. What features of the DARPA model and the overall national innovation system help technology transition?¹⁰

DARPA's willingness to challenge incumbent technologies. DARPA is willing to challenge existing technologies and the organizations that produce and use them – it sees its job as changing people's minds about what is possible. For example, it showed that a computer network using open standards could replace proprietary networking systems. It created and then, with support from the Office of the Secretary of Defense (OSD), pushed for the adoption of stealth, unmanned aerial systems, precision strike, and night vision. It uses conferences, prize competitions, “technology insertion projects” (demonstrations of new technology in actual military systems), and other techniques to demonstrate and publicize new technical capabilities.

A community of technology advocates. DARPA and its performers have created new technical communities. Besides helping DARPA undertake new research, investigators in these new communities also often become knowledgeable, enthusiastic advocates. Some of these experts work in government, some in universities, some for large firms, and some start new entrepreneurial companies. They share an overall vision of what can be done, and they often become “communities of change-state advocates” – people who are willing and able to change

Researchers working in DARPA's Microphysiological Systems program have been devising “organs on a chip,” which rely on cell-lined microfluidic labyrinths for rapid, efficient, and animal-free testing of biomedical countermeasures to natural and human-made threats, including emerging infectious diseases and bioweapons.

the technology world. This is one reason that DARPA has been so influential.

Close ties to DOD leaders. DARPA works on problems important to high-level leadership. The agency's close ties to the Secretary of Defense and other senior officials not only help DARPA maintain its independence, but these officials also can become “champions” who want to further develop the technologies and see them transitioned and deployed. For example, senior DOD officials pushed the U.S. Air Force to adopt both stealth aircraft and unmanned aerial vehicles, which DARPA programs had developed and demonstrated.

Connection to technically sophisticated customers. Turning radical new technologies into new operational capabilities is usually risky, difficult, and expensive. DARPA and the Department of Defense deal with this in two ways. First, the DOD itself can turn new prototype technologies into actual products. The department's senior leaders may want advanced technologies, and through its laboratories, contractors, and acquisition system, it can refine and buy these new products. However, “transitioning” new technologies from DARPA to the military services is often difficult because DARPA-developed capabilities are frequently “disruptive,” challenging the current ways of conducting operations. Successful transition into a military capability, therefore, is usually preceded by considerable time and effort by DARPA managers.

A second pathway for transitioning is when the agency works with the private sector in commercializing DARPA-demonstrated technologies. DARPA repeatedly has worked with the semiconductor industry to adopt advanced chip technologies, for example, silicon-on-insulator technology



ABOVE: Former DARPA Director Arati Prabhakar speaks during the award ceremony after the world's first all-machine hacking tournament in Las Vegas on Aug. 4, 2016. Seven teams competed in the capture-the-flag event and three of them won cash prizes. RIGHT: The Running Man robot of Team IHMC Robotics from Pensacola, Florida, clears a doorway during the DARPA Robotics Challenge. Twenty-five teams from around the world competed in 2015 for \$3.5 million in prizes as they navigated a simulated disaster-response course.

and Monolithic Microwave Integrated Circuit signal processing chips. The new commercial frontier of self-driving vehicles is another example of an industry adopting and building upon DARPA-funded research.

Sustaining the DARPA Vision

DARPA's higher-risk, longer-term R&D agenda distinguishes it from other defense R&D organizations. Perhaps the most important effect of DARPA's work is to change people's minds as to what is possible. DARPA's 60-year history reveals an institution driven by a constant imperative to create novel, high-payoff capabilities by pushing the frontiers of knowledge. DARPA has many of the same features as its research. DARPA began as an experiment aimed at overcoming the usual incremental processes of technology development. Like the research it is chartered to develop, DARPA consistently has been purposively "disruptive" and "transformational." Over the decades, there have been various efforts to tone down DARPA, make its research more compatible and integrated into the rest of DOD R&D, and have it focus more heavily on nearer-term, more incremental applications – that is, to shift its focus away from disruptive possibilities. Also, there have





Former Defense Secretary Ash Carter, center, speaks with a group of early-career “rising stars” over lunch on Sept. 9, 2015, as he attends the DARPA *Wait, What?* future technology forum in St. Louis, Missouri.

been efforts to broaden its charter into prototyping systems beyond the proof-of-concept demonstrations DARPA traditionally has carried out. However, with strong internal leadership, both within DARPA and in the OSD, as well as with support from Congress, DARPA has been able to perform a truly unique role for six decades. It has been and continues to be DOD’s “Chief Innovation Agency,” pushing the frontiers of what is possible for the benefit of national security and the nation.


DARPA remains an impressive “opportunity farm.” For example, DARPA helped move AI from an inchoate notion with almost no technological underpinnings into pervasive capabilities affecting our everyday lives and supporting real-time military operational decision-making. It is now pursuing similar advances in electronics and autonomous technologies. It is pursuing fundamental advances in materials, such as biomaterials, and accelerated materials development. The agency has revolutionized the realm of distributed sensing. Among many current DARPA research topics that populate the opportunity farm are heterogeneous electronics, engineering biology, agile access to space, and hypersonic systems.

Looking to the future, the question is not whether DARPA can still pursue new change-state prospects. The question is this: How can DARPA and the Department of Defense identify and focus on what these should be in the changing geopolitical and technological environments? DARPA

has been adroit in addressing emerging technological prospects – but in today’s world, it has to be even more focused on where it can have leverage as others are investing, often massively, in the very technologies that DARPA initially championed.

With global investments in robotics, AI, synthetic biology, quantum computing, and advanced materials, on what should DARPA focus? From a military applications perspective, what should DARPA do to harness and promote the potential use of such emerging technologies into defense uses? To what needs, as opposed to today’s defined requirements, should DARPA seek to employ these technologies? Crucially, how should DOD go about achieving the fruition of these efforts?

Today DARPA faces new challenges that raise a key issue concerning its future success – the ability to draw upon extraordinary technical talent for program managers. The commercial high-tech sector, particularly in such areas as information technology, autonomous systems, advanced biology – areas in which DARPA is focused – is aggressively spending vast sums and hiring the very best. These firms attract this talent with high salaries and relatively unfettered work environments in locations far from Washington, D.C. Foreign nationals make up a growing proportion of the hires. Moreover, many leading tech companies are now outside of the United States, in Asia and Europe. Thus, there is greater competition for technical talent and greater competition worldwide in advanced technologies. Even as DARPA now must confront a tougher recruiting context than it has had to in the past, it still presents prospective program managers unique opportunities to affect the future that few other organizations



DARPA's research and development in stealth technology during the 1970s and 1980s led to the world's most advanced radar-evading aircraft, providing strategic national security advantage to the United States. Today, hypersonic technologies have the potential to provide the dominance once afforded by stealth to support a range of future national security missions.

can offer. These talent dynamics are crucial to understanding today's DARPA and its ongoing mission to identify, demonstrate, and develop the technologies of the future.

These questions emphasize a crucial point – DARPA does not succeed by itself. Its success resides in the opportunities it creates that others bring to fruition. Thus, its success must build upon the larger U.S. innovation infrastructure. That innovation ecosystem has changed fundamentally over the past 25 years. For DARPA research to be successful, it must eventually culminate in transition, whether in an operational military capability or a new field of technology that expands frontiers for decades. DARPA itself is not responsible for executing transitions, but it depends on effective transition paths being there. These paths need to be better understood, and other stakeholders, beyond DARPA, need to support the measures that foster transition – whether within DOD or within industry. Some worry that military transition mechanisms within the DOD have eroded. In the broader commercial economy, transition paths have become more uncertain and diffuse. For DARPA to continue to have transformative impacts, it must exist within an economic and policy environment that encourages implementation. These are critical technology policy concerns that the United States must address to ensure that DARPA can continue to deliver breakthrough technologies in the decades to come.

1. This article draws upon the author's collaboration over the years with several others, especially William Bonvillian, Seymour Deitchman, Michael Lippitz, Sidney Reed and Patrick Windham, as well as numerous discussions with many DARPA and Department of Defense officials.
2. The key features of the "DARPA model" were laid out in Richard Van Atta, "Fifty Years of Innovation and Discovery," *DARPA: 50 Years of Bridging the Gap*, DARPA 2008, the introduction to DARPA's 50th anniversary publication.
3. ARPA was formed by Secretary McElroy on Feb. 7, 1958, through DOD Directive 5105.15. Congress, through an amendment by Senator Mike Mansfield, renamed "ARPA" as "DARPA" in 1972, adding the word "Defense." Congress, through Senator Jeff Bingaman, renamed it "ARPA" again in February 1993 because of its "dual-use" role in creating technologies with commercial as well as military applications. The name reverted to "DARPA" in March 1996.
4. Richard J. Barber Associates, Inc., *The Advanced Research Projects Agency, 1958-1974*, December 1975. p. VII-3.
5. Richard H. Van Atta and Michael J. Lippitz, with Jasper C. Lupo, Rob Mahoney, and Jack H. Nunn, "Transformation and Transition: DARPA's Role in Fostering an Emerging Revolution in Military Affairs, Volume 1 – Overall Assessment," Alexandria, Virginia: Institute for Defense Analyses, IDA Paper P-3698, April 2003.
6. This transformation is detailed in M. Mitchell Waldrop, *The Dream Machine*, New York: Viking 2001.
7. The agency also has approximately 100 other government employees who provide important services, such as contracting, legal services, human resources, and security, and at any one time it also has several military liaisons. Additionally, contractors support these government employees. Some of these contractors are highly trained PhD scientists and engineers who provide valuable technical assistance to program managers, and others are support staff.
8. There was one significant exception. DARPA did develop operational technology for seismic detection of Soviet underground nuclear tests. DARPA was only able to transition this seismic detection network to the Air Force after running it for approximately 20 years.
9. Care should be taken to differentiate between transitioning such radical, "high-risk, high payoff" technologies and transitioning incremental improvements in existing technologies. DARPA explores many prospective advanced technologies and concepts and often its projects will demonstrate that many of these will not work given current knowledge and state-of-the-art. One of the most important aspects of DARPA management – from the program manager to the office director to the agency's director – is determining whether the concept should be further pursued, and in what way, if a particular program is not bearing fruit. Another major management responsibility is determining how best to move projects into transition and use when they show promising results.
10. This section draws largely from Bonvillian and Van Atta, "ARPA-E and DARPA: Applying the DARPA Model to Energy Innovation," published in 2011.

DARPA TILES TOGETHER A VISION OF MOSAIC WARFARE

Banking on cost-effective complexity to overwhelm adversaries

By Stew Magnuson

Start with a selection of tiles of different colors and shapes and you can assemble a thousand different mosaic images. Mosaic composition has become an important metaphor for Dr. Thomas J. Burns, who with his colleagues at DARPA envisions combatant commanders of the future plucking from a list of manned and unmanned weapon systems and tiling them together into a battle plan.

The concept is called “Mosaic Warfare.” Like the ceramic tiles in mosaics, these individual warfighting platforms are put together to make a larger picture, or in this case, a force package.

The idea will be to send so many weapon and sensor platforms at the enemy that its forces are overwhelmed. The goal is to take complexity and to turn that into an asymmetric advantage, said Burns, who retired this past May as director of DARPA’s Strategic Technology Office (STO), handing the leadership baton off to Timothy Grayson.

“When you attack in parallel across a wide front and you have distributed your sense-decide-and-act systems across a wide number of platforms, you can mass your firepower without having to mass your forces,” said Burns, who is credited with getting talk about the Mosaic Warfare concept going in Department of Defense circles.

The DARPA-hard problem is that today’s weapon systems are not built to function this way, Burns said. “They are more like pieces of a puzzle than tiles for a mosaic. They are exquisitely engineered to fit into a certain part of the picture and one part only. You can’t pull it out and put in a different puzzle piece. It won’t fit,” he said.

One way Mosaic Warfare might work in a ground battle would be to send an unmanned aerial vehicle or ground robot ahead of the main ground battle force. It might spot an enemy tank. The unmanned system passes the coordinates back, which are then relayed to a non-line-of-sight strike system in the rear, which in turn launches its munitions and takes out the target.

“It sounds like it should be something very doable, but it’s not right now,” said Burns. “The interfaces are not made to communicate that kind of information and the Army doesn’t have air and ground vehicles that it can send forward,” he added.

In the air domain, four F-16s might be going head-to-head with four rival jet fighters. However, in a Mosaic Warfare context, the U.S. Air Force might also deploy four relatively inexpensive, somewhat expendable unmanned aerial systems ahead, each with different

weapons or sensor systems. The combatant commander can treat these assets like a football coach who chooses team members and then positions them on the field to run plays. The added aircraft make the situation much more complex and can overwhelm the opponent’s decision-making.

“It makes us more lethal and a lot more survivable,” Burns said. But like a football play, things don’t always unfold as planned. The autonomous systems and pilots must be able to adapt, especially as the mission changes or unexpected events occur. And commanders in a Mosaic Warfare context would have the option of substituting new components and systems as parts of the initial mosaic composition are lost or they want to deploy a new tactic that requires different capabilities.

John Waterston, a Program Manager in the STO and a Navy Reserve officer, said Mosaic Warfare may impose even more complexity on the adversary in the maritime domain, because it encompasses a diversity of environments: air, land, sea, and undersea. His charge now is to figure out how ships, submarines, aircraft, and unmanned systems all can work together to achieve a mission.

The organization and war-planning task almost surely would cross services as well, as combatant commanders mix and match assets. This fits into recent joint multi-domain battlefield concepts that Pentagon leaders have been talking about. These decision-makers acknowledge that going up against peer and near-peer competitors means having to protect forces from threats that could be coming at them from any domain — ground, air, space, sea, and/or cyberspace.

“You want to leverage the best characteristics of different platforms,” Waterston said. “It all gets down to where do you have access and capacity, and distributing them properly so all your eggs aren’t in one basket.”

“We keep making awesome stealth fighters, or better submarines, and better and better unmanned systems,” Waterston elaborated. “The thinking is: Why don’t we take simpler systems and then network them



together, have them share, collaborate – sense their world in their own unique way – and put them together?”

Expendability (attritability in military speak) is key, Burns said. Conventional wisdom says U.S. forces shouldn’t fight in the open. “You’ll be killed. But if you have large numbers of expendable platforms, you can fight in the open,” he said.

Again, the problem is that to create these systems of systems, they need to be linked together, Waterston said, highlighting the challenges here by pointing to recent reports that the new F-35A and the F-22, the Air Force’s two most sophisticated fighters, cannot stealthily share data.

Burns said the Strategic Technology Office’s goal is to create the interfaces, communications links, and the precision navigation and timing software – the technology backbone – to allow these exquisite systems to work together. On PowerPoint illustrations of battlefields, these communication links are often portrayed with lightning bolts. “One of our mottos is to make lightning bolts real,” Burns said.

For a concept that welled up from DARPA, rather than from the services, think tanks, or war colleges, Burns said that Mosaic Warfare – a term coined by Burns and his former deputy director Dan Patt – has been relatively well received during briefs to military leaders.

For Patt, Mosaic Warfare is “about an effective warfighting whole made up of many diverse and fluid pieces. ... How can you get all these little pieces all aligned toward a common objective without perfect communications and without planning everything in advance? It’s really hard. And that’s the idea of Mosaic Warfare.”

There is a direct line in thinking from the Chinese military strategist Sun Tzu and his treatise, *The Art of War*, to the Mosaic Warfare concept, said Patt, who is now CEO of Vecna Robotics and a non-resident Senior

DARPA’s Strategic Technology Office (STO) is focused on technologies that enable fighting as a network to increase military effectiveness, cost leverage, and adaptability. STO’s areas of interest include battle management; command and control; communications and networks; intelligence, surveillance, and reconnaissance; electronic warfare; positioning, navigation, and timing; and foundational strategic technologies and systems.

“We keep making awesome stealth fighters, or better submarines, and better and better unmanned systems,” Waterston elaborated. “The thinking is: Why don’t we take simpler systems and then network them together, have them share, collaborate – sense their world in their own unique way – and put them together?”



Named for the imaginary, mischievous imps that became the good luck charms of many Allied pilots during World War II, DARPA's Gremlins program envisions launching groups of UAVs from existing large aircraft such as bombers or transport aircraft – as well as from fighters and other small, fixed-wing platforms – while those planes are out of range of adversary defenses.



Fellow at the Center for Strategic and Budgetary Analysis, a Washington, D.C.-based think tank.

"All ideas are present in Sun Tzu. But when these ideas are applied for the first time, it can give an asymmetric advantage," he said, citing Germany's blitzkrieg tactics in World War II as an example, where an overwhelming force of armor, motorized infantry, artillery, and air power combined to force a local breakthrough that could then be exploited to continue the advance.

The so-called "Second Offset" strategy following the Vietnam War, which matured into the air-land battle concept, called for airborne sensors and missiles that could work together to overpower a large Soviet army without having to escalate to nuclear warfare. Also known as Assault Breaker, the strategy is centered on the deployment of a system of systems, Patt said.

The problem with that approach was that it was "very brittle," he said. It took years of engineering to ensure one system could link with another system.

COURTESY OF DYNETICS

DARPA IMAGE



M1A1 Abrams main battle tanks and M2/M3 Bradleys of the 3rd Armored Division move out on a mission during Operation Desert Storm. The air-land battle concept was triumphant during Desert Storm, but Mosaic Warfare promises to greatly expand upon the concept, including many more systems.



“They brought together a couple pieces and it did offset Soviet capability,” Patt said. “But it wasn’t a particularly easy or scalable approach. There were a lot of challenges to make that work fluidly. Today, it is still difficult for systems to share information with each other.”

“There has to be a better way. And the technologies DARPA is developing are about that better way,” Patt said.

Another benefit of Mosaic Warfare is that it makes the kill chain more resilient, Patt said. The sense-decide-act decision loop has also been around since the days of Sun Tzu, or longer, he said. More recently, the U.S. military refined the idea to the observe-orient-decide-act decision cycle, or OODA loop.

If a commander could unbundle those functions, everything that has a sensor could be connected to everything that can make a decision, and then to anything that can take an action. “That is really powerful, because mathematically, you expose all the possible combinations and create thousands upon thousands of connections,” Patt said. Those thousands upon thousands of connections force an enemy to contend with many possible combinations of attacks as well.

“That gives resilience. It doesn’t matter what the enemy does, the [blue force] still has options for completing a kill chain.”

Many of the platforms that could be used for Mosaic Warfare already exist. Nevertheless, work continues on developing unmanned platforms that could be applied to the concept.

As for expendable platforms, that falls to DARPA’s Tactical Technology Office (TTO), which is working hand in hand with STO to push the Mosaic Warfare technology forward, Burns said.

The TTO robotics program, Gremlins, is developing some of the hardware needed for Mosaic Warfare, Patt said. The program is named after the imaginary flying imps that were blamed for equipment malfunctions or mishaps by Allied aviators during World War II. The heart of the program is groups of small unmanned aerial systems that can be launched from “motherships” such as bombers or transport aircraft – as well as from fighters and other small, fixed-wing platforms – while those planes are out of range of adversary defenses.

When the gremlins complete their mission, a mothership would retrieve them in the air and carry them home, where ground crews would prepare them for their next use within 24 hours.

“The gremlins’ expected lifetime of about 20 uses could provide significant cost advantages over non-expendable systems by reducing payload and airframe costs and by having lower mission and maintenance costs than conventional platforms, which are designed to operate for decades,” according to a statement released by DARPA.

Waterston said the next step for Mosaic Warfare, and an all-important one, will be demonstrating how it all works.

“Operational commanders aren’t going to use these systems if they haven’t been tested and demonstrated,” Waterston said. “They have to trust them.”

Many of the platforms that could be used for Mosaic Warfare already exist. Nevertheless, work continues on developing unmanned platforms that could be applied to the concept.

FURTHERMOORE

The intertwined history of DARPA and Moore's Law

By William Chappell

In 1965, the legendary technology pioneer Gordon Moore set us on a 50-year odyssey so consequential that it is defensible to think of our times as the “Microelectronics Age.” In a short paper published that year in *Electronics* magazine titled “Cramming more components onto integrated circuits,” Moore predicted a trajectory of progress in which the transistor count of integrated circuits would roughly double every two years while the cost per transistor would decrease. When the three-page paper was first published in this niche trade magazine, early readers couldn’t possibly have imagined the impact it would have on the electronics industry. However, from these humble beginnings emerged the line of progress that we today know as Moore’s Law.

Moore did not likely predict that he would set a course for the investment of hundreds of millions of federal research dollars, and even more from industry. He, along with many leaders in the field from government and industry, predicted this line of progress’ demise. However, through ingenuity, funding, and partnerships, the prophecy continues to be fulfilled.

DARPA, like so many other research institutions, has turned to Moore’s Law as a means of charting a continued path forward in electronics innovation. For decades, the agency has invested heavily in the advancement of electronics, yielding many industry-changing technologies, while fulfilling Moore’s prophecy.

Since its inception, DARPA often has relied on an open research model that involves pairing with non-defense-oriented partners. Rather than relying on secrecy, which often is required in military research, the investments the agency has made in the fundamentals of semiconductors, as a result of this inclusive and collaborative model, have allowed the country to take the lead in pioneering this technology. We have helped build communities that allow ideas to be rigorously developed, and then perfected and manufactured by industry, generating advancements that have brought both economic and defense gains. Correct navigation of Moore’s Law has been a defining factor for our position of global leadership.


One of DARPA’s earliest investments in the advancement of integrated circuit technology was an ambitious effort called the Very Large Scale Integrated Circuits (VLSI) program. During the 1970s and 1980s, VLSI brought together entire research communities to create significant advances in computer architecture and system design, microelectronics fabrication, and the overall cycle of design fabrication, testing, and evaluation. These R&D commitments helped overcome early barriers to the transistor-scaling trends that Moore articulated. The progress achieved under VLSI helped propel the field of computing, furthering U.S. military

capabilities and enhancing national security, all the while helping to usher in a new era of commercial applications¹.

Among the resulting technologies from the VLSI program were Reduced Instruction Set Computing (RISC) processors, which have provided the computational power undergirding everything from supercomputers and the NASA Mars Pathfinder to today’s cellphones and mobile devices². Because of the development of RISC processors, the performance of graphics hardware grew 55 percent per year, essentially achieving a doubling in performance every 18 months³. Although Moore’s observation solely described the inverse relationship between an increasing number of transistors and cost, performance improvements quickly became synonymous with transistor scaling and a prime motivator for continued scaling.

The VLSI program underscored the need for continued collaboration across the U.S. electronics community as well as the role DARPA could play in opening doors for further innovation. To help foster the pursuit of new chip designs, DARPA established the Metal Oxide Silicon Implementation Service (MOSIS) in January 1981. MOSIS provided a fast-turnaround, low-cost capability to fabricate limited batches of custom and semi-custom microelectronic devices. The service opened opportunities to researchers who otherwise would not have had direct access to microelectronics fabrication facilities. Over the course of its more than 35-year run, MOSIS fostered a steady pace of innovation in microelectronics design and manufacturing.

While the U.S. accelerated the pace of microelectronics innovation throughout the 1970s and early 1980s, Japan took the lead in advanced semiconductor production and manufacturing toward the end of the 1980s. To regain dominance, the Semiconductor Manufacturing Technology (SEMATECH) consortium was founded with support and funding from



DARPA and the U.S. semiconductor industry. Throughout the decades that followed, the consortium fostered stronger community engagement among manufacturers and suppliers, and significantly enhanced R&D of next-generation production tools and equipment. By 1992, the United States was responsible for 82 percent of the semiconductor production yields, which is in part attributable to this cross-community effort⁴.

During the late 1980s and early 1990s, new and evolving military and commercial applications, including advanced weapons systems, networks, and the Global Positioning System (GPS), continued to drive the need for powerful, low-cost microelectronics. The persistent transistor scaling needed to make this happen, of course, required that innovation in semiconductor materials, device integration schemes, and other technical areas continue unabated.

During this time, DARPA funded a program that ushered in the state of the art in semiconductor lithography. Working with academia and industry, the program advanced the development of new lens materials and photoresists capable of pushing past technical barriers that had previously limited the technology to 248-nanometer (nm) lithography and of supporting new-generation technology produced with 193-nm lithography. These advances in miniaturization and circuit density had a dramatic effect on the semiconductor industry. The new lithography capabilities quickly became mainstream and industry players used it for advanced commercial and military microelectronics.

Building on the exploration of new materials and integration schemes from the early 1990s, DARPA launched a program to develop transistors beyond the 25-nm-size threshold in 1995. The research efforts completed under the program led to FinFETs (Fin Field Effect Transistors) based on a novel 3-D transistor design that leverages protruding fin-like silicon structures, which allow multiple gates to operate on a single transistor. Today, leading chipmakers continue to use FinFET technology to scale transistors down to 7 nm.

While Moore's predictions helped chart the course for transistor scaling over the past 50 years, it was the ingenuity and dedication of industry, academia, and government organizations, like DARPA, that brought Moore's Law to life. DARPA's investments have helped industry and the Department of Defense (DOD) overcome the barriers of traditional transistor scaling through the

From humble beginnings in a 1965 trade magazine, few would have imagined that Gordon Moore's prediction – that the transistor count of integrated circuits would roughly double every two years while the cost per transistor would decrease – indeed would characterize a 50-year journey into the age of microelectronics.



discovery of new materials that exceed current limitations and can attain future performance and efficiency requirements. This has only been possible by fostering an environment for collaboration and innovation around novel design schemes and architectures, and by opening pathways for experimentation within the manufacturing and production of microelectronics.

It is because of the intertwined history of commercial and defense support for the semiconductor industry through programs like VLSI, MOSIS, and SEMATECH, that the United States has enjoyed the distinct advantage of global leadership in microelectronics innovation. This has resulted in consumer electronics having benefited from components with heritage in the DOD, such as GPS, as well as military systems that are leveraging the processing power of leading-edge commercial processors alongside purpose-built integrated circuits.

Looking Back to Look Forward: Moore's Inflection

The U.S. semiconductor industry plays a uniquely outsized role to the U.S. economy, substantially contributing more than any other major domestic manufacturing sector to the economy⁵. Over the last 30 years, growth in the semiconductor industry has increased rapidly, outpacing the U.S. GDP growth rate by more than a factor of six⁶.

Not all good things can go on forever, however. Today, semiconductor technology continues to progress according to Moore's Law, but that march forward is showing signs of slowing down. In addition to the fundamental technological limits that apply as the size of devices continues to shrink, unintended consequences associated with the economics of continuing down this path are surfacing. Increasing circuit complexity and the associated development costs have kept many commercial and government organizations from participating at the cutting edge of electronics R&D.

Today, U.S. electronics development and manufacturing is facing a trio of challenges that threaten the future health of the industry, as well as our military capabilities:

1. **The cost of integrated circuit design is skyrocketing, which is limiting innovation.** Only large, global, multinational entities backed by massive commercial demand can innovate and compete in today's electronics landscape. This severely limits the complexity of circuits that cash-strapped startups and DOD designers can produce.
2. **Foreign investment is distorting the market and driving a shift outside of the United States.** China's plan to invest \$150 billion into developing its manufacturing capabilities is luring foreign interest. Even by 2015, China already had begun building 26 new 300-mm semiconductor foundries⁷ and had launched 1,300 fabless startups⁸. These global, economic forces are placing a premium on transformative semiconductor invention to stay ahead.
3. **The continued move toward generalization and abstraction is stifling potential gains in hardware.** The rising cost of managing the complexity of a modern electronics system – from manufacturing and designing circuits to programming – has led to increased layers of abstraction. The numerous steps from the invention in the bottom of the stack (for example in new materials) to the money-making portion higher up the computing stack leaves a reluctance to invest significantly. Coupled with the predictable benefits of continued transistor scaling, this has created an ecosystem where only generalized electronic hardware can be economically successful, and much of the value has moved closer to the application higher up the software stack. As a result, hardware has become closer to a commodity, reserving much of the potential gains in performance from specialized hardware for only select situations.

At a time such as this, it is instructive to go back to the origins of the industry and look to the leaders of the field for clues on how to move forward. Even while setting the course in 1965, Moore himself foresaw the end of scaling. In his seminal paper in which he conveyed the projection we know as Moore's Law, Moore predicted that economic limitations, in addition to technical and engineering challenges, could eventually become an impediment for scaling. Equally important, on the third page of his article, he predicted that progress in areas that today we know as design automation, materials science, packaging, and architecture specialization could keep the pathway open for increasingly capable electronics.

Based on his observations 50 years ago, Moore accurately predicted the point we are reaching today. In honor of Moore's ongoing presence in electronics, we at DARPA refer to this point as "Moore's Inflection" – a point where the priorities we set today will determine whether the state of the electronics ecosystem becomes stagnant, rigid, and traditional, or grows to be dynamic, flexible, and innovative.

The Electronics Resurgence Initiative: A Response to Moore's Inflection

As Moore's Inflection approaches, the U.S. government has decided to take large-scale action by investing some \$1.5 billion over the next five

years in the DARPA-led Electronics Resurgence Initiative (ERI). ERI seeks to build a specialized, secure, and heavily automated innovation cycle that will enable the U.S. electronics community to move from an era of generalized hardware to specialized systems.

Building on DARPA's legacy of electronics invention, ERI aims to foster forward-looking collaborations and novel approaches to usher in this new era of circuit specialization. The large-scale initiative will apply DARPA's open research model to the future of microelectronics and bring together government, academia, industry, the defense industrial base, and the DOD to create the environment needed for continued and dramatic advancement.


In deference to the guidance provided on page three of Moore's 1965 paper, ERI seeks to create an ecosystem where smarter design automation tools will be able to directly take logic diagrams and turn them into physical chips without requiring any special engineering intervention in between. This would make it economically feasible to produce small batches of custom circuits – or accelerator cores – designed for specific functions rather than only producing large volumes of general circuits. The ability to construct and interconnect arrays of custom circuits to form larger systems would enable the rapid and highly efficient creation of a considerable variety of unique electronic products.

ERI is comprised of several DARPA programs – many of which kicked off after the official announcement of the initiative in June 2017 – that focus on three primary research thrusts: architectures, materials and integration, and design. Teams in the design thrust seek to develop an open framework that enables researchers and design teams to apply machine-learning algorithms that can rapidly and automatically translate high-level functions and requirements into physical layouts of custom circuits. To ensure that a variety of custom circuits, materials, and device technologies can be used together to build larger systems, the materials and integration thrust will investigate new interconnect standards and the integration of novel memory and logic circuits. Lastly, the architectures thrust will explore circuit-level coordination and hardware/software co-design methodologies to help create modular and flexible systems able to adapt and optimize combinations of new devices and accelerator cores into systems tailored for any application.

Design

"Perhaps newly devised design automation procedures could translate from logic diagram to technological realization without any special engineering." – Gordon Moore, 1965

Although Moore could not have predicted the extent to which his observations on transistor scaling would be stretched, he did understand how the increase in the number of transistors would eventually create circuits too complex for designers to lay out by hand and that automation tools would need to be developed. At the time when Moore published his observations, integrated circuits had around 50 transistors; today, that number is around 21 billion⁹. The electronics community began developing electronic design automation (EDA) tools to help automate the process as the number of transistors continued to increase. As powerful as these tools are at helping designers manage the complexity of laying out billions of transistors, they have not kept pace with physical manufacturing capabilities and the rise of analog circuits, which are still manually



designed. As a result, the size of design teams has exploded, and the need for specialized technical expertise has never been greater.

The development of modular design methodologies has helped mitigate some of the limitations of EDA. One technique designers employ is to capture frequently used circuit functions into discrete, modular blocks, called intellectual property (IP) blocks, which can be used and reused to create larger, more complex systems. For comparison, in 2000, more than 90 percent of a chip had to be specially designed. Today, that number has reversed with designers reusing already-designed IP blocks for more than 90 percent of a chip¹⁰.

Even with the growing use of IP blocks, however, the rapid rise in the cost to design and verify new hardware has made access to leading-edge electronics prohibitively expensive to all but the largest companies. The Circuit Realization at Faster Timescales (CRAFT) program was conceived to explore solutions to this problem through the use of automated generators to rapidly create new circuits and accelerate the design cycle. Recently, researchers in the CRAFT program demonstrated a design flow that leveraged automated generators to produce digital circuits seven times faster than that achieved by traditional methods. Put in another way, these tools enabled small design teams to be just as productive as teams seven times their size.

Maintaining continued forward momentum beyond the imminent Moore's Inflection will require pushing the limits of machine learning to extend automation into every aspect of circuit design. Two new programs in the ERI Design thrust, inspired by Moore's prescience, aim to explore machine-centric hardware design flows that can support the physical layout generation of complex electronic circuits with "no human in the loop" and in less than 24 hours. To facilitate the reliable reuse of circuit blocks and to engage the collective brain power of the open-source design community, these efforts will seek to leverage new simulation

"Moore's Inflection" – a point marked by arrows on the diagram where priorities set today will determine whether advances in electronics will begin to slow and stagnate or where new innovations will catalyze another long run of dynamic and flexible technological progress.

technologies and applied machine learning to verify and emulate circuit blocks. With enhanced design automation tools like these, the barrier to entry for a growing number of innovators will shrink and thereby unleash an era of unprecedented specialization and capability in electronics technologies.

Materials and Integration

"... build large systems out of smaller functions, which are separately packaged and interconnected." – Gordon Moore, 1965

A central challenge to managing modularity is how to properly interconnect the growing number of functional blocks without affecting performance. Since 2000, not only has the number of transistors per chip grown from 42 million to 21 billion¹¹, but the number of IP blocks on that same chip has increased more than 10 times¹² as well. In addition, these functional blocks are increasingly becoming a mixture of digital and analog circuits, and are often made from vastly different materials, further complicating the challenge of integration.

To realize Moore's vision of building larger functions out of smaller functional blocks, we need to find new ways for various dissimilar blocks to connect and communicate with each other.

Moore's predictions regarding materials and integration are already being realized in DARPA's Common Heterogeneous Integration and IP Reuse Strategies (CHIPS) program. This research effort seeks to develop modular chip designs that can be rapidly assembled and reconfigured through the integration of a variety of IP blocks in the form of prefabricated chipleths. These chipleths will leverage standard layouts and interfaces to easily link together. The program recently announced that Intel® will be contributing their proprietary interface and the relevant IP associated with it to the CHIPS program to be used as the program's standard interface. Intel's direct participation will help ensure that all the various IP blocks in the program can be seamlessly connected together. This is a huge step toward the creation of a national interconnect standard that will enable the rapid assembly of large modular systems.

What is often left out of the story behind the growing number of transistors is the parallel rise of the number of interconnects required to shuttle data back and forth across the chip. The explosion of wires have not only complicated the design process but have also created longer and more convoluted paths for data to travel through. To get a sense of scale, if all the wires in a modern chip were laid out end to end, they would span more than 21 miles. For most computing architectures, which separate the central processing unit (CPU) and the memory, moving data across this growing tangle of wires severely limits computational performance. The conundrum even has its own name, the "memory bottleneck." For instance, to execute a machine-learning algorithm on a leading-edge chip, more than 92 percent of the execution time is spent waiting to access memory.

With the vast number of circuit combinations made possible by new standard interfaces and the performance limitations of current interconnects, we must ask the question: What role can new materials and radically new architectures play in addressing these challenges? In response to this question, one of the new programs under the ERI Materials and Integration thrust plans to explore the use of vertical, rather than planar, integration of microsystem components. By leveraging the third dimension, designers can dramatically reduce wire lengths between essential components such as the CPU and memory. Simulations show that 3-D chips fabricated using older 90-nm nodes can perform 50 times better than circuits fabricated on 7-nm nodes using planar integration. Furthermore, another program will investigate new materials coupled with architectures that rethink the flow of data between processors and memory to provide new solutions for processing the growing volume of scientific, sensor, social, environmental, and many other kinds of data.

Architectures

"The availability of large functions, combined with functional design and construction, should allow the manufacturer of large systems to design and construct a considerable variety of equipment both rapidly and economically." —Gordon Moore, 1965

The relentless pace of Moore's Law ensured that the general-purpose computer would be the dominant architecture for the last 50 years. When compared to performance gains achieved under

Moore's Law, exploring new computer architectures and committing the years of development and hundreds of millions of dollars required to do so just did not make economic sense. As this trend starts to slow down, however, it is becoming harder to squeeze performance out of generalized hardware, setting the stage for a resurgence in specialized architectures.

Imagining what the future would look like, Moore suggested a framework for delivering specialized architectures by focusing on "functional design and construction" that would lead to manufacturable systems that also make economic sense. In other words, he was envisioning flexible architectures that can take advantage of specialized hardware to solve specific computing problems faster and more efficiently.


Last year, DARPA started the Hierarchical Identify Verify Exploit (HIVE) program to explore the optimization of a specialized integrated circuit that could analyze the various relationships between data points in large-scale datasets, such as social media, sensor feeds, and scientific studies. Working with industry partners such as Qualcomm and Intel, the HIVE program aims to develop a specialized integrated circuit capable of processing large-scale data analytics 1,000 times faster than current processing technology. This advanced hardware could have the power to analyze the billion- and trillion-edge datasets that will be generated by the Internet of Things, ever-expanding social networks, and future sensor systems.

While HIVE is an example of current progress, it will take much more innovation to bring Moore's vision of specialized hardware to fruition. One of the key challenges to employing more specialization is the tension between the flexibility of general-purpose processors and the efficiency of specialized processors. If designers find specialized hardware too difficult to use or program, they are likely to forgo the efficiencies the hardware could deliver.

The two new ERI Architectures programs seek to demonstrate that the trade-off between flexibility and efficiency need not be binary. These programs seek to develop methods for determining the right amount and type of specialization while making a system as programmable and flexible as possible.

One of the programs will investigate reconfigurable computing architectures and software environments that together enable data-intensive application performance near that of single application specialized processing implementations without sacrificing versatility or programmability. The resulting capabilities will enable the real-time optimization of computational resources based on real-time introspection of incoming data. The program will realize processing performance 500-1,000 times better than state-of-the-art, general-purpose processing and provide application-specialized processing performance while maintaining flexibility and programmability.

The second program under the Architectures pillar of ERI will explore methods for combining a massive number of accelerator cores. Although accelerator cores can perform specific functions faster and more efficiently than is possible in software running on a general-purpose processor, programming and coordinating applications on many heterogeneous cores has been a big challenge. One solution is to take a vertical view of the computing stack, which cross-cuts from the application software to the operating system and all the way down to



the underlying hardware. By exploring the concept of a domain-driven approach to identify the appropriate accelerators; working on better languages and compilers to optimize code for these accelerators; and implementing intelligent scheduling for the applications running on such a complex processor, this program is looking at a new concept in customized chips that can rapidly utilize myriad accelerators to address multiple applications.

Toward a More Innovative Future

The gains that came from Moore's Law were not guaranteed, but realized through ingenuity and close collaboration between commercial industry, academia, and government. Today, the rising cost to design integrated circuits, increasing foreign investments, and the commodification of hardware threaten the future health of an innovative and dynamic domestic microelectronics community. Facing these challenges, the Electronics Resurgence Initiative will build on the long tradition of successful government-industry partnerships to foster the environment needed for the next wave of U.S. semiconductor innovation.

The CHIPS program is pushing for a new microsystem architecture based on the mixing and matching of small, single-function chiplets into chip-sized systems as capable as an entire printed circuit board's worth of chips and components.

1. National Research Council. 1999. Funding a Revolution: Government Support for Computing Research. Washington, DC: The National Academies Press. <https://doi.org/10.17226/6323>.
2. National Research Council. 1999. Funding a Revolution: Government Support for Computing Research. Washington, DC: The National Academies Press. <https://doi.org/10.17226/6323>.
3. IBM. RISC Architecture. <http://www-03.ibm.com/ibm/history/ibm100/us/en/icons/risc/>
4. National Research Council. 1999. Funding a Revolution: Government Support for Computing Research. Washington, DC: The National Academies Press. <https://doi.org/10.17226/6323>.
5. SIA, "The U.S. Semiconductor Industry: A Key Contributor to U.S. Economic Growth," 2014.
6. McKinsey industry analysis
7. EE Times, "China Fab Boom Fuels Equipment Spending Revival," March, 2017. https://www.eetimes.com/document.asp?doc_id=1331492
8. EE Times, "Much Ado About China's Big IC Surge," June, 2017. https://www.eetimes.com/document.asp?doc_id=1331928
9. PC World, "Nvidia's monstrous Volta GPU appears, packed with 21 billion transistors and 5,120 cores," May, 2017, <https://www.pcworld.com/article/3196026/components-graphics/nvidias-monstrous-volta-gpu-appears-packed-with-21-billion-transistors-and-5120-cores.html>
10. SEMICO Research Corporation, 2014
11. The Economist, "Technology Quarterly: After Moore's Law," <https://www.economist.com/technology-quarterly/2016-03-12/after-moores-law>
12. SEMICO Research Corporation, 2014

THE STUFF OF NATIONAL SECURITY

DARPA and the invention of new materials

By Ivan Amato and Steve Wax

About a decade ago, when DARPA Program Manager Dr. Jan Vandenbrande was a technical fellow and senior manager of the Applied Math, Geometry, and Optimization group at Boeing, he visited an assembly line where he could examine two of his company's iconic aircraft – the aluminum-body 777 and the 787, half of which is made of carbon-fiber-reinforced composite.

“What was absolutely remarkable to me was that the two structures were almost identical,” Vandenbrande said, noting that more than a decade passed between the 777's maiden flight in 1994 and the 787's maiden flight in 2009. “We had gone from metals to a dramatically different material, a composite, yet we had not changed how we design and built the planes. Things were still being riveted together and we still shaped the two planes pretty much the same way.”

While Vandenbrande's comments highlight the current challenges in material science and design from the vantage point of 2018, it is easy to forget that the path to even considering load-bearing composite aircraft structures was long and arduous. It is equally easy to take for granted all of the capabilities made possible by advances in materials science over the last 50 years. Lightweight aircraft and ship structures, high-temperature jet engines, satellite mirrors and optics, and even the ceramic body armor that has saved so many of our warfighters' lives would not exist were it not for these advances. Composites, high-strength plastics, and lightweight metals have likewise revolutionized the commercial world.


More than 60 years ago, during the heart of the Cold War, such capabilities did not exist. Many in the technology community viewed materials as “the single most limiting factor on progress in military, space and nuclear system development,” according to a detailed study by Richard J. Barber Associates of DARPA's early years. However, it was not until the job of overcoming those limitations was given to a fledgling agency with a charter for innovation that things really got going. And so began a nearly 60-year relationship between DARPA and the science community.

And that relationship began in a disciplinary scrum of sorts. On July 11, 1960, just a little more than two years after it was chartered, DARPA (then ARPA) announced the creation of three interdisciplinary laboratories (IDLs) at Cornell University, the University of Pennsylvania, and Northwestern University. Over the next several years, the IDL roster would increase to a dozen universities and serve as a model for the establishment of materials programs all over the country. DARPA provided the IDLs with stable funding for research and equipment and a tolerance of – actually a demand for – an interdisciplinary approach to developing materials. With the formation of the IDLs, DARPA is widely recognized for its catalytic role in the creation of the field that became known as materials science and engineering. While DARPA handed the IDL program off to the National Science Foundation (NSF) in 1972, which changed the program's name to the Materials Research Laboratories (MRL) program, this was only the beginning of the DARPA material science story.

From the 1960s through the 1970s, DARPA led the charge of developing materials capabilities for defense, most often with significant support from the Department of Defense's service laboratories – including the Naval Research Laboratory, and what are now the Army Research Laboratory and the Air Force Research Laboratory. Among the long list of materials that were becoming available for defense systems were high-temperature, nickel-based superalloys for high-performing jet engines, ceramic tank armor to combat advances in Soviet anti-armor weapons, and large, precise, stable, low-weight and low-scatter beryllium mirrors for space and missile guidance. Also included in these materials were composites of all ilks – metal matrix, ceramic matrix, carbon-carbon, and early versions of those same polymer composites in Vandenbrande's 787.



Boeing's 787 on the assembly line. In the 15 years that passed between the earlier 777's maiden flight and the 787's first flight in 2009, composites became a major part of airframes, making aircraft lighter, stronger, and more efficient.




As the services increased their role in materials development, DARPA turned to a problem that emerged as a direct result of those advancing materials capabilities. The specific properties, strengths, and weaknesses of a material are derived from both its composition and the way that composition evolves into the material's underlying microstructures during a variety of process treatments such as heating, cooling, pressing, or extruding. Processing these new, complex materials in a way that meets performance specifications at high yield with predictable properties proved to be a significant challenge.

To attack this problem, DARPA convened an extensive set of discussions and brainstorming sessions involving DARPA program managers, professionals in the manufacturing trenches, and advisers on the Defense Sciences Research Council (DSRC) to determine how to advance the field of materials processing for these emerging complex materials. The result was

the Intelligent Processing of Materials (IPM) program, that in the words of Dr. Haydn Wadley, Edgar Starke Professor of Materials Science and Engineering at the University of Virginia, “was sufficiently general that it could — and eventually was — applied to almost every materials process, yet sufficiently precise that it could reduce the variability of each of these processes.”

The concept of IPM is to use *in-situ* sensors to monitor a material manufacturing process as it unfolds, and to compare what you are measuring at each step with what the process models are predicting the state of the material to be. If there is a variation between what is measured and what is predicted, then engineers can use the process models to determine what processing parameters they need to adjust to get the process under control. IPM was a perfect fit for DARPA, as it was enabled by a confluence of game-changing capabilities in microelectronics, computer



DARPA's Defense Sciences Office (DSO) is aggressively pursuing the development of novel materials with the potential to boost national security. Shown here are images depicting (left to right) a cold-plasma materials deposition process, an ultra-low-density structural material, and precision molecular modeling for designing new materials.

twin” of a manufactured material, that is, a digital record of the specific evolving environment that brought a batch of material into existence. This is just the sort of data that researchers subsequently could summon to address weaknesses in their materials models and in the expert systems and other artificial intelligence (AI) tools that the IPM adventure revealed. IPM changed how materials were made.

As IPM took hold in the materials and manufacturing communities in the 1990s, it presaged how the coming Information Age would continue to drive the evolution of how materials are invented, made, and deployed. Emblematic of this phase was the emergence of solid free form (SFF) manufacturing, now known more familiarly as 3-D printing or additive manufacturing. This opened the way to forming material structures layer-by-layer and even point-by-point straight from a digital file.

“The focus was on ceramic components, which were notoriously hard to prototype because they required expensive molds or machining, typically with diamond abrasives,” explained Dr. William Coblenz, a ceramic scientist and former program manager. “Only designers who absolutely needed the properties the ceramic provided, or those with negative risk aversity, would design with ceramics.” As a result, many of ceramics’ fantastic properties, among them high-heat operation and toughness, were lost on designers, who were more apt to rely on materials that were easier to work with.

Even with the ascent of IPM, SFF manufacturing, and other data-reliant advances, there was so much art, artisan knowledge, and empirical trial and error in the mix that advances in materials still remained a slow-go exercise. At the same time, designers of engines, aircraft, and other systems were leveraging new computational tools to speed up their design cycle. A new engine design could take a few years, while it could take a decade or two to develop a new material to the point where engineers and designers actually would embrace it. This temporal mismatch meant that designers, unable to wait out the time it took to certify a potentially more capable new material, were constrained by the materials that were already available.


Taking on the challenge, DARPA began to look into what it would take to shorten the time for insertion of a new material into a DOD application, allowing designers to be less constrained in their selection of new materials. As with IPM, DARPA enlisted the entire materials research community, including the DSRC, and from that the Accelerated Insertion of Materials (AIM) program was born. The concept of AIM was to integrate system design into the development of a new material right from the start.

“In AIM, we turned the question around and asked: What does the designer need to know to use a new material?” said Dr. Leo Christodoulou, who was the program manager for AIM, a former director of the Agency’s Defense Sciences Office, and who is now Enterprise Technology Strategy Leader at Boeing. It was an approach designed to mitigate the impasse caused by the lack of the performance data that a designer needed to greenlight the use of a material. “The aim of AIM was to reduce the time it takes to introduce a new material into service,” said Christodoulou. “It was all about managing the uncertainty in the

science, and a new concept, artificial intelligence — all technologies championed by DARPA.

The vision of IPM could not have been more alluring to materials scientists, who at that time still generally relied on measurements made after the process was done to determine if a material they were trying to deliver was likely to meet specifications for a given application. “The idea was not just to measure the heating cycle, pressure course, chemical environment, and other relevant variables in a materials manufacturing process, but to also sense the attributes of the material that dictated its performance for its intended application,” Wadley explained.

“[IPM] has had a transformational impact upon the processing of materials. It rapidly spread and is pervasively used everywhere materials are processed — from integrated steel mills to microelectronic foundries,” he added. What’s more, Wadley noted that the practice of IPM yields a “digital



The Rapid Low Cost Additive Manufacturing (RLCAM) effort, part of DARPA's Open Manufacturing program, aims to use first-principles and physics-based modeling to predict materials performance for direct metal laser sintering (DMLS), using a nickel-based super alloy powder. In DMLS, a laser melts the metal powder to additively build a 3-D product.


use of a new material. Nobody wanted to be the first to use a material, because something might go wrong."

This new approach called for relying more on powerful assets of data, computation, and modeling and far less on the sluggish, trial-and-error process of the past. The idea behind AIM was to make the distinction between computation and experiment go away so that computational and lab experiments become interchangeable data sources.

Along with work at the Department of Energy and the National Science Foundation, the AIM program led to a sea change in materials development and ultimately to the discipline of Integrated Computational Materials Engineering (ICME). "The DARPA AIM program

on the acceleration of materials technology is probably the single research program with the most impact on the field," according to Dr. Greg Olson, Walter P. Murphy Professor of Materials Science and Engineering at Northwestern University and founder of a company based on computational materials science. "I see computational materials engineering as the greatest innovation in materials since the iron swords of the Hittites," added Olson, referring to the anomalously early accomplishment of the Hittite culture to smelt iron some 3,800 years ago — while the world was still in the Bronze Age.

Follow-on work beyond the AIM program, supported by both DARPA and the Navy, helped to establish 3-D characterization and simulation techniques for materials structure. The computational and simulation tools that emerged from these efforts made for a natural fit with DARPA's championing of Solid Freeform Fabrication in the 1990s, a precursor of today's rapidly advancing infrastructure of additive manufacturing and 3-D printing. "The adaptation of the AIM acceleration approach to 3-D printing technology to rapidly qualify 3-D printing technology is a tremendous advance," Olson said, referring to the parallel need



DARPA's Atoms to Product (A2P) program is developing means such as microelectromechanical (MEM)-based manipulation for assembling nanometer- to micron-scale components into larger human-scale systems.

to certify the usability of 3-D parts in actual systems along with the ability to make the parts in the first place.

With the growing ability to use 3-D printing to combine form and functions in ways that had not been possible before, said Vandenbrande, "we have opened up the design space significantly." This ushers in another set of challenges, of course. 3-D printing provides the possibility of engineering the properties of each volumetric pixel, or voxel, of a part. The amount of data and computation required to do that is enormous. "What we need is a new set of maths and algorithms that can describe materials and the shapes they comprise in one kind of cohesive way," said Vandenbrande. Coming, he added, is the ability to design multiple functions – think here of aircraft skins that combine structural, sensing, cloaking, and antenna functions – in a fully integrated way.

If materials engineers and manufacturers could master this exquisite degree of control over material structure, they could responsibly imagine such things as strong and featherweight aircraft made out of structural materials designed and grown to have massive

amounts of empty interior space like the bones of birds. The current DARPA Atoms to Product (A2P) program, which is all about developing means for assembling nanometer- to micron-scale components into larger human-scale systems, heads in these directions.

Which brings us full circle to Vandenbrande's observations over a decade ago – and highlights the next great challenge that DARPA is taking on. Vandenbrande refers to this as solving "the inverse problem in design." This is where the designers would first specify the performance values for a system and its parts – whether it is a missile, ground vehicle, or prosthetic limb – and hand those off to computational and modeling tools to generate multiple solutions that balance shape with the detailed material structures that could deliver the performance and functionality required by the design specifications. His current programs, Fundamental Design (FUN Design) and Transformative Design (TRADES), are developing the foundational mathematics and algorithms to initiate this vision.

Vandenbrande envisions a time when computers become a true partner in design, capable of integrating advanced and emerging material architectures, to create systems currently unimaginable today – perhaps an aircraft additively built up of ultra-light micro-structures that resemble bird bones. When there is seamless integration of materials, design, and manufacturing, Vandenbrande will know that his vision has become reality.

GIANT STEPS

DARPA's X-planes and the quest to redefine the boundaries of flight

By Jan Tegler

The story of DARPA's work on experimental aircraft programs began quietly – with a strange airplane that, ironically, wasn't designated an "X" plane. The QT-2 was the answer to a noisy problem.

American and South Vietnamese troops in Vietnam fought primarily during daylight, while the opposition preferred to operate mostly under cover of darkness. To counter enemy activity, U.S. forces conducted nighttime reconnaissance and surveillance by air. But too often, flights after dark yielded little information. The engine noise of American aircraft gave the enemy advanced warning and they simply vanished into the night.

To solve the problem, DARPA (then ARPA) provided funding to Lockheed Martin in April 1967 to develop an airplane that would be nearly silent. Lockheed came up with the "Quiet Thruster," a small, odd-looking propeller-driven aircraft based on the Schweizer SGS 2-32 sailplanes (X-26As) that the U.S. Navy used to train novice pilots.

The program yielded a prototype within six months. An engine mounted on the fuselage behind the cockpit drove a 10-foot shaft connecting to a four-bladed propeller on a pylon at the nose of the aircraft. Employing a heavily muffled engine exhaust and a very slow-turning propeller, the QT-2 was virtually inaudible flying at 1,000 feet – the world's quietest airplane at the time. The follow-on YO-3A placed the engine in a more conventional location at the nose, and the U.S. Army ordered 14, 13 of which entered service. Humble as the little plane was, it might be considered an early venture into the development of stealth aircraft.

The leap in capability produced by DARPA's first X-plane program is emblematic of the agency's approach to X-plane development.

"If you're going to make an improvement, you might as well take a giant step," James Allburn said. "If it works, it works. If it doesn't, then you've shown it was too big a step."

The former X-29 program manager explained that pushing boundaries is part of DARPA's DNA. Some X-plane programs, like QT-2, have been striking successes, yielding new technologies and even aircraft that have directly influenced military aviation.

Others were less successful, but "it was inculcated in me at DARPA," Allburn said, "that we're here to make mistakes, to have failures. If we have a failure, that's fine. Later on, it might be a success."

X-plane Success

Like the QT-2, some of DARPA's most successful X-plane programs made little noise, at least initially.

"Have Blue" was a proof-of-concept program launched in the early 1970s to achieve radar cross-section (RCS) reduction – an effort aimed at making it difficult for the advanced air-defense missile systems proliferating during the period to detect tactical aircraft. By 1977, the world's first practical combat stealth aircraft flew as a DARPA X-plane. DARPA research supporting Have Blue enabled not only RCS reduction through "faceted" stealth shaping, but the development of radar-absorbent materials (RAM), visual-signature-reduction technologies, and jet-exhaust shaping and cooling as well as other infrared-shielding technologies. These DARPA efforts produced an aircraft that revolutionized tactical air power more than 40 years ago. Procured in secret by the U.S. Air Force, its distinctive angular dart shape would be developed into the famous F-117 Nighthawk stealth aircraft, which flew its first combat mission in 1989 and was one of the transformational technologies employed during Operation Desert Storm.

During the same period, DARPA was secretly at work on second-generation radar cross-section reduction technology with "Tacit Blue," employing curved surfaces and advanced materials. The resulting X-plane, nicknamed "the whale" and "Shamu" by program insiders, flew in 1982, and could operate radar sensors while maintaining a low RCS. The aircraft laid foundations for development of the B-2 stealth bomber that debuted in 1989. This second-generation stealth aircraft became operational in 1997 and flew its first combat mission in 1999.

DARPA's development of stealth revolutionized air warfare, and it is important to remember not only the initial shock the technologies caused, but that potential adversaries had no answer to stealth for more than three decades. The Su-57, Russia's first stealth combat aircraft, first flew in 2010, and delivery to the Russian air force is not expected until 2019. China's J-20 stealth aircraft likewise first flew in 2011, and entered squadron service just this year. That these nations took so long to develop

The F-117A Nighthawk is the world's first tactical fighter aircraft designed to exploit low-observable stealth technology, one of DARPA's premier disruptive technologies.





their own stealth aircraft is a telling indication of just how giant a step the development of stealth represented.

Rotary Wings and UAVs

Less well known, but equally successful, was DARPA's work on no-tail-rotor (NOTAR) helicopter technology. The agency's support helped Hughes Helicopters mature the NOTAR concept and show its operational advantages with the NOTAR flying demonstrator in 1981. This led to a NOTAR series of military/commercial helicopters from McDonnell Douglas after it acquired Hughes Helicopters in 1984.

In 1984, DARPA began a rapid prototype program that produced the world's first long-endurance unmanned aerial vehicle, the Amber UAV, which

CLOCKWISE FROM TOP: Tacit Blue employed second-generation stealth shaping and materials to enable a battlefield reconnaissance and surveillance aircraft to operate near the forward edge of the battlefield with a high degree of survivability; DARPA's Amber UAV was the world's first long-endurance UAV, and led to the famed Predator UAV; the RQ-3 Darkstar was developed to a DARPA requirement for a low-observable high-altitude long-endurance UAV, and contributed to later stealthy UAV designs; the beginning of it all was Have Blue, the prototype for the first stealth aircraft. Note the early inward-canted twin tails, later modified to outward-canted tails on the F-117.

in 1988 flew for 38 hours continuously. The technology pioneered with Amber led to the famed medium-endurance MQ-1 Predator UAV.

A decade later, two more unmanned X-plane programs were up and running with DARPA support. Lockheed Martin's RQ-3 Dark Star was developed for DARPA's requirement for a low-observable high-altitude long-endurance (HALE) UAV. A highly survivable intelligence, surveillance, and reconnaissance (ISR) platform, Dark Star's technology contributed to the design of the stealthy Lockheed Martin RQ-170 Sentinel, famously dubbed the "Beast of Kandahar."

The RQ-4 Global Hawk also originated as a DARPA high-altitude long-endurance program in 1994. The agency issued an Advanced Concept Technology Demonstration contract to Teledyne Ryan Aeronautical (later acquired by Northrop Grumman) in 1995, which produced several RQ-4 prototypes. The system transitioned directly to operations with the USAF. Air Vehicle Number 3 from the demonstration program flew thousands of hours of combat missions over Iraq. The demonstrator flew just 3 percent of the imagery intelligence missions over Iraq, but located 55 percent of the time-sensitive targets generated to destroy air defense equipment.

DOD PHOTO

DARPA PHOTO

DARPA PHOTO

NASA PHOTO BY CARLA THOMAS



In 2005, DARPA sponsored development of another helicopter-based X-plane, the unmanned A-160 long-endurance, high-speed helicopter. In 2008, the A-160, intended for ISR missions, set a world record in its weight class, flying for 18.7 hours.

While some of the most famous DARPA X-plane programs were developed behind a veil of secrecy, three of DARPA's most successful aircraft became highly visible examples of successful X-plane programs.

X-29

The X-29 program, featuring a spectacularly odd-looking fighter-type aircraft with forward-swept wings, created quite a stir in the 1980s. While Have Blue and Tacit Blue were conducted behind closed doors, the X-29 flew DARPA's X-plane flag publicly as the most significant experimental aircraft in nearly 15 years.

The effort derived from research on aircraft with forward-swept wings dating back to 1931. As aircraft performance began to rise toward transonic speeds in the World War II era, it was widely recognized that wing sweep could delay a rise in drag that occurred as airplanes approached Mach 1. Aircraft designers knew even then that the benefits of wing-sweep were present regardless of which direction wings were swept – backward or forward.

A high-altitude, long-endurance program run by DARPA led to the RQ-4 Global Hawk. A prototype went directly into operations, flying thousands of hours of combat missions over Iraq.

But forward-swept wings were problematic because their leading edges tended to twist upward when loaded aerodynamically, increasing their angle of attack, which then caused even more twisting. This “structural divergence” made forward-swept wings impractical. Adding stiffness to limit twisting wasn’t possible as the materials then available made the wings too heavy.

Fast forward to 1975, when lighter, stronger materials were available. That year, Lt. Col. Norris Krone (USAF) authored a doctoral dissertation on the use of advanced composites to solve the structural divergence problem. One of the earliest methods of fabricating composite structures in some ways mimicked the construction of plywood. Layers, or plies, of composite fabric were “laid-up” to form a stack. The structure was then usually baked in an autoclave to cure resins permeating the layers.

Krone proposed that the “lay-up” orientation of composite plies and their thickness could be tailored to resist wing twist. The concept was labeled aeroelastic tailoring, and formed the basis for DARPA's research



Krone proposed that the “lay-up” orientation of composite plies and their thickness could be tailored to resist wing twist. The concept was labeled aeroelastic tailoring, and formed the basis for DARPA’s research on the potential of forward-swept wings.

The X-29’s forward-swept wings might have been the aircraft’s claim to fame, but it also pioneered investigations into high angle of attack and post-stall supermaneuverability.

on the potential of forward-swept wings. In 1977, the agency and the Air Force Flight Dynamics Laboratory at Wright-Patterson AFB, Ohio, issued proposals for a research aircraft.

James Allburn was in the midst of a successful career as a fighter pilot, flying F-4 Phantoms at Nellis AFB with the 429th TFS in the late-1970s, when he was introduced to DARPA by a friend who had, like Allburn, been an instructor at the Air Force Academy.

“I was intrigued, interviewed with DARPA, and was selected, but before the orders could be written I got orders to go to Korea!” said Allburn.

The situation was resolved when Dr. William Perry (Undersecretary of Defense for Research and Engineering in the Carter administration) was informed that the Nellis-based major wouldn’t be able to join DARPA because of the deployment.

“Perry decided that I would be better at DARPA than in Korea,” Allburn said.

In the spring of 1980, he was assigned to DARPA’s Tactical Technology Office. Allburn was already familiar with the research activities the X-29 program would undertake via colleagues like Rockwell International engineer Michael Robinson.

“I thought X-29 was a fascinating concept from the standpoint of super-maneuverability, high angle of attack, and post-stall activities,” Allburn said.

Krone served as the initial program manager for the X-29, and was in place when Allburn joined the program. By this time, DARPA had requested information from contractors who could build an aircraft capable of supersonic performance with forward-swept wings. Three firms responded – Grumman, General Dynamics, and Rockwell International.

The contractors’ proposals arrived at about the same time Allburn did. Shortly thereafter, Krone left DARPA and Allburn became DARPA’s X-29 program manager.

“Norris said, ‘Shake the stick. You’ve got it!’” Allburn remembered.

“So with no background in acquisition, I scratched my head to figure out what to do. I got a couple people to help me and went through the proposal

evaluation process. It only took a couple months. Dr. [Robert] Cooper (DARPA Director 1981-1985) made the selection in terms of what airplane to build.

General Dynamics put forward a swept-wing F-16 design, which was “the least risky” of the three proposals but also offered “the lowest degree of wing-sweep,” Allburn explained.

Rockwell “designed a whole new airplane with a highly-swept wing” and systems/sub-systems donated by other manufacturers. Allburn considered the Rockwell “Sabrebat” design “the highest risk.”

Grumman offered a medium-risk design incorporating “an F-5 forward fuselage, along with bits and parts of other airplanes – F-16 actuators and nose wheel, and flight computers from the SR-71.”

“I thought the most interesting aspects of the design were the canards, the flaps and the beaver-tail, and integrating them to reduce the risk overall of transonic drag. They really focused on technologies that could pay off,” he said.

Grumman’s design was chosen, and made its maiden flight on Dec. 14, 1984, from Edwards AFB. So began a nearly eight-year test program during which two X-29s conducted flight test research in three phases, completing nearly 280 flights – an unprecedented number for an X-plane program.

The first phase was dedicated to investigating divergence at transonic speed. In December 1985, the X-29 became the first forward-swept wing aircraft to fly at supersonic speed in level flight, proving that aeroelastic tailoring worked. As the program progressed, the initial airframe was complemented by a second X-29, which first flew in May 1989.

This second aircraft conducted high angle-of-attack testing and a vortex control study that kept it flying until 1992. Allburn said the success of the X-29 program was made possible by its people – a team including members from the Air Force, NASA, and Grumman “who were unbelievable in terms of their support for the program.” Allburn left DARPA in 1986, summing up his involvement with the X-29 as an “exciting” experience.

“People wanted to be part of it at Dryden,” he said. “And the X-29 was the subject of more technical papers by one of the Air Force Academy classes than anything else.”

Visible in magazines, on television, and on static display at air shows (Dayton, Oshkosh), part of the X-29’s contribution was that it re-familiarized the public with X-planes. Technically, the forward-swept wing fighter proved



ABOVE: The X-31 flight testing included an exploration of extremely short takeoff and landing (ESTOL) as part of the U.S. Navy Vector program. The X-31 was equipped with three paddle-like tail vanes that, when activated, redirected the engine's exhaust plume in response to pilot input. This "thrust vectoring" provided control and lift at dramatically low speeds and high angles. RIGHT: In addition to proving thrust-vectoring supermaneuverability and "pointability" in a variety of post-stall regimes, the X-31 could fly at up to 70 degrees angle of attack (AoA), demonstrated extremely short take-off and landing capabilities, and employed helmet-mounted displays for pilots.

the value of composite construction and an advanced digital/analog flight control system.

Allburn had a more modest assessment: "Well, first — and we had confidence this could be done — the wing didn't break off! Second, the computers were working and controlling the airplane so it didn't go end over end. I guess it was the validation of some of the design activities."

X-31

Dr. Michael Francis, the fourth program manager of the attention-getting X-31 Joint Enhanced Fighter Maneuverability (EFM) program, described the history-making performance of the X-31 at the Paris Air Show in 1995 as "our E.F. Hutton moment." The E.F. Hutton financial firm was running a series of commercials at the time with the catchphrase: "When E.F. Hutton talks, people listen."

"We're the only X-plane that ever did an actual performance, an eight-minute routine that we did every day of the airshow," Francis stressed. "Everybody stopped to watch it!"

What stunned show-goers was the X-31's ability to maneuver at extremely high angles of attack in post-stall flight. The trick that gripped them most was called the Herbst Maneuver.

"It's equivalent to a swimmer doing a flip-turn in a swimming pool," Francis explains. "You go through zero forward flight speed. X-31 was capable of pointing independent of its flight path." Amazingly, the airplane could fly at up to 70 degrees angle of attack (AoA).



"It would pop to very high angle of attack as a giant speed brake. As it would get there, it would start to do a velocity vector roll. And while it was doing that it would reverse its flight path in the opposite direction. The airplane would slow to between 40 and 50 knots as it's turning around and pointing itself," Francis said.

The X-31 would then accelerate "out the back end" employing its powerful General Electric F404-GE-400 engine. Thrust vectoring made the Herbst maneuver possible, allowing the X-31 to maneuver even as its aerodynamic control surfaces became ineffective beyond the stall at high angles of attack.

"Thrust vectoring is just another way to generate the forces and moments that you have to have to keep an airplane controllable and stable," Francis said. "We used the canards on the front as an insurance policy if we lost control at high AoA. Those surfaces could articulate enough to bring the airplane back under control." The successful performance of this unprecedented maneuver was the result of work carried out by a team including NASA, the U.S. Navy, the U.S. Air Force, Rockwell Aerospace, the Federal Republic of Germany, and Daimler-Benz (formerly Messerschmitt-Bölkow-Blohm, MBB).

DARPA program managers guided the effort, with Francis being the driver behind many X-31 program breakthroughs.

His four “maneuver milestones,” crafted in 1992 to accelerate what had been “a slow-moving program,” emphasized combat effectiveness and tactical utility.

“The first was steady flight at high angle of attack,” said Francis. “The second was rolling around the velocity vector in both directions through 360 degrees. The third was getting into post-stall conditions rapidly because the airplane had to be able to do a high-speed dynamic entry. The fourth was putting all the pieces together to do the Herbst Maneuver.”

Dr. Wolfgang Herbst, a noted engineer with MBB, originated the idea behind the X-31. While working on a fighter design in the 1970s that would become the Eurofighter Typhoon, he recognized the potential of extreme maneuvering to defeat short-range air-to-air missiles. Herbst theorized that an airplane capable of flying in the post-stall regime with full control at high angles of attack could gain a significant tactical advantage in close-in aerial combat.

The idea, dubbed “supermaneuverability,” made its way to the United States and to Rockwell’s Mike Robinson. Robinson began a collaboration with Herbst that led to the EFM program. In 1983, Allburn, already at work on the X-29 program, called then-U.S. Air Force Capt. Michael Francis to a meeting at DARPA.

“Jim Allburn asked if I would listen to a pitch on an interesting project with the name ‘SNAKE.’ That was the acronym for Super Normal Attitude Kinetic Enhancement,” Francis recalled.

Wolfgang Herbst and Robinson gave the presentation.

“I thought, ‘that was nice,’” Francis said. “I’ll read about it in the papers.”

Nearly a decade after Herbst first envisioned supermaneuverability, Allburn was looking for the right person to push the X-31 program forward. Greenlighted by Allburn in 1986, X-31 was the first international X-plane. Funding from the Nunn-Quayle amendment (an international defense cooperation initiative) provided the foundation for the American-German research program.

Initially, the Air Force had been sought as the program executive agent. But by the late 1980s, USAF interests had turned to the F-117 and the Advanced Tactical Fighter program – the effort that would yield the F-22 Raptor. Convinced that stealthy fighters with advanced sensor fusion would conduct aerial combat mostly from beyond visual range, the service had lost interest in the X-31, an airplane dedicated to research for close-range combat.

So DARPA turned to the U.S. Navy. But the Navy was moving slowly with X-31, averse to the kind of risk such a program came with.

“I had been working on [Michael Francis] for a couple years to try to get him to come to DARPA,” Allburn remembered. “The program was experiencing uncertain funding and delays. Mike did a bang up job.”

Francis came to DARPA from the Air Force Space Division (now Space and Missile Systems Center, Los Angeles AFB) in 1991 and grasped what needed to be done.

“Transferring the airplane from Rockwell at Palmdale [California] to NASA Dryden [Edwards AFB] was the way to go,” he said.

With NASA collaboration and the maneuver milestones created by Francis in place, progress came swiftly.

Said Francis, “Instead of a test plan calculated to be 12 years long we accomplished most of our original goals in six months!”

The X-31 EFM program logged a staggering 559 research flights. Along the way, two X-31 airframes graphically demonstrated the

advantages Herbst had theorized in close-in dogfights. The X-31s soundly outperformed aircraft lacking thrust vectoring, using post-stall maneuvers to achieve a potential kill ratio of 30 to 1.

“We proved that thrust vectoring really was an advantage for this form of combat,” Francis said.

In addition, the X-31 program proved the viability of a (simulated) tailless fighter design and was one of the first experimental aircraft programs to introduce and employ helmet-mounted displays for pilots.

Unfortunately, one of the X-31s was lost in an accident in January 1995. The pilot, Karl-Heinz Lang, ejected safely before the aircraft crashed in the desert just north of Edwards AFB. An investigation indicated the crash resulted from ice accretion on the aircraft’s nose boom.

Though the EFM program concluded in 1995, the X-31 lived on. By 2000, the remaining airframe was transferred to NAS Patuxent River, Maryland, to undertake ESTOL (extremely short takeoff and landing) research for the U.S. Navy and the German government under the VECTOR program.

“I think we did more with X-31 than has been done with any X-plane,” Francis said. “We conquered the stall barrier more completely than any other airplane.”

X-45 J-UCAS

In 1992, Francis “started dabbling” with the “UTA concept” (Unmanned Tactical Aircraft) that would later yield the X-45 and become the \$4 billion Joint Unmanned Combat Air Systems (J-UCAS) program.


As the X-31 EFM program came to a close, Mike Leahy was an Air Force major working at Wright Patterson AFB with the Directorate of Science and Technology, the acquisition arm that governed the Air Force Research Laboratory (AFRL). He had a background in manned robotics and taught robotics at the Air Force Institute of Technology.

“I saw what DARPA was doing with Global Hawk and Darkstar,” Leahy said, “and I kind of became the staff go-to guy for UAVs.”

Maj. Gen. Richard Paul, commander of AFRL at the time, told Leahy he wanted AFRL to be involved with X-45, and wanted Leahy to be the one doing it.

Still attached to AFRL, Leahy spent 1996 and 1997 working with initial DARPA X-45 program manager Larry Birkelbaw to craft what is arguably the most ambitious and capable unmanned combat aerial vehicle to date.

“This was really a systems program. We were integrating existing things into a system in a way that had never been done before. You had to create something that could not be ignored.”



"We didn't want to do a tech demonstrator," Leahy explained. "Larry and I wanted something that was on a direct path to being a production system. We did operational studies and built them into the program. The contractor candidates needed to first prove that their system could accomplish a useful military mission. The one we picked was suppression of enemy air defenses [SEAD] because that was challenging and needed."

"We picked a cost target for the program initially at one-third the cost of the Joint Strike Fighter," Leahy said. "You had to be able to do the mission, do it affordably and use simulation for a lot of research." They selected Boeing to build the X-45.

The year was 1999 and Birkelbaw decided to leave DARPA. Offered the opportunity to become the X-45 program manager, Leahy leapt at the opportunity.

"If you were in the field I was in at the time, this was your dream job. There are not many flight test programs where you can say you were there from concept to something that flew," Leahy said.

Conducted in three phases, the X-45 program aimed to produce autonomous UCAVs that eventually would "fly in packs, searching for enemy anti-aircraft missile launchers, working together to destroy them under the supervision of a human operator, who could be located anywhere in the world," said Leahy following the first flight of the X-45A on May 22, 2002, at NASA's Dryden facility in California.

He remained DARPA's X-45 PM for just one more flight, leaving to complete his senior military education at the Industrial College of the Armed Forces later in 2002. Reflecting on the program, he noted that X-45 was much more than a flight test vehicle.

"This wasn't just an X-plane program," Leahy said. "This was really a systems program. We were integrating existing things into a system in a way that had never been done before. You had to create something that could not be ignored."

Ultimately, that's just what X-45 did, becoming J-UCAS in 2003. By 2005, the effort had demonstrated formation flight with two X-45As, culminating in a "graduation combat demonstration" over the California desert

DARPA's X-45 program produced a UAV with tactical aircraft performance, which demonstrated formation flight, manned/unmanned teaming, and a coordinated strike against simulated radars and missile launchers. X-45 could be seen as the proof-of-concept vehicle for the unmanned combat air vehicle (UCAV) concept.

in which the X-45As simulated a coordinated strike against radars and missile launchers.

Unfortunately, the J-UCAS program and the X-45A were consigned to history following the 2006 Quadrennial Defense Review, which unexpectedly ended the program.

Looking back, Leahy considered his time as X-45 program manager a highlight of his career.

"Seeing paper transform into reality with a living, breathing X-45 was really satisfying. And the people we had... everybody pulled together and was equally committed on both the government and Boeing sides. They were committed to success and being there."

Each of these highly visible programs has contributed to future combat and civilian aircraft designs. For combat aircraft, the X-29 and X-31 demonstrated the advantages of unconventional planforms, tailless designs, thrust vectoring, and other technologies that contributed to high angle of attack maneuvering and "pointability" in a dogfight scenario. The technologies and systems demonstrated by DARPA directly influenced the requirements, and ultimately the designs for the F-22 and F-35 operational aircraft. The X-45 brought tactical aircraft performance, stealth qualities, and autonomy to the slow, pedestrian aircraft of the UAV world at the time, representing a giant step toward the realization of an unmanned, autonomous, combat air vehicle.

DARPA and experimental aircraft share a storied history, and the story surely has chapters to come. Consider the current Experimental Spaceplane program. If things go well, the nation eventually will have an entirely new class of spacecraft system that provides short-notice, low-cost access to space with aircraft-like operability, reliability, and cost-efficiency.

DARPA AT THE TACTICAL EDGE

TIGR & RAA: Information from the bottom up

By Eric Tegler

Marching within a Greek phalanx, a hoplite soldier could really only follow the momentum of the formation in which he strode, set in motion by the generals commanding it. He had little idea of the disposition of the enemy and no say regarding the form in which he met it.

Flash forward almost three millennia, and a parable of a far more informed, flexible soldier had become ingrained in America's Army culture. Yet as it began to engage in Iraq and Afghanistan in the 1990s and 2000s, the U.S. Army was, in fact, a top-down organization where intelligence was concerned.

As the situation in both theaters morphed into counterinsurgency, it was becoming obvious that something else was needed. On the complex, ever-changing ground in Iraq and Afghanistan, those with the most important, timely information were in the lowest, boots-on-the-ground echelons.

On patrol after patrol, they observed the environment, identifying the patterns of their adversaries and potential allies, and discerning the urgent needs of coalition partners. The experience they accumulated was priceless, but passing it on to their command, to new units, and to partners was difficult at best.

A new paradigm in which intelligence, information, and action emanated from the bottom up required new tools. DARPA took on the task of building and deploying those tools at the tactical edge.

A Virtual Notebook – TIGR

In 2002, Maj. Patrick Michaelis was an assistant professor at the United States Military Academy at West Point, New York, teaching a course on organizational change. He indulged his interest in the subject by running a peer-to-peer web community for Army officers then called *Platoonleader.org*.

Michaelis saw in it, and in a similar site called *CompanyCommand*, the potential to share with their peers and those who routinely replaced them the hard-won knowledge that units gained on operations. It was a break from traditional Army practice, in which the operational acumen a unit built up typically lived and died with its deployment.

The idea of developing a platform to collect and share operational knowledge laterally wasn't well received by those Michaelis ran it by, with one exception. Col. Paul E. Funk, brigade commander with the 1st Cavalry Division, and a former commander of Michaelis, thought it had merit. He put it in front

of Maj. Gen. Peter W. Chiarelli (then-commander of the 1st Cavalry Division), and Lt. Gen. Thomas F. Metz (former III Corps commander).

With the support of these senior officers, a platform called CAVNET was created and fielded in Iraq in 2004. Michaelis went with it as battle command and chief knowledge officer with 1st Cavalry. The CAVNET website resided on SIPRNet (the U.S. military's secure internet), cataloging the latest after-action reports and intelligence from units in Iraq.

The information was available to a swath of personnel broader than the typical intel officers at division or brigade level. Captains and lieutenants in companies and platoons could use CAVNET to prepare and plan for their next patrol. By 2004, Iraq had become a small-unit war and CAVNET gave those units the power of aggregated operational knowledge.

Today, Michaelis is executive officer to the under secretary of the Army. He emphasizes that the idea behind CAVNET was central to the fight in Iraq. "We knew that sharing ground-level information and data was more important than [focusing on] top-down information."

Bringing TIGR to Life

Even before CAVNET made it into the field, DARPA was developing another command-and-control platform called Command Post of the Future (CPOF), a networked information visualization system that facilitated communication between dispersed commanders. CPOF was part of an ambitious broader vision within DARPA to leverage and share data from a suite of sensors worn by a soldier. Beginning in the early 2000s, CPOF saw extensive use in wartime in Iraq and Afghanistan, and it was continuously updated in the field in response to many users' inputs. CPOF became an official program of record for the Army in 2006. As successful as CPOF was, Mari Maeda, Ph.D., a Program Manager in DARPA's Strategic Technology Office at the time, thought a simpler knowledge-sharing solution that could be fielded even more quickly could be even more relevant.

After speaking with a variety of captains, lieutenants, and sergeants returning from Iraq, she realized that they still shared mission information via

notebooks, digital cameras, and laptops with Microsoft Office. Operations information was saved locally and sent to headquarters.

"When a new unit would rotate in, the local command would say, 'Here's a hard drive with all the information we've collected about people who live in this area and its particulars,' Maeda said. "It was woefully inadequate."

Maeda was thinking along the same lines as Michaelis. In 2005, the latter had returned from Iraq. He was promptly invited to speak to DARPA about CAVNET. Upon meeting, the pair began to brainstorm.

"She pulled me aside and said, 'If you could create something for your next deployment, what would it look like?,'" Michaelis remembered. His response was a combination of CPOF, CAVNET, and Sharepoint, the web-based collaborative framework, oriented to lower-level units.

"It would be a virtual notebook of information and events so that soldiers on patrol could do their own mission planning and debriefs," Michaelis said.

"There was an 'Aha!' moment when we realized we could create a map-based application that could disseminate this information," Maeda recalled.

To do so at the "speed of DARPA," she asked Joseph Evans, a professor of electrical engineering and computer science at the University of Kansas, to help assemble a small team to bring to life what would become known as the Tactical Ground Reporting (TIGR) system.

By dint of their previous work together at the National Science Foundation, Maeda knew that Evans had a background in online games. Young soldiers, she realized, understood video games intuitively. Maeda wanted TIGR to have that kind of instantly comprehensible user interface (UI).

"Even before TIGR got started, we had chatted about various ideas about collecting information for the warfighter," Evans said.

To develop TIGR into a web-based, visually arresting information-sharing application, Evans reassembled the team with which he'd developed online games, including software developers Steve Pennington, Benjamin Ewy, and Mike Swink. Borrowing ideas from CAVNET, CPOF, and other command and control (C2) applications, they put together a demo for DARPA in three weeks.

"The demo basically looked like the initial TIGR system would look," Evans said with amusement. "It was just that if you clicked on anything that wasn't exactly where we pointed, it would crash."

The prototype integrated the kind of video-game UI that DARPA was after. Swink said that the developers made this a priority, even using some of the same artists with whom they'd designed video games to achieve the effect Maeda envisioned.

"What she said made sense to me," Swink said, "but, not having worked with the military before, I was stunned that they didn't have it."

The other main TIGR selling point Swink emphasized was giving small units the ability to access and update maps, and use them in planning patrols. Intelligence officers at brigade level and above already had the capability, but those going outside the wire did not.

"They found it really difficult to get up-to-date maps of the areas they were going to patrol," Swink explained. "They loved that they could load in the latest maps at platoon level, print them, and walk out the door with them right before patrol."

Swink knew this from experience, from sitting beside soldiers in Camp Victory or Camp Taji in Iraq and talking to those who went on patrol. With that input, Swink could make coding improvements to TIGR on the fly. The 1st Cavalry, 1st Brigade Combat Team, Funk, and Michaelis had made this possible, teaming with DARPA to experiment with TIGR during exercises in 2005-2006. The feedback from soldiers was immediate, Maeda remembered.



Paratroopers from the 3rd Brigade Combat Team, 82nd Airborne Division, use Joint Tactical Radio System radios and prototype Joint Battle Command-Platform handhelds to communicate during a recent field exercise at Fort Bragg, North Carolina. The handhelds ran a Tactical Ground Reporting (TIGR) application for soldiers who dismount from their vehicles but need to remain in communication with higher headquarters.

Said Maeda, "We had a construct [in the TIGR environment] including 'Events, People, Reports.' One of the soldiers pointed out the need for a [repository] called 'Collection,' saying that when about to embark on the next patrol, he wanted to look for relevant events in a single folder, i.e. Collection."

The feedback was embraced by TIGR developers, who brought an agile development mind-set to the system. Unlike more formal Army software, TIGR didn't have to be fully perfected before it rolled out. Eighty percent was good enough, with successive iteration in the field improving the product more effectively.

Todd Hack, a project manager with General Dynamics Mission Systems (which later acquired TIGR) worked on TIGR development stateside. "If anything took more than a few months, it never made Maeda happy," he said.

Timeliness was key. The 1st Brigade Combat Team was redeploying to Iraq in late 2006 with Funk commanding and Michaelis as brigade S-3 (operations officer). TIGR would follow it in early 2007 in the form of a small DARPA developer and training team. Operator feedback moved quickly from Swink and others in theater to Maeda and Evans back at DARPA.



LEFT: A screenshot of the Tactical Ground Reporting System (TIGR), a virtual notebook developed by DARPA that could be used by soldiers in the field to share tactical information. BELOW: Soldiers learn to use Force XXI Battle Command Brigade and Below, also known as FBCB2, while attending Digital University. The second generation of FBCB2, Joint Battle Command-Platform (JBC-P), will allow soldiers to track friendly forces with greater accuracy, better maps, and a new user interface.



"The thing about working with DARPA at the time was that they were very tightly connected to the warfighter, whether in the field at Fort Hood or in Iraq," Michaelis said. "The programmers were sitting right next to the folks on the ground. We were providing capability from the bottom up rather than the top down. That was pretty revolutionary for the time."

Revolutions always meet resistance, and TIGR was no different. In Washington, D.C., and Iraq, mid-level managers were uncomfortable with it.

Becoming an Indispensable Tool

"We had fights inside the Pentagon and inside the theater," Michaelis acknowledged.

"We got big push-back," Maeda agreed. "There were two big objections: One was that this was going to break the network; the other was that this was not an Army system – it was not officially derived from stated requirements."

Concerns over the resources TIGR would require were valid, Swink admitted.

"The networks were just awful early on in Iraq and Afghanistan for any smaller base with 100 guys or less," he said.

The outposts from which patrols frequently launched saw periods of up to 8 hours without a SIPRNet connection. TIGR's strength as a web application, accessible from anywhere, was network-vulnerable. When it was down, troops still needed information, such as maps and post reports, on their operational environment.

DARPA found a way to get around the problem by linking dispersed servers. When connectivity was down in one location, other servers were caching TIGR data. While the network was up, servers constantly loaded information in the background, making it accessible even when connections dropped. DARPA developers prioritized and filtered TIGR traffic, building compression algorithms for simplified documents.

"The [software] architecture that we designed limited information shared to what was necessary," Maeda explained. "For example, only reduced-size photos would get pushed out. We really thought hard about designing the application so it didn't burden the network."

Michaelis refers to TIGR as a sort of precursor to the Cloud, with federated servers providing real-time operational intelligence to widely spread formations.

A company commander could log in to TIGR and, using its game-like, map-referenced interface, plot out patrol routes noting changes to buildings, terrain, roads, or bridges. TIGR included before-and-after photos of the changing landscape, updated continuously by its users. As it developed, it allowed not only for notes and maps but also for voice recordings, digital photos, and GPS tracks to be collected and searched.

"One of the many reasons people liked TIGR was because we made sure it got the best, freshest satellite imagery," Maeda said. "The decision-makers said, 'Oh, that imagery is available to everybody.' That didn't mean that someone in an outpost could download it."



ABOVE: Maj. Summer Favors, Capt. John (Mack) Turner, and Capt. David Gerdes demonstrate the Command Post Computing Environment prototype at Aberdeen Proving Ground, Maryland, on May 16, 2017. LEFT: Soldiers connect to OEF SIPRNet via PRC-117G during RTO/Operator Training for the first time.



In Iraq, TIGR grew from 25 trainers and 10 active users in 2007 to more than 1,000 users in early 2008. By October 2009, it was fielded to 15 brigade combat teams and 23,000 users. Ultimately, TIGR would expand to 97,000 user accounts, with many non-account holders regularly leveraging its information.

"This was a juggernaut," Michaelis said. "Once you were exposed to TIGR, there was an epiphany. You understood its potential. But you had to have leaders who understood that when you're building something on the fly, there will be failures."

Fortunately, Chiarelli was one of those. The general went out of his way to support the TIGR team, according to Maeda. "He said, 'Whenever you encounter a wall, let me know.' And we did encounter many walls."

From getting onto SIPRNet to expediting the Army's extended information assurance process, command support at the highest levels was vital. An early meeting with Multinational Force Iraq Commander Gen. David Petraeus was crucial, Swink said.

"That's what got us onto SIPRNet instead of just our own small network."

As the use of TIGR expanded in Iraq and Afghanistan, more capability was added, including access to broader Army databases, a link-analysis tool, and

a highly prized Google Street View-like feature called Patrol View. Being in theater allowed Swink and other developers to update TIGR code and have soldiers test it the same day as opposed to going through a six-month Army development/test/deployment cycle.

By the time General Dynamics Mission Systems acquired TIGR in 2010, it was firmly part of the Army tactical toolkit, with an unexpected viral quality.

"The beauty of TIGR," Hack explained, "was that because it was on SIPRNet, even when soldiers rotated out of the theater, they still kept up with what was going on. It was kind of like an Army social media tool where a guy might walk into a master sergeant's office and say, 'Let me show you what happened on TIGR yesterday.' They'd log back into a server in Baghdad and he could say, 'This is where I was operating, and I knew these people.'"

"It was neat to see how this had become ingrained in the tactical culture of the Army as the way to share and distribute information," Michaelis said.

TIGR transitioned to a formal Army program in 2012, and as the service resets to face new, more sophisticated threats, the system is being folded into a common operating environment. The Army's Command Post Computing Environment (CPCE) will integrate TIGR data with that of several other C2 systems into a single user interface overlaid on a single map. Despite its consolidation, TIGR's significance remains.

"It pushed intel out to the edge," Evans said. "Mari [Maeda] really wanted to promulgate TIGR as a battle command system, not an intel system."

Even today, it is the "best of breed" for disseminating an operating picture to small units, Hack added. In fact, a civilian version of TIGR is in use by a group called African Parks that works to protect animals and drive poachers from 10 nature reserves in Africa. Former British special forces experts running the effort knew and valued TIGR from



their experience in Afghanistan. They contacted Swink to facilitate deployment of the system.

For Maeda, TIGR was also a step beyond the work that DARPA typically engages in.

“This wasn’t just publishing research,” she said. “We actually had people using the system. The captains, lieutenants, and soldiers who patrolled every day needed [tactical] knowledge as much as the intel guys. That’s where things changed in the wars in Iraq and Afghanistan.”

Coalition Capability – RAA

The rise of ISIS and its subsequent territorial grab in Iraq produced a new kind of tool with DARPA roots. Remote Advise and Assist (RAA) arose from a chaotic situation in Baghdad and an American administration politically opposed to intervention to reverse it.

By summer 2014, ISIS had overtaken the northern and western provinces of Iraq and was headed for Baghdad as Iraqi regular forces folded. A special operations element that included fewer than 50 U.S. Army Special Forces soldiers had been deployed to Baghdad International Airport to support the U.S. embassy.

They realized their most effective partners were Iraqi special operations forces (ISOF) whose decade-long training by U.S. Special Forces had given them the skill and will to fight ISIS. ISOF was desperate for American technical and tactical assistance. But U.S. policy meant that Army special operators could only assist them peripherally. The question was how to do it?

In the years following the initial development of TIGR, DARPA remained active in researching and developing other distributed command and control technologies, particularly those that could be actuated from hand-held devices.

In 2012, DARPA launched a program called Transformative Apps, which sought to develop a military mobile-apps marketplace similar to Apple’s App Store. Applications would be available on Android phones and tablets, cellular or satellite-enabled. “Trans Apps,” as it was called, built on the methodology developed for TIGR and aimed to nurture a new model for acquiring, introducing, and maintaining software.

An artist’s concept depicts the goal of DARPA’s new Secure Handhelds on Assured Resilient networks at the tactical Edge (SHARE) program: Create a system where information at multiple levels of security classification can be processed on a single hand-held device using a resilient secure network that links devices without needing to route traffic through secure data centers.

Maeda was again a program leader, and she admits that her team naively thought their success with TIGR would ease acceptance of Trans Apps. But the same battles over resources and Army oversight were revisited. As such, the DARPA Trans Apps team started small, rolling out select applications for trials in Afghanistan.

“The DARPA thread, through Trans Apps, to place new functionality on hand-helds led to several apps that became popular in Afghanistan,” Evans said.

One such application was the Android Tactical Assault Kit (ATAK) based on NASA/Air Force Research Laboratory-developed WorldWind software. DARPA provided early funding for ATAK through Trans Apps.

ATAK allows Android phone/tablet users to maintain collective situational awareness, coordinate with other users, quickly issue commands, text messages, enemy/friendly locations, and make full “9-line” calls for fire/close-air support. The software and operational logic, which grew out of Trans Apps, influenced the ATAK concept, Evans said.

“These activities at the tactical edge have been offshoots of the systems that led to Trans Apps and RAA. Newer initiatives with tactical handhelds stem from the same tree as well.”

Another branch of that tree is represented by DARPA’s own PCAS (Persistent Close Air Support) system, a fire-support application with both air and ground components, which the Agency began developing in 2012. Like ATAK, PCAS resided on mobile devices, improving situational awareness and automating the process for coordinating close air support. The PCAS ground element was tested on 500 tablets in Afghanistan in 2013, and the combined system was demonstrated with the Marine Corps in 2015.



Maj. Graham Perry, the officer in charge of 4th Air Naval Gunfire Liaison Company (ANGLICO), Force Headquarters Group, checks his Android Tactical Assault Kit to see possible target areas in Cape Wrath, Scotland, on April 23, 2018. 4th ANGLICO was in Scotland to take part in Joint Warrior 18-1, an exercise that furthers their readiness and effectiveness in combined arms integration, small-unit tactics, and land navigation.

“The idea was to have something like ATAK on a phone,” Evans said, “and use that for sharing coordinates for strikes and getting that information to the air-support layer.”

Back in Iraq, ISOF forces were already using Android applications (Offline Maps, Google Earth) on mobile phones and tablets to help communicate, conduct reconnaissance and targeting, track their movements, and navigate.

Maj. Eric Roles, a company commander with the 1st Battalion, 3rd Special Forces Group, is a DARPA fellow who led the RAA effort. In an essay he co-authored for the National Defense University, he explained that U.S. Special Forces soldiers in Baghdad quickly realized that a program like ATAK combined with the ISOF's existing hand-held coordination could yield battlefield results while keeping American SOF operators at the required arm's length from their ISOF partners.

They made an urgent request to Special Operations Command, Central (SOCCENT) to improvise “Remote Advise and Assist kits” using American, ISOF, and locally procured components. Combining Frontier Iridium GPS trackers with cell phones and apps including ATAK on the limited Iraqi terrestrial cellular networks, the kits enabled U.S. special operations forces to communicate, track, and share limited data with ISOF. Such back-up support quickly facilitated the capture of several high-value ISIS targets.

Later RAA kits with wireless routers and additional hand-held satellite communications capability (some of the funding for which came from DARPA, Evans added) ranged in capacity from supporting a handful of users up to more than 100. These kits are regarded as instrumental in the success of ISOF forces in turning back ISIS in Anbar, Tikrit, and other areas with remote U.S. support.

A DARPA study in cooperation with SOCCENT and the Naval Postgraduate School assessed conditions favoring the application of RAA. The study identified necessary prerequisites (a trained, capable partner force, for example) and lessons learned. It also examined the technical framework of RAA.

“That way of working with coalition partners was pretty atypical, but our study showed that it worked,” Evans said. “We found that being able to share information at that level was unbelievably helpful.”

Though deployed in circumstances specific to Iraq, RAA's use of a hand-held application to coordinate remote fire support from U.S. forces merged with nearly a decade of DARPA development to empower small, dispersed frontline units with up-to-the-minute knowledge.

DARPA's latest C2 development — the Secure Handhelds on Assured Resilient networks at the tactical Edge (SHARE) program — maintains its continuum of soldier-led battlefield technology. SHARE will add a new level of information assurance to tactical operations, allowing information at multiple security classification levels to be processed on a single hand-held device leveraging existing networks.

This ability to share classified tactical data securely is yet another step in DARPA's mission to provide the warfighter on the ground with information and tools that were once only available to the upper echelons of command, ensuring that the rank and file possess a level of situational awareness that makes them safer, more lethal, and more efficient than ever before.

SEMICONDUCTOR SAFARI

Exotic materials beyond silicon

By Dan Green

At this moment, the semiconductor industry is justifiably focused on the impending inflection in Moore's Law, the famous technology projection that underlies the astonishing progress of the microelectronics era, and its potential impact on the continuing advance and dominance of silicon technology. That means it's worthwhile to consider another visionary perspective in the history of microsystems. As the nascent Advanced Research Projects Agency (ARPA) passed its one-year anniversary in 1959, Prof. Richard Feynman at Caltech delivered one of his most famous and consequential talks, titled "There's plenty of room at the bottom."

Similar to Gordon Moore, Feynman anticipated many of the opportunities for technological advance that lie within the realm of microscale systems. However, Feynman took a much broader view that highlighted the exotic possibilities that would emerge with the ability to manipulate structures at the atomic scale. DARPA has played a central role in bringing many of these "exotic" structures, including semiconductors, to life with capabilities beyond the binary-processing feats that silicon electronics have been pulling off for half a century.

Feynman's talk inspired a resurgent interest in the 1980s as his speculative notions of nanotechnology and the ability to tailor materials at the atomic scale were becoming tantalizingly close to realization. At that time, emerging crystal growth techniques were enabling the creation of a class of materials known as compound semiconductors, where the exact chemical composition or alloy could be varied at the atomic level on a layer-by-layer basis. In particular, gallium arsenide (GaAs) and its alloys emerged as new wonder materials that allowed transistors to operate well beyond the performance limits of silicon. DARPA identified the potential for the new GaAs transistors to move electrons faster and therefore operate at higher frequencies in the electromagnetic spectrum. While this new technology would not displace silicon technology for highly integrated digital logic, DARPA anticipated its value to enable the next generation of radar and communications systems. To that end, DARPA in 1988 took over the baton

DARPA's MIMIC technology, particularly the techniques of integration that came out of it, enabled the DOD to make radios and radar systems that engage the spectrum at higher frequencies and bandwidths than ever before.



A U.S. Air Force F-16C Fighting Falcon breaks off and flies back toward its patrol area after concluding refueling from a KC-135 Stratotanker over Afghanistan in support of Operation Freedom's Sentinel on March 11, 2018. Key to the development of the Joint Direct Attack Munitions carried by this aircraft are high-performance chips developed through DARPA's Monolithic Microwave Integrated Circuit program, which also enabled the radio frequency (RF) and millimeter-wave circuits needed in precision weapons.



A soldier with 3-6 FA, 1BCT, 10th Mountain Division, navigates the new Precision Fires-Dismounted system that allows soldiers to view live-streaming full-motion video from unmanned aerial vehicles through an app on approved smartphones. The ubiquitous cellphone technology used worldwide owes its existence in part to the development of the GaAs semiconductor industry made possible by DARPA-funded research.

from the Office of the Secretary of Defense (OSD) to run the Microwave and Millimeter Wave Integrated Circuit (MIMIC) program, which OSD had stood up two years earlier.



The MIMIC program, which ran until 1995, had a profound impact on industry, as it sought to develop the ways and means of integrating higher-frequency materials and components into military-relevant technologies, such as radio and radar, and to establish a reliable industrial base to do those things. In fact, the MIMIC program was able to realize GaAs transistor technology that led to a new class of RF (radio frequency) “front end” components. The front end of an RF system is the amplifier technology that sends and receives signals in the electromagnetic spectrum. DARPA's MIMIC technology, particularly the techniques of integration that came out of it, enabled the Department of Defense (DOD) to make radios and radar systems that engage the spectrum at higher frequencies and bandwidths than ever before. The use of GaAs technology in DOD systems continues to this day.

Beyond defense applications, the high-frequency GaAs amplifiers provided a key piece of the puzzle to the commercial sector as it sought to establish newly developed cellular phone technology in the 1990s. GaAs transistors enabled handheld phones with small batteries to establish the critical communications link to the towers. To this day, every smartphone contains a small piece of GaAs to perform this critical function, and the United States enjoys a dominant share of the suppliers of this multi-billion-dollar semiconductor industry as a result of DARPA's investment in the MIMIC program.

The success of GaAs technology proved the defense relevance and the commercial viability of semiconductor technology beyond silicon and made a once-exotic research material into a commodity technology. However, even as GaAs was maturing into an industry, researchers sponsored by the Office of Naval Research (ONR) and elsewhere had already begun to identify the next leap in semiconductor materials. Wide band gap semiconductor (WBGs) materials were identified as promising due to their ability to move electrons rapidly like GaAs but also handle large electric fields as well. This combination of high current capability and high voltage drives the ability to deliver more RF power. While several candidate materials were being developed around the world, DARPA considered gallium nitride (GaN) and its alloys the most promising and established the Wide Band Gap Semiconductor-RF (WBGs-RF) program to rapidly advance the technology.

The WBGs-RF program sought to mature an unproven material with obvious potential into an industrially relevant technology that could further the cause of national defense. Launched in the early 2000s, the program started with GaN material that was delivered on small semiconductor wafers (2-inch diameter) that had large numbers of micropipes or holes, similar to a slice of Swiss cheese. From this inauspicious state, the WBGs-RF program systematically addressed the materials challenges before progressively and successfully taking on the device and circuit-design challenges. Ultimately, the GaN technology delivered on its promise and is now being used in the next generation of radar technology, such as the Navy's Air and Missile Defense Radar (AMDR). And there is way more to come: GaN is now part of the technology portfolios of all major RF semiconductor players. Once again the United States has a dominant role in this emerging market.

DARPA's efforts have enabled compound semiconductors to move from the research fringes to a mainstream semiconductor industry. They have also driven mainstream silicon technology to embrace variants that include alloys of silicon. In particular, the mixing of silicon with germanium is a technology that DARPA championed during the 2000s with the Technology for Efficient, Agile Microsystems (TEAM) program. Germanium (Ge) was the material basis for the original Bell Labs transistor created in 1947; however, this material was soon abandoned in favor of silicon due to germanium's reliability problems and the processing advantages of silicon. The insight that brought Ge back was that while it was not useful on its own, a materials stack that included a mix of Ge with Si, or SiGe, allowed for the atomic-level engineering of devices with enhanced RF performance to be built right alongside conventional silicon logic devices in high density. This technology did not possess the complete performance advantages of the other compound semiconductors, such as GaAs and GaN, but it had the ability to produce chips that mixed analog and digital functions. This trait proved highly useful, and SiGe technology is now dominant for delivering low-power commercial solutions for applications like local WiFi amplifiers and now potentially for phased-array systems for 5G communications radios.



LA203 PHOTO VIA WIKIMEDIA COMMONS

TRIQUINT SEMICONDUCTOR IMAGE

UNIVERSITY OF MICHIGAN PHOTO

The high-profile successes of these GaAs, GaN, and SiGe transistor technologies exemplify the ongoing innovation that is possible through manipulation of the crystal structure at the atomic scale. Yet, even these efforts took place within the relatively well-understood paradigm of transistor physics that was established by the silicon semiconductor community. The wider frontier of microsystems goes beyond the electronic properties of materials, as illustrated by some of the more exotic technologies that have emerged along the way. For instance, DARPA championed microelectromechanical systems, or MEMS, in the 2000s through a series of programs that leveraged semiconductor processing to create tiny structures that move and flex rather than just conduct electrons. MEMS technology blossomed with DARPA support into a multi-billion-dollar industry today. MEMS motion sensors and actuators are the heart of protective air-bag systems, navigational and gaming products, and even chips with millions of micro-mirrors that project movies onto theater screens.

In more recent years, DARPA has pioneered work to leverage so-called phase-change materials to create RF switches that operate by a toggle in a material's crystal structure rather than through conventional transistor action. This wholesale shift to another physical basis and set of materials for digital switching has enabled the demonstration of RF switches with cut-off frequencies in the terahertz (THz) region of the spectrum, which is about 1,000 times the frequency of cell phone operation.

The menagerie of semiconductor technologies that has emerged, in part through sustained DARPA investments, reinforces Feynman's notion of the broad opportunity that exists in the realm of the microsystem. While these micro- and now nano-landscapes are not the *terra nova* they were at DARPA's inception, there's still plenty of room at the bottom!

TOP LEFT: A pristine 2-inch single crystal gallium arsenide wafer. Purple features are a reflection of a nitrile glove. TOP RIGHT: DARPA's GaN-on-diamond high electron mobility transistor (HEMT) demonstrated improved thermal properties that may lead to better performance for RF systems. ABOVE: DARPA-supported researchers at the University of Michigan made significant progress with a timing and inertial measurement unit (TIMU) that contains everything needed to aid navigation when GPS is temporarily unavailable. The single-chip TIMU prototype contains a six-axis IMU (three gyroscopes and three accelerometers) and integrates a highly accurate master clock into a single miniature system, smaller than the size of a penny. This IMU chip integrates breakthrough devices (clocks, gyroscopes, and accelerometers), materials, and designs from DARPA's Micro-Technology for Positioning, Navigation and Timing (Micro-PNT) program.



JUMP-STARTING INNOVATION

DARPA's Grand Challenges

By Ivan Amato

When Congress released its 515-page Defense Authorization Bill for Fiscal Year 2001, a few lines about 40 pages into the document set DARPA onto what would become one of the agency's most effective and, with increasing likelihood, world-changing technology quests. "It shall be a goal of the Armed Forces to achieve the fielding of unmanned, remotely controlled technology such that ... by 2015, one-third of the operational ground combat vehicles are unmanned," the tome stated. It was, in effect, an audacious technology challenge – involving advances in robotics, artificial intelligence, sensors, and other components – to get U.S. military personnel out of harm's way. DARPA leadership took the challenge on and its resolve strengthened in the early years of the millennium as coalition convoys were coming under attack and suffering casualties on roads in Iraq.

The first public display of that resolve unfolded on March 13, 2004, when DARPA hosted its first Grand Challenge, designed to accelerate technology development in autonomous ground vehicles. Even though no contestants came even close to completing the Challenge's 142-mile desert course between Barstow, California, and Primm, Nevada, to claim the \$1 million prize, the event – and two follow-on challenges to accelerate the development of autonomous, ground-vehicle technology – would prove to be a transformative moment in the history of driverless vehicle technology.

"That first competition created a community of innovators, engineers, students, programmers, off-road racers, backyard mechanics, inventors, and dreamers who came together to make history by trying to solve a tough technical problem," Lt. Col. Scott Wadle, DARPA's liaison to the U.S. Marine Corps, said in 2014 on the 10-year anniversary of the event. "The fresh thinking they brought was the spark that has triggered major advances in the development of autonomous robotic ground vehicle technology in the years since." It also would stand as the initial embrace by DARPA of a powerful and effective tactic and incentive for mobilizing the research community to deliver audacious new technological capabilities.

"Since 2004, DARPA has consistently used prized-based Challenges to develop innovative solutions to some of the most difficult national security problems," DARPA Director Steven Walker said in March 2018 at a White House roundtable on American innovation.

A month later, on April 18, at the 34th Space Symposium in Colorado Springs, DARPA formally announced its Launch Challenge in which

competitors will vie for a top prize of more than \$10 million based on their performance in two separate launches of small satellites at different launch sites on short notice – changing the paradigm from years of pre-planning and preparation to days. Beyond realizing the technical goal of demonstrating agile responsive launch capability, the new Challenge also is about nurturing a nascent industry and community of space-launch providers, Fred Kennedy, director of DARPA's Tactical Technology Office (TTO), told the gathering.

"We've been operating in a mode where every space launch we execute is a national event – requiring large amounts of planning, coordination, and financial resources," noted Todd Master, manager of the DARPA Launch Challenge. "We look to a future use of space that shifts from strategic to tactical, with dramatically more frequent, low-cost access to orbit becoming the norm. The DARPA Launch Challenge is intended to show how that future can be realizable."

High-profile competitions, challenges, and prize incentives have been conspicuously effective at motivating innovators to expand the boundaries of human performance and technology. Awards like the Emmys and Oscars have inspired creative individuals and communities to take their respective arts to new heights. The prestige and medals of the Olympic Games have driven thousands of competitors to deliver faster, higher, farther, and more capable athletic performance than ever. With its prize of helping the Allies win World War II, the Manhattan Project to beat German scientists and engineers in the race to nuclear weapons amounted to one of the most consequential technology challenges of all time.



Technology challenges have been providing incentive to inventors and innovators for centuries. Among the earliest, still most famous (surely partly because of Dava Sobel's 1995 blockbuster book *Longitude*), and most transformative of the early technology challenges was the British government's offer in 1714 of £20,000 (equivalent to more than \$2 million today) to whomever would deliver a solution to the disastrous inability to determine longitude precisely at sea, a failing that routinely led to ship navigators being uncertain of their locations and thereby to extensive loss of life and property. A primary impetus for the prize and challenge was a wartime disaster on Oct. 22, 1707, in which 1,550 British sailors aboard four weather-defeated warships in the Isles of Scilly off the Cornish coast lost their lives, allegedly due in part to the navigators' inability to accurately calculate their positions. No contender ever took home the award, but in several portions and most of it more than a half-century after the initiation of the prize, John Harrison, a carpenter and clockmaker, received compensation from the British government for his maritime chronometer that could provide navigators with at-sea timekeeping of unprecedented precision and stability.

It took Nicolas Appert, a chef, confectioner, pickle maker, and culinary innovator, only 15 years to collect the 12,000 franc prize (equivalent to more than \$40,000 today) offered by the French government in 1795 for his invention of techniques to preserve large amounts of food. This was a capability that Napoleon Bonaparte, who famously said "an army marches on its stomach," coveted to solve the problem of food shortages that had been limiting his military ambitions. Appert collected the prize in 1810 for what amounted to a sterilizing bottling technique that not only

The winner of DARPA's 2005 Grand Challenge – "Stanley" – and the team that built the autonomous vehicle. DARPA's Assured Autonomy program leverages some of the research pioneered more than a decade ago in DARPA's 2005 Grand Challenge.

aided Napoleon's hungry troops but also marked the beginning of modern food packaging.

A famous technology prize hastened the day of the first New York-to-Paris flight. After six aviators died and three others were injured in attempts to claim the \$25,000 prize (equivalent to some \$300,000 today) first offered in 1919 by New York City hotel owner Raymond Orteig for that transatlantic first, Charles Lindberg, in his plane, the *Spirit of St. Louis*, captured the prize and world admiration in 1927. Like many technology prizes, the Orteig Prize triggered technology development investments that exceeded the value of the prize itself while also raising public interest and awareness of the wonders and opportunities associated with new technology, in this case advances in aviation.

The list of technology prizes is a long one. In 1829, a trio of inventors shared 550 British pounds for building a locomotive that weighed in under 6 tons but could pull 20 tons at 10 miles per hour. In 1865, John W. Hyatt collected \$10,000 from the firm Phelan & Collender in Albany, New York, for developing a cellulose plastic alternative to ivory for manufacturing billiard balls, a materials innovation that presaged the ongoing age of polymers. In 1979, a team led by Paul MacCready, Jr., won the 100,000-pound Kremer Prize for Human Powered Flight for crossing the English Channel

DARPA's Network Challenge. The challenge was to be the first to submit the locations of 10 moored, 8-foot, red weather balloons at 10 fixed locations in the continental United States.



in their Gossamer Albatross. In 1997, IBM's Deep Blue chess team took the \$100,000 Fredkin Prize for developing the first computer program that could beat a reigning world chess champion. And, to finish here with just one of many other possible examples, the firm Mojave Aerospace Ventures took the \$10 million Ansari X Prize in 2004 – the same year that DARPA got into the Grand Challenge game – for its achievement of two manned suborbital flights to an altitude of 100 kilometers within two weeks.

DARPA's maiden adoption in 2004 of the prize-based challenge approach to accelerating technology development resided in the call from congressional leaders for innovations that would make it possible to substitute automated vehicles for warfighter-operated ones as a

way to reduce U.S. casualties in conflicts. The DARPA Grand Challenge initially attracted 106 applications from teams wishing to compete for the \$1 million prize, but in the end, none of the 15 finalists made it farther than 9 miles into the 142-mile desert course. Eighteen months later, however, in a follow-on competition, four teams of an original 195 completed a 132-mile desert course in Nevada with a Stanford University team taking the \$2 million prize for the winning performance of its vehicle, Stanley. In 2007, just two years later, six teams completed the Urban Challenge, a second follow-on competition to spur innovations in automation relevant to the cityscape battlespaces that were becoming more prevalent in current conflicts. With its vehicle,



Two autonomous vehicle entrants of the DARPA Urban Challenge handle an intersection and each other.

“Boss,” a team from Carnegie Mellon University in Pittsburgh crossed the finish line first and won the \$2 million carrot. The second and third place teams from Stanford and Virginia Tech took home \$1 million and \$500,000, respectively.

“The progress, from Challenge to Challenge, was truly spectacular,” said Brian Pierce, current Director of DARPA’s Information Innovation Office. Perhaps an even more important and consequential result of this trio of autonomous vehicle Challenges than the ability of teams to finish the courses was that the Challenges effectively recruited a diverse talent pool into a collective goal that required the honing of novel skill sets and mind-sets. The Challenges created a new community with targeted technology goals. Said Pierce, “Many of those who are leading self-driving car development today competed in the Challenges.” A decade after the Urban Challenge, this technology community is engaged in what now seems like a sprint toward a new era of transportation in which the occasional human drivers will be what turns heads in massive fleets of otherwise automated driverless vehicles.

Since the Urban Challenge in 2007, DARPA has rolled out an average of more than one new technology challenge each year. In the 2009 “Red Balloon Challenge” — more formally “The Network Challenge” — the winning team from MIT managed to leverage the internet and social networking techniques to crowdsource a national network of balloon searchers who successfully located all 10 secretly released weather balloons. The feat, which earned the team the \$40,000 award, took just under nine hours from the surprise announcement of balloon release to the pinpointing of all ten balloons.

During the fall of 2011, DARPA’s Shredder Challenge had computer scientists, puzzle enthusiasts, and other interested competitors striving to provide

U.S. troops with the capability to reconstruct documents that had been destroyed in war zones while also identifying intelligence vulnerabilities that might reside in the shredding practices of the U.S. national security community. With \$50,000 in prize money for motivation adding to the satisfaction of trying to boost national security, 69 entrants wrestled with “five separate puzzles in which the number of documents, the document subject matter, and the method of shredding were varied to present challenges of increasing difficulty,” the archived Shredder Challenge website states. “To complete each problem, participants were required to provide the answer to a puzzle embedded in the content of the reconstructed document.” By racking up the highest overall score based on the number and difficulty of the problems solved, the team that called itself “All Your Shreds Are Belong To U.S.” was named the winner of the Shredder Challenge on Dec. 2, 2011.

One of the Agency’s higher-profile challenges, the DARPA Robotics Challenge (DRC), culminated in early June 2015 on the Fairplex grounds in Pomona, California. The humanoid robots of eight of the 23 participating teams managed to complete a course designed with eight disaster-relief tasks — such as driving a car, climbing stairs, turning a valve, and handling power tools — within the allotted 60 minutes. The top three teams took home the \$2 million, \$1 million, and \$500,000 prizes. The press had a field day with the inevitable falls and foibles of a far-from-mature robotics technology, but like the autonomous vehicle challenges before it, the DRC helped to create a can-do research community whose own ambitions and confidence were bolstered by the competition. The same can be said of the Cyber Grand Challenge, or CGC, which ran in real time in Las Vegas on Aug. 4, 2016, in front of 5,000 spectators, many of them hackers who were attending DEF CON, a storied annual gathering of hackers.

“I’m enormously gratified that we achieved CGC’s primary goal, which was to provide clear proof of principle that machine-speed, scalable cyber defense is indeed possible,” said Mike Walker, the former DARPA program



manager who launched the Challenge in 2013, at the conclusion of the live competition. “I’m confident it will speed the day when networked attackers no longer have the inherent advantage they enjoy today.”

Foundational to DARPA’s initial and ongoing embrace of incentive prizes and challenges was the legislative permission granted by Congress, as specified in 10 U.S. Code § 2374a, to even offer cash prizes. Otherwise known as the “DOD Prize Authority,” the law enables the Agency “to award cash prizes and other types of prizes that the Secretary determines are appropriate to recognize outstanding achievements in basic, advanced, and applied research, technology development, and prototype development that have the potential for application to the performance of the military missions of the Department of Defense.”

Adding to the federal-level authority to offer money-backed incentive Challenges was the America COMPETES Reauthorization Act of 2010, which gave a legal nod to all federal agencies to conduct competitions and award prizes. According to Challenge.gov, the government’s clearinghouse for federal prize challenges, since the website was launched in 2010, federal agencies have offered more than 825 prizes in which more than 250,000 participants have vied for “more than \$250 million in prize money along with other valuable and unique incentive prizes.” For its part, DARPA is closing in on its 20th technology challenge with the latest additions including the Spectrum Collaboration Challenge, the Subterranean (SubT) Challenge, and the Launch Challenge.

The civilian and nonprofit sectors are rife with ambitious incentive challenges. Among the most well-known of these are the Grand Challenges Explorations grant program that the Bill & Melinda Gates

Team Tartan Rescue’s robot, CHIMP, earned the third-highest score, 18 points, during the Dec. 20-21 DARPA Robotics Challenge Trials 2013. The team’s lead organization was the Carnegie Mellon University National Robotics Engineering Center.

Foundation has been running since 2008 and the XPRIZE Foundation, which describes itself on its Twitter landing page this way: “Massive, world changing competitions. Disruptive, breakthrough technology and innovation. We are a catalyst for the benefit of humanity. An innovation engine.”

There are numerous arguments that champions of particular Challenges have summoned to justify moving forward. A short list of reasons for considering organizing a Challenge that DARPA leadership compiled in 2017 zeroed in on these:

- Solve a problem that has eluded solutions through traditional means
- Spur innovation by way of nontraditional or unexpected solutions
- Develop a new R&D community that attracts diverse talent and widens performer base
- Encourage STEM (science, tech, engineering & math) in student populations
- Attract the kind of public attention that can assist recruiting, inform Congress, and generally enhance the DARPA “brand”

DARPA has tended to embrace a Challenge approach when that tactic stands a better chance of achieving all of these desired outcomes in



concert than do the agency's traditional technology-development approaches – such as a standard DARPA program and contract research process, convening of industry days, and familiar social media, public affairs, and outreach activities. One template DARPA leadership has developed to help qualify a potential “Challenge-appropriate” problem centers on the following five questions:

- Can the problem be stated in enough detail to bound the parameters into a competitive paradigm?
- Can this problem be solved through more traditional means?
- If the problem is solved, who benefits and how significant is that benefit?
- Is the problem feasible, that is, can it be solved without violating laws of physics?
- Are there clear metrics by which success in solving the problem can be assessed?

The current momentum in autonomous vehicle technology is retroactively bestowing wisdom on the organizers of the three autonomous vehicle challenges over a decade ago. It is too early to say if the same eventually will apply to, say, the more recent Cyber Grand Challenge, but that only is a reminder that the problems these Challenges are trying to solve are hard and take time to solve. “In the same way that the Wright brothers’ first flight – although it didn’t go very far – launched a chain of events that quickly made the world a much smaller place, we now have seen for the first time autonomy involving the kind of reasoning that’s required for cyber defense,” said CGC director Walker. “That is a huge advance compared to where the cyber defense world was yesterday.”

The recently launched Subterranean (SubT) Challenge “aims to provide previously unimaginable situational awareness capabilities for operations underground” by warfighters and first responders, said Timothy Chung, program manager in DARPA’s TTO. “We’ve reached a crucial point where advances in robotics, autonomy, and even biological systems could permit us to explore and exploit underground environments that are too dangerous for humans,” added TTO Director Fred Kennedy. “Through the DARPA Subterranean Challenge, we are inviting the scientific and engineering communities – as well as the public – to use their creativity and resourcefulness to come up with new technologies and concepts to make the inaccessible accessible.”

The seven automated contenders, or cyber-reasoning systems, of the Cyber Grand Challenge finals stand ready to compete on Aug. 4, 2016.

When he was deputy director for policy in the Obama White House’s Office of Science and Technology Policy, Tom Kalil, now Chief Innovation Officer with Schmidt Futures, began what for him has been a continuing campaign to instill a Challenge ethic into the broad public mindset. His aim, he said in an interview earlier this year, has been to increase, diversify, accelerate, and otherwise elevate the adoption of audacious innovation as a value that can further the public good. Challenges and incentive prizes, Kalil noted in 2012 during a speech at the Information Technology and Innovation Foundation, “shine a spotlight on an ambitious goal without having to predict which team or approach is most likely to succeed. Incentive prizes help us reach beyond the ‘usual suspects’ to increase the number of minds tackling a problem, bringing out-of-discipline perspectives to bear and inspiring risk-taking by offering a level playing field.”

DARPA added the Grand Challenge approach to its bag of technology development and community-building tricks only in 2004. As the current trajectory toward a world rife with autonomous vehicles indicates, the Agency has enjoyed some potentially world-changing success with its first three challenges, all of them designed to advance autonomous vehicle technology. Other challenges, among them the DARPA Robotics Challenge and Cyber Grand Challenge, are showing signs that they too, in time, could prove to be as prescient and consequential in their respective domains as the first DARPA Challenges.

As Kalil defined them, Grand Challenges are “ambitious yet achievable goals that capture the public’s imagination and that require innovation and breakthroughs in science and technology to achieve.” That definition maps well onto DARPA’s own institutional way of thinking and doing. It also reveals that DARPA itself amounts to an ongoing Grand Challenge with the ambitious yet achievable goal of providing a culture of innovation that consistently identifies and helps realize those not-yet technologies that need to become here-now technologies.

DARPA'S QUEST FOR A BENEFICENT CYBER FUTURE

The innovations it will take to assure our cyber security

By Brian M. Pierce

Much of modern life takes place online in cyberspace. A variety of threats have emerged to imperil the future of the cyber domain, where more and more of our commercial and social activity unfolds. Securing cyberspace has become a priority at the highest levels of government, including the White House. DARPA is at the leading edge of our nation's pursuit of a beneficent cyber future.

Legacy of Cyber Past

Roughly a half-century ago, DARPA (when it was known as the Advanced Research Projects Agency, or ARPA) program managers conceived, implemented, and demonstrated the ARPANET, the proof-of-concept precursor of the internet. The principal advance of the ARPANET resided in the introduction of packet-switched communications to achieve a dramatic advance in the efficiency by which the information packets wound their way through available wires and other inter-computer linkages. By 1977, the ARPANET had grown into a national network connecting university, commercial, and government research activities.

The initial engineering of the ARPANET focused on providing basic communication services, and accountability was not included as a primary design goal. These design choices were intentional. According to an historical account of the original network protocols developed at ARPA, "... since this network was designed to operate in a military context ... survivability was put as a first goal and accountability as a last goal ... An architecture primarily for commercial deployment would clearly place these goals at the opposite end of the list." Instead of building in reliable attribution mechanisms, user behavior was loosely governed by rules and norms. For example, at the MIT AI Lab circa 1982¹, it was simply decreed that "It is considered illegal to use the ARPANET for anything which is not in direct support of government business." The critical need for stronger security and control mechanisms was not yet apparent.

During the 1980s, a commercial version of the ARPANET, dubbed the internet, grew rapidly, and by 1988 served roughly 60,000 internet-connected computers ("hosts")². Nov. 2, 1988, was an important reality-check day in the history of the internet, with the release and rapid spread of the Morris worm to approximately 10 percent of internet

hosts³. It was a before-and-after moment. After the Morris worm, users became aware of the need for stronger security in computers and networks.

During the 1990s, the Department of Defense (DOD) recognized networking as a critical enabler for military operations⁴. As such, DOD aggressively integrated networking, computation, and automation throughout military systems. Adversary nation-states responded by exploiting these networks, and it was during this period that the notion of an "electronic Pearl Harbor" entered the public consciousness⁵ and the status of cyberspace ascended to that of a domain of active conflict alongside the traditional domains of sea, land, air, and space. The connected nature of cyberspace enables an adversary to strike at any geographic location in the United States, and at a wide range of targets, including the power grid, refineries, chemical plants, airline reservation systems, enterprise and wide area networks, the financial markets, the bulk power markets, communications systems, natural gas pipelines, and water and wastewater utilities.

Affecting the Cyber Present

During the 2000s, DARPA funded selected efforts in cybersecurity. In 2010, DARPA centralized cyber R&D within a new office – the Information Innovation Office (I2O). The creation of I2O included the hiring of several well-known computer science and cyber experts from academia and industry, and greatly expanded DARPA's cyber program portfolio and investment. To stimulate the interest and involvement of the cyber community, DARPA hosted, on Nov. 7, 2011, a Cyber Colloquium that was attended by approximately 700 researchers, operators, and other stakeholders from industry, academia, and government^{6,7}. The DARPA Cyber Colloquium was a bright signal to the broad U.S. cyber R&D community that DARPA would bring its unique project-centric approach

army@DARPA:~\$ PLAN X

to the development of future cyber capabilities, both defensive and offensive.


In the seven years since the colloquium, DARPA's programs have had great impact⁸. Recognizing that cyber threats to physical systems such as vehicles could be devastating for the military, DARPA in 2012 created the High-Assurance Cyber Military Systems (HACMS) program to secure embedded computing systems in mission-critical commercial and military assets against cyber attacks. DARPA's work in this threat space resulted in heightened awareness of the need for improved automobile cybersecurity⁹ and led to changes throughout the industry. DARPA's technical approach to secure software featured formal methods (based on mathematical techniques) to ensure that software reliably does what it is specified to do, and nothing else. DARPA demonstrated these formal methods by developing a secure mission system for an autonomous helicopter. The agency now is working with the DOD to transition tools for building software with much greater cyber resiliency, and envisions a day when formal methods and other advanced tools for creating provably secure software will be adopted by the defense procurement process.

Understanding that cyberspace had become a warfighting domain, DARPA initiated the Plan X program in late 2012 to create a mission command system on which the military can plan, conduct, and assess cyberwarfare in a manner similar to kinetic warfare. DARPA coordinated closely with multiple DOD cyber stakeholders, most notably U.S. Army Cyber Command (ARCYBER), to develop the Plan X prototype system. In 2017, according to an article posted on the U.S. Army website

"Plan X" was a \$120 million program at the Defense Advanced Research Projects Agency, which currently is making it easier for humans to, among other things, visualize a network and its components, to automate the task of identifying as hostile or benign the anomalies that might appear on that network, to provide intuitive symbology that accurately conveys to users the status of various components of a network, and to make it easier for even inexperienced users to take action to prevent hostile parties from gaining access to and causing damage to a network.

by ARCYBER, "Plan X is a battle command system for cyberspace operations which possesses technology that firmly places our forces at significant advantage in cyberspace."¹⁰

Recognizing the need to engage cyber attackers at machine speed rather than human speed, DARPA created and executed the Cyber Grand Challenge (CGC) program to automate the process of finding, fixing, and exploiting software vulnerabilities. CGC featured a capture-the-flag-style competition in which so-called Cyber Reasoning Systems devised by the CGC teams competed to find and patch flawed code and to prove the inefficacy of opponents' defenses. The Cyber Reasoning Systems were put to the test at DARPA's CGC Final Event – the world's first automated hacking tournament – in front of the DEF CON conference in Las Vegas on Aug. 4, 2016. All of the seven CGC Cyber Reasoning Systems competing in the finals succeeded in automatically identifying and fixing software



Capt. James McColl and Capt. Justin Lanahan, both cyber officers at U.S. Army Cyber Command, took part in a weeklong "hackathon" in Arlington, Virginia, in July 2015, in support of the continued development of Plan X.

flaws, sometimes within seconds of the introduction of the software by the competition organizers.

U.S. commercial and government networks are subject to nearly continuous cyber attack. DARPA is developing automated, scalable algorithms that identify anomalous behavior in networks indicative of these threats and the security compromises that can result. These methods triage events, classify known threats, and identify novel threats to dynamically detect attacks. DARPA is also tackling the challenge of real-time monitoring and defense of even the largest enterprise networks such as those in the DOD.

Vulnerabilities in the cyber domain are only increasing. Consumer imaging products, such as smartphones, have become ubiquitous, and it is estimated that about 2 billion images and videos are uploaded to social media every day. At the same time, a growing proportion of this visual media has been manipulated. Many manipulations are benign, performed for fun or for artistic value, but some are for adversarial purposes, such as propaganda or disinformation campaigns. The forensic tools that are available today for detecting manipulation lack robustness and scalability and address only some aspects of media authentication; an end-to-end

platform to perform a complete and automated forensic analysis does not exist. DARPA is leveling the playing field, which currently favors the image manipulator, by developing technologies for the automated assessment of the integrity of an image or video¹¹.

Vision of Cyber Future

Deterrence of any attack depends on several factors, most importantly the adversary's estimate of their probability of successful attack and the anticipated benefits, and the adversary's estimate for costs that will result from the response. It is prudent to assume that potential adversaries work to refine these estimates by probing our defenses and observing our offensive capabilities. In cases where deterrence is reliable, that is, where the costs to the adversary are likely to exceed the benefits, we may find it advantageous to assist the adversary in estimating these probabilities. We might achieve this by way of, for example, a demonstration of offensive capability. In cases where deterrence is not reliable, that is, where the probability that the benefits outweigh the costs is high enough that an adversary might rationally contemplate proceeding with an attack, it will be advantageous to hinder the adversary's ability to estimate these probabilities.

Deterrence in the cyber domain is proving to be even more complex than in the traditional warfighting domains due in large part to the following three factors:

1. **An exponentially growing domestic attack surface:** Our modern society depends on information and information systems, and information technology (IT) is deeply embedded in critical



DARPA hosted the Cyber Grand Challenge (CGC) led by I2O Program Manager Mike Walker on Aug. 4, 2016, in Las Vegas, Nevada. Seven prototype machines squared off against each other and competed for nearly \$4 million in prizes in a live network competition. Team Mayhem took home the top prize of \$2 million, which it used to further its research.

infrastructure, commercial services, cyber-physical systems, and other components of the constructed landscape. Our dependence on IT and the cyber domain is growing exponentially both in terms of scale (i.e., number of users/hosts, number of networks and network nodes, volume of storage) and in terms of the complexity of the applications (e.g., self-driving cars and other autonomous systems). Few of these systems are resilient to cyber attack, and so they present an inviting attack surface for potential adversaries. Metaphorically, we have built for ourselves a “cyber glass house” at which adversaries may freely cast stones, and we want our house to be resilient against these attacks. The way we construct our cyber structures, therefore, plays a foundational role in determining how resilient these structures will be amidst those who will throw stones.

2. Lack of visibility and limited intelligence: Many cyber attacks, the so-called advanced persistent threats, remain undiscovered for extended periods, while other attacks have never been conclusively attributed even with significant forensics effort. As a result, it is difficult, if not impossible, to estimate with confidence the cyber capabilities of a potential cyber adversary. Moreover, while defensive cyber technology development is a large and growing commercial activity, offensive cyber technology is typically developed in secret by both nation-states and diverse criminal enterprises; the potential for technological surprise by one or more of these entities cannot be ignored. In the cyber domain, we need far greater visibility into and situational awareness of adversarial activity. We need to know who is throwing stones against our house.

3. Empowered adversaries that act with impunity: Software is the ultimate democratic technology. It is proving to be a facile weapon for adversaries ranging from so-called “script kiddies” to peer-nation intelligence agencies. These adversaries are empowered by the ability to re-use readily available malwares, access large-scale computing resources – both legal (commercial cloud) and illicit (botnet) – and hide their activities in the flood of internet communications and transactions. Few cyber attackers ever suffer any consequences, and so they act without restraint.

If we are to deter cyber attacks, we must develop accurate and calibrated cyber response capabilities. This is how we will stop the stones from hitting our house.

Here is what a reliable cyber deterrence capability would have to include:

- **Cyber resilience:** In the event of a cyber attack, the information and operational technology the United States uses to manage and control its critical networks and systems must operate through the attack and be rapidly recoverable afterward.
- **Cyber situational awareness:** The United States must be able to detect, understand, and attribute in a timely fashion any subtle or overt escalations in the intensity of cyber conflict and adversary attacks on



Because a widespread, prolonged loss of electrical power would pose a grave threat to the United States, DARPA is developing technology that cyber security personnel, power engineers, and first responders can employ to rapidly restore electrical systems subjected to cyber attack.

our critical networks and systems involving the cyber domain.

- **Cyber accurate response:** The United States must have the capability to mount an accurate, timely, effective, and appropriately scaled cyber response to any cyber attack in a calibrated way that discourages further escalation.

DARPA is sponsoring a broad portfolio of programs to develop the technologies necessary to realize these elements required for an effective cyber deterrence.

In the cyber-resilience category, DARPA's approach includes techniques to harden systems against cyber attack and techniques to enable systems to operate correctly even when subject to cyber attack. DARPA's initial investments featured formal methods but this portfolio has expanded to include other approaches. Inspired

by big-data approaches that have impacted numerous industries, DARPA saw that big-data technology could improve the way we build, debug, verify, and maintain software. DARPA created a corpus of hundreds of billions of lines of code – so-called “big code” – and data-mining engines (specialized for use in software) to extract useful properties, behaviors, and vulnerabilities of the program components in the corpus. This knowledge opens the way toward new mechanisms for automatically constructing, improving, and repairing complex software.

DARPA is also developing the necessary design, analysis, and verification tools to allow system engineers to design in cyber resiliency and to manage trade-offs as they do other nonfunctional properties when designing complex embedded computing systems. In addition, DARPA is developing techniques that address the need for long-lived, survivable, and complex software systems that are robust to changes in the physical and logical resources provided by their underlying computational ecosystem. These will reduce high software maintenance costs and stave off premature obsolescence of otherwise functionally sound, legacy software systems. In principle, these abilities could enable the creation of “100-year software.”



U.S. commercial and government networks are subject to nearly continuous cyber attack. DARPA is developing automated, scalable algorithms that identify anomalous behavior in networks indicative of these threats and the security compromises that can result.

A substantial and prolonged disruption of electric power would have profound economic and human costs for the United States. From a defense perspective, a major power outage could hamper military mobilization and logistics and impair the capability to project force. DARPA is developing technology that cybersecurity personnel, power engineers, and first responders can utilize to accelerate restoration of cyber-impacted electrical systems.

In the category of situational awareness in the cyber domain, the goal is to achieve a comprehensive ability to detect and monitor cyber attacks in the making. At present, cyber adversaries are often able to operate on U.S. networks for extended periods without discovery. DARPA is developing a number of technologies to enhance situational awareness of attacks on networks and systems by, for example, providing high-fidelity visibility into component interactions during system operation across all layers of software abstraction. These techniques will automatically or semi-automatically “connect the dots” across multiple activities that are individually legitimate, but collectively indicate malice or abnormal behavior. This should enable the prompt detection of advanced persistent threats. Malicious actors in cyberspace currently operate with little fear of being caught. This is because it is extremely difficult, in some cases perhaps even impossible, to reliably and confidently attribute actions in cyberspace to individuals. DARPA is developing techniques to enable reliable attribution of malicious cyber actions and to increase the government’s ability to reveal publicly the actions of individual malicious cyber operators without compromising sources and methods.

In the category of deploying accurate and calibrated cyber response capabilities, DARPA envisions high-intensity cyber operations executed by computers under human supervision. Such semi-automated response systems would enable operators to create and analyze cyber effects more rapidly and accurately than unaided human operators. Fully automated cyber defense capabilities, such as those developed in DARPA’s Cyber Grand Challenge, will help in this cause. These will be integrated with human-centric cyber operations planning and execution capabilities, such as those developed under DARPA’s Plan X program. This technology will automatically evaluate the defensive readiness of software and networks during operations, triage and verify system security issues, determine adversary intent, and guide operator responses. Because botnets pose

a significant threat to national security, DARPA is exploring the feasibility of countering malicious botnets and similar large-scale malware.

Toward a Cyber-Safe Era

The cyber domain has become central to our modern way of life, and it is a matter of national security¹². As such, the ability to deter cyber attacks has become a strategic technology priority. For its part, DARPA is working to develop technologies to enable U.S. cyber deterrence and is collaborating with DOD cyber stakeholders to deploy and improve cyber deterrence capabilities. This includes a variety of efforts with USCYBERCOM and the military Services to participate in exercises, develop concepts of operation, evolve prototype systems, mature the technology base, and transition cyber-deterrence technologies to operations. In addition, DARPA is developing technologies to create software systems that are secure by design rather than by constant patching in response to newly discovered vulnerabilities; provide greater visibility into network operations for enterprises and service providers; and enable cyber response capabilities that are accurate, robust, and safe. Taken together, the new cyber technologies DARPA is developing hold promise for a cyber future in which the benefits of the cyber domain are assured.

1. https://dspace.mit.edu/bitstream/handle/1721.1/41180/AL_WP_235.pdf, Section 2.4.2 Etiquette
2. <http://www.computerhistory.org/internethistory/1970s/>
3. https://en.wikipedia.org/wiki/Morris_worm
4. https://en.wikipedia.org/wiki/Network-centric_warfare
5. <http://www.washingtonpost.com/wp-srv/washtech/daily/may98/cyberattack052498.htm>
6. <http://archive.defense.gov/news/newsarticle.aspx?id=65988>
7. <https://www.wired.com/2011/11/darpa-hackers-cybersecurity/>
8. <https://www.youtube.com/watch?v=mgAjvmgr08w>
9. <https://cyber.stanford.edu/spotlight-research-automobile-cybersecurity-policy>
10. https://www.army.mil/article/182571/plan_x_harnesses_collaboration_innovation_to_build_mission_command_system_for_cyberspace
11. <https://www.nbcnews.com/tech/tech-news/inside-government-agency-designing-tech-fight-fake-news-n865586>
12. Richard Danzig “Surviving on a Diet of Poisoned Fruit: Reducing the National Security Risks of America’s Cyber Dependencies,” 2014 (https://s3.amazonaws.com/files.cnas.org/documents/CNAS_PoisonedFruit_Danzig.pdf?mtime=20161010215746)

ENTERPRISE DISRUPTION

Shaking off entrenched ways is hard, but it is the gateway to better technology frameworks for all of the military domains

By Fred Kennedy

“... the first guy through the wall – he always gets bloody. Always. It’s the threat of not just the way of doing business, but in their minds, it’s threatening the game. But really what it’s threatening is their livelihoods. It’s threatening their jobs. It’s threatening the way that they do things.”

– Boston Red Sox owner John Henry, in *Moneyball*, 2011

In early January 2018, many news outlets reported that a classified spacecraft, dubbed Zuma, failed to reach orbit and reportedly plummeted to Earth. If accurate, it would be just the most recent and visible loss in a long line of satellite launch failures since the dawn of the Space Age.

Zuma illustrates the undeniably harsh reality of space access: the ever-present possibility of catastrophic failure. Zuma and the SpaceX booster that was supposed to take it to orbit undoubtedly endured years of component, subsystem, and systems test, validation, and certification. Even after such exhaustive precautions, accidents will happen, and our assets won’t always work as designed. We’ve known this for decades, and yet our approach to creating new space-based capability hasn’t changed much, if at all. We continue to construct exquisite, one-off systems that often cost billions of dollars. In recognition of the high risks associated with this strategy, the Air Force’s senior officials including the Secretary, have been calling for radical changes in the way we conduct spacecraft development and deployment. But their progress is severely hampered by lower-level bureaucracies, which are themselves fine-tuned for risk aversion through decades of statutory and regulatory accretion.

Decision-makers are constrained by the need to sustain immense quantities of legacy hardware and software, slow innovation time-scales, a reliance on small numbers of high-value assets, and (not surprisingly!) a deeply systemic aversion to risk. And space is by no means the only place where these kinds of endemic limitations exist. Each Service has its own rigidly defined order of battle, which governs and constrains the way the respective Service plans for and creates the warfighting capabilities they require for future battles. Our industrial base is set up to address the same questions we’ve been asking for decades, and is very good at delivering incremental improvements

and upgrades to our current capabilities. Perhaps we need to revisit the most basic assumptions that underlie our current force structure, strategy, and tactics, and ask ourselves, “Is there a better way?” I suspect there is – but it’s going to be disruptive, not just to our adversaries, but also to ourselves.

Consider again the space domain. Space operations and acquisition have emerged as the enterprise most visibly in need of urgent reconfiguration. Senior leaders from the military departments and the intelligence community have decried our inability to respond to adversary threats to our space systems in a timely fashion. Gen. John Hyten, commander of U.S. Strategic Command, has stated, “I’m worried about the future. Somehow, this country lost the ability to go fast.” He’s right. By the time we’re able to field a new space-based communications or ballistic missile early warning capability – the last of which required almost a quarter-century to go from requirement to objective constellation – our adversaries have had more than enough time to deconstruct our approach and evolve their countermeasures accordingly.

Similarly, the Air Force and Navy are facing a slew of challenges arising from emerging and unprecedented threats in the air and at sea. The proliferation of anti-ship ballistic missiles, long-range strike capabilities, ubiquitous ISR (intelligence, surveillance, and reconnaissance), and integrated air defense systems have threatened our ability to reliably dominate the sea and airspace domains. For seven decades, the United States has dominated airspace and projected power through the multi-ship naval formation known as the Carrier Strike Group – this is how America fights “from the sea.” Tom Ricks, in a September 2017 essay in *Foreign Policy*, asserts that “... our carriers must anticipate a future threat environment that will be characterized by massive barrages, whether of cruise and ballistic missiles, swarms of micro-drones, subsurface drones, speedboat borne IEDs, or some combination – that seek to overwhelm escorts’ ability to engage.” Ricks



U.S. Air Force F-35A Lightning II aircraft. This state-of-the-art stealth fighter took more than 20 years to develop and deploy. DARPA aims to both shorten the development and improve the efficacy of future air platforms by introducing disruptive new platforms and employment strategies.

sees parallels between today's carrier and the World War II era battleship, noting that carrier-based fighters nullified the formidable firepower of the battleship, making it obsolete in the same manner that volleys of cruise missiles and hypersonic weapons threaten today's carrier. Given the amount of time our adversaries have had to analyze our systems and strategies, and the emerging threats we're facing, wouldn't it behoove us to consider what comes next? And do it now, before the crisis is upon us?

Our armed services are built around the idea that maintaining our strategic advantage depends on our ability to stay ahead of the evolutionary trajectories of our adversaries' threat stances, that is, to be in a position to counter their moves with relevant and effective countermoves. This evolutionary outlook is the defining characteristic in every one of the Armed Services' strategies for the future. And why not? If you expend enormous resources and time on a capability such as the F-35, emergence of a threat tends to focus the jet's operators not on what could replace the F-35, but instead on how best to improve the F-35 and its supporting technologies, given the extraordinary outlay of money and long developmental period that preceded its fielding. It's plain vanilla sunk cost, but a seductive argument nevertheless.

Therefore, while adversary nations rapidly field technologies that could make our traditional approach to air dominance impossible, our services are scrambling to develop a sixth-generation jet fighter and

Long Range Strike Bomber (LRS-B, or B-21 Raider), doubling down on the attributes and capabilities of the previous models. Envisioned to replace our existing air platforms, we're reliably informed that these prospective aircraft feature designs with improved strike, enhanced sensing, and stealth capabilities that will continue to allow us to conduct operations inside contested airspace. These systems are likely to be expensive, and we cannot expect to see them fielded in the next decade. Will our adversaries be as slow?

For its part, the U.S. Navy recognizes the danger posed by swarms of cheap anti-ship cruise missiles (ASCMs), a proliferating threat which current defense systems will have difficulty handling at scale. In 2016, the destroyer USS *Mason* successfully used the Standard Missile 2 (SM-2) and Evolved Sea Sparrow Missile (ESSM) to counter two incoming ASCMs fired by Houthi rebels off the coast of Yemen. At a cost of \$3 million per SM-2 and \$1 million per ESSM, we see how an adversary can asymmetrically impose cost on U.S. forces. The Houthi missiles — likely Chinese C-802 variants — have been estimated to set back their owners no more than \$500,000 each. Michael Armstrong, writing in *The National Interest* in 2017, noted, "*Mason's* attackers fired ASCMs one or two at a time. Countries employing anti-access area-denial strategies could have aircraft, ships, submarines, and/or coastal batteries firing them in dozens." What then?

Our deliberate, methodical, work-until-you-meet-the-letter-of-the-specification approach to solutions fails to accommodate the real, fundamental challenges in each domain. The existing culture and approach in the Department of Defense (DOD) and its performer base is partial to developing capabilities that are evolutionary extensions of known warfighting constructs and that rely primarily on monolithic high-value assets such as aircraft carriers or exquisite satellites. These assets are expensive to design, develop, field, and



oped, fielded, and sold ever more sophisticated air defense systems, designed to pierce our aircraft's stealthiness, across the electromagnetic spectrum. As Dave Majumdar wrote in a January 2018 article for *The National Interest*, "While the Russians – and the Chinese – have not yet cracked the problem, it is clear that stealth is becoming much less of an advantage over time, though perhaps no less expensive an acquisition. Eventually, Moscow will find a solution to the stealth problem as the cyclical struggle between offense and defense continues *ad infinitum* – it's just a matter of time."

As adversary-sensing capabilities have improved, however, we've relied on those familiar evolutionary approaches to keep our stealth a step ahead. Unfortunately, developing that next iteration is expensive, and now yields only marginal results. If our adversaries' sensors can detect (and shoot down) F-35s, we'll either need an upgraded version of the F-35 or another fighter. The Air Force and Navy predict that it would take until 2030 for a traditional development and procurement approach to deliver their sixth-generation fighter jet. Meanwhile, the Air Force doesn't expect to be flying the B-21 until 2025 – an optimistic assessment, given historical precedent.

In the face of tightening development timelines by adversaries, the tone of our response to emerging threats – incremental, deliberate, focused on exquisite outcomes – is unacceptable and self-defeating. Predictable countermeasures offered up on glacial schedules won't work. We have to surprise our adversary.

Operating autonomously and networked together, a fleet of low-cost, unmanned, expendable platforms such as DARPA's Anti-Submarine Warfare Continuous Trail Unmanned Vessel (ACTUV) might represent a future approach to assembling a maritime force, as an alternative to today's model of fewer high-value assets of great power.



The Disruption

"If you challenge the conventional wisdom, you will find ways to do things much better than they are currently done."

– Michael Lewis, *Moneyball: The Art of Winning an Unfair Game*

Questioning the status quo – whether in baseball or warfighting – carries with it significant risks. The community is heavily invested (monetarily and culturally) in the gear it buys, the heuristics it uses, and the activities or operations it conducts. Questions are inconvenient and carry an implication of disrespect – so the system rejects the heretic and begs off answering the question. This is tolerable when the nation faces no existential threat. Since 1991, the United States has enjoyed a singularly unipolar moment. Our platforms and systems, networks, and plans were built to

overcome a massive Soviet invasion of Europe. Instead, they've been co-opted in the service of combat operations on a lesser scale — perhaps not perfect for counterinsurgency but more than adequate to the task when suitably modified.

The next war may not afford us this luxury. Our traditional approaches to developing, fielding, and using future warfighting capabilities are not up to the task. They don't account for well-resourced and technologically superior adversaries, but they should. Let's take a look at four assumptions that define how we prepare for and fight today, in each of the physical domains in which we operate:

- Achieving air dominance is essential to securing victory in combat. This tenet is so central to the American way of war that it is now nearly impossible to imagine a time when we faced a truly contested air regime. Yet 26,000 U.S. airmen were lost in the European theater during World War II — more than the number of Marines who perished in the Pacific. Our adversaries fully understand how critical control of the skies is to American victory, and they are working to deny us that control. We must ask ourselves if success in combat can be achieved without devoting immense resources to suppression of air defense systems and enemy aircraft, and what technologies would be needed to allow for this. Might it be possible for us to forgo the next round of stealth appliques and focus on those enabling capabilities to permit our air forces to be detected and even tracked by the adversary, but make them simply impossible to shoot down? Rather than air dominance, we would achieve an undeterrable air presence — a disruptive capability that our adversaries will not have prepared for and which can be expected to impose considerable cost to mitigate. One option: Hypersonic weapons, which, even if detected, are difficult to defend against owing to their speed and maneuverability. DARPA's Hypersonic Airbreathing Weapon Concept (HAWC) and Tactical Boost Glide (TBG) programs will demonstrate the technologies for enabling hypersonic cruise missiles and hypersonic glide vehicles, respectively, with flight tests slated to occur over the next two years.
- The aircraft carrier is the only platform that can adequately project power in the maritime domain and must therefore be protected at all cost. I've already discussed how Japan continued building powerful battleships long after the battleship's deficiencies were understood and exploited by the United States and its carrier strike groups. To believe that the carrier can be defended indefinitely, in the face of capabilities such as China's DF-21 anti-ship ballistic missile, which could be fielded in quantity from long range, is to ignore reality. The carrier is an extremely high-value asset and will be targeted. Like the obsolete battleship, they will eventually be forced to retreat to safe distances or risk destruction. This requires us to posit a future without the power-projection capacity of the supercarrier — and to ask how we will win a maritime engagement in its absence. Perhaps proliferation and disaggregation of the carrier's air wing into many smaller, easier-to-produce platforms is an alternative to ever-more-exquisite defensive schemes. It behooves us to ask what technologies are required to achieve this end, and to begin working on them now. DARPA's Anti-Submarine Warfare (ASW) Continuous Trail Unmanned Vessel (ACTUV) may point the way to

one possible future: Small, autonomous naval vessels that can operate without human intervention over ranges of thousands of miles, linked together, and operating as a distributed combat force, with no precious “center.”

- Ground combat is, at its essence, close combat — you must meet the enemy symmetrically, gun against gun, tank against tank, to achieve victory. No other domain is as defined by its penchant for “stand in” combat as the land domain. Our troops sit heroically atop armored personnel carriers to draw fire from snipers hiding on urban rooftops. They drive over improvised explosive devices (IEDs) to ferry materiel and man checkpoints to draw out the suicide bomber. They interface intimately with an enemy that wishes to deny us our inherent technological advantages by forcing face-to-face engagements. Our goal should be to refuse the gambit offered by the adversary and work to achieve standoff effects in the same manner that we take for granted in the other physical domains. We must buffer our personnel on the ground with ultra-long-range munitions and networks of autonomous agents that detect and discern (and defeat) the enemy long before physical contact is ever made. DARPA's Urban Reconnaissance through Supervised Autonomy (URSA) program — a necessary and responsible first step — aims to use unmanned agents to autonomously detect hostile forces and establish positive identification of combatants in the complex, uncertain, and unsafe urban battlefield.
- The space domain must evolve in a manner analogous to that of the air and maritime domains, with combat functions emerging to defend our satellites and degrade the space capability of the enemy. Ringing our exquisite spacecraft with defensive and offensive capabilities is a highly predictable countermeasure that will force the United States to incur substantial additional costs — in addition to the large sums of money the country already spends on the assets we operate today. The stark similarity between the space and sea domains becomes clear: You can accept the inherent advantages of the high-value asset — powerful, concentrated warfighting capability — and be forced to defend it — or you can opt out and disrupt the enemy calculus, which assumes we will adhere to sunk cost and “the way it's always been done.” One pathway of such disruption is to proliferate and disaggregate capability and present our adversaries with an entirely new problem set — hundreds or thousands of commodity platforms, easily replenished and individually expendable. The value of this approach? Deterrence. Why start shooting if your shots cause minimal damage and your activities are instantly attributable? Recreating our space order of battle without high-value assets (and their associated cost and deployment schedule problems) will allow us to tighten our own innovation cycle and give our enemy pause. DARPA's Blackjack program will leverage low-cost commercial satellites and the emerging space-based internet to provide ubiquitous sensor coverage and change the space community's incentive structure for risk acceptance. Our Launch Challenge will demonstrate an on-demand small launch capability, allowing us to constitute, reconstitute, and technologically refresh our capabilities and respond to (if not pre-empt altogether) emerging threats.



DARPA's Blackjack program might provide the means to transition many of today's national security space missions from big, expensive satellites to a networked constellation consisting of many inexpensive, maneuverable satellites in low Earth orbit.

It has been a very long time since the Department of Defense and its industrial base were forced to re-examine the fundamental assumptions that underlie the American way of war. This is understandable: Without an event as momentous and threatening as Pearl Harbor, one that forces a reassessment of everything we thought we understood about the world, it is difficult to make a case for a change that requires wholesale revamping of acquisition, requirements, operational planning, and tactics. While some might argue that 9/11 represented such an event, the threat was never existential, and it couldn't force the level of soul-searching that dragged our nation into World War II.

For the past three decades, we've chosen to sharpen the saw, doubling down on the approaches we've grown comfortable with — stealth and precision weaponry in the air domain, concentration of power in high-value assets in the sea and space domains, and increasing the proficiency of our small units on land, mostly against irregulars and insurgents. This tendency to stagnation in thinking and action is the real challenge we're facing in all of the warfighting domains: We are so intent on improving what we have in hand already that we prevent ourselves from fully seeing and addressing the actual and rapidly evolving threat we face. And the real threat is an ad-

versary that can anticipate our moves and countermoves and act inside our decision cycles, be they tactical or strategic. Only disruptive moves — moves that the enemy could not have expected based on past behavior — will upset the enemy's plans and make it expensive for them to act. If the disruption is sufficiently extreme, the cost of building or using the countermeasure will be too high and we will have achieved our goal: deterrence.

Our senior leaders are cognizant of the danger and are speaking out. Dr. Michael Griffin, the Defense Department's Under Secretary for Research and Engineering, has stated that American military superiority "was bought for us by our forebears." It is not a birthright. Gen. John Hyten has noted that our satellite systems are now "big, fat, juicy targets." He goes on to say, "I watch what our adversaries do. I see them moving quickly into the space domain. They are moving very fast, and I see our country not moving fast, and that causes me concern." We have to re-assess our decades-old strategy, founded on a belief in a perpetual Pax Americana and U.S. technological superiority, and make needed changes.

The recent reported loss of the Zuma spacecraft is a painful reminder that the tried-and-true path is faltering. Disruptive change is essential if we hope to confound (and deter) potential adversaries. Slow and deliberate must give way to rapid and risk-accepting. Complex, exquisite systems must be shelved in favor of commodities, and "good enough" platforms and networks. To disrupt, to sow chaos in our adversaries' plans, requires us to embrace that selfsame chaos and disruption within our own culture — and we can't wait any longer.

The Biological Technologies Office develops new tools to support warfighter health, fitness, and training so that the force can perform at peak capacity no matter the mission. Technologies like embedded hydrogel sensors continuously measure body chemistry and give individuals and commanders real-time visibility into squad readiness.



SECURITY AND SURPRISE AT BIOLOGICAL SCALES

The evolution of defense technology

By Justin Sanchez and Jacob Jordan

The power and challenge of biology as the basis of technology reside in its adaptability and resilience, and from those attributes spring the seeds of surprise. DARPA's mission is to master surprise, and so the agency has embraced biotechnology and focused its investments on mitigating threats to human health and global stability, improving military training and readiness, and rethinking current approaches to traditional defense missions. In this seventh decade of DARPA-based innovation, the 4-year-old Biological Technologies Office (BTO) is set to transform biology from a niche specialization within the Department of Defense (DOD) into a fundamental source of high-impact capabilities, among them force protection, sensing, command, control, and communications.

DARPA is not alone in this game. Across the board, the growing affordability and accessibility of tools to modify biological systems means that other nations, and even individuals, have the resources to engineer systems and organisms to have altered or new functions. Consider the dramatic decrease in the cost of sequencing a genome, paired with widely available tools for editing genes. Now, virtually any person anywhere might experiment with genetic modifications. Much of that research has positive ends, but all biotechnologies have the potential to be dual-use. And there is always the risk of unintended consequences.

The application of biotechnologies by an adversary is an area where the United States could be most surprised as a nation, but it is also a source of great potential, where the United States could develop a host of new surprises of its own.

Incorporating biotechnologies into DOD operations requires a deliberate effort, beginning with familiarizing potential end users with emerging technologies, demonstrating capabilities in contexts relevant to DOD missions, and creating opportunities to integrate those capabilities into DOD organizations. Part of DARPA's challenge is breaking through subtle institutional preconceptions about what is and is not possible. A generation of military leaders raised on a diet of science fiction may be inclined to perpetually relegate certain biotechnologies to a vague future. But, in fact, BTO's talented roster of program managers and research performers not only operates at the frontiers of science, but thinks beyond them, and the seeds of those tools and capabilities that might sound fantastical already exist today.

The emerging importance of biological technologies for national security is an important decision-making driver for Dr. Steven Walker, Director of DARPA. "Biology percolates throughout every aspect of the Department of Defense. The people who serve, the missions they carry out, and the technologies that facilitate their performance all exist within an organic world, and we're building the tools to engage with that world in new ways," Walker said. "As the DOD accelerates its modernization push, the technologies coming out of BTO could transform how the services carry out their missions, and DARPA is committed to facilitating the integration of biological capabilities into the total force."

Since its establishment in 2014, BTO has optimized its approach, operations, and research portfolio to create opportunities to transition defense-relevant capabilities to military end users and, in some cases, to society at large. That strategy involves considering broad classes of challenges rather than individual threats, expanding the scope of biotechnology to missions that have not traditionally had an overt biological component, and reducing the various risks involved with embracing new technologies.

Reinventing Approaches to Bio Threats

The United States confronts a potential national security threat each time a new outbreak of infectious disease occurs anywhere in the world. The nationwide spread of H1N1 flu in 2009, the spread of Middle East Respiratory Syndrome in Indiana in 2013, and the presence of Ebola in Texas in 2014 are recent reminders of this reality. A single viral particle has the power to replicate billions of times over, introducing communicable disease to large cities, overwhelming hospital systems, inciting mass terror, and dramatically reducing the readiness of our military, governmental functions, and infrastructure elements, all of which require healthy personnel to run them. And in this connected world, a virus can circle the planet on jetliners in less than 36 hours.



The threat to force health and readiness posed by infectious disease is persistent and evolving. BTO is developing a range of countermeasures to natural and engineered threats, such as technologies to predict viral evolution and spread, contain diseases within animal reservoirs, disrupt disease vectors' ability to transmit viruses, and control or remediate the effects of gene editors.

Risks posed by naturally occurring outbreaks are only one concern. Recent advances in synthetic biology and gene editing have reduced long-standing barriers to entry for nefarious actors wishing to develop engineered biothreats. Doing so no longer requires a detailed understanding of molecular biology, wet lab skills, or access to live pathogens. The widespread availability and reduced cost of DNA sequencing, and more importantly DNA synthesis, have made it possible to reconstitute eradicated pathogens—U.S. and Canadian researchers reported in 2017 that they had synthesized horsepox virus—and to create new ones *de novo*. Advances in gene editing also open the way for direct modifications to the human genome and the construction of "gene drives" that override traditional inheritance in offspring in ways that can propagate specific traits and biological functions throughout an entire population.

There is currently a mismatch between the rapidity at which bio threats can emerge and proliferate and the response time for developing and deploying effective countermeasures. Traditional biodefense technologies are designed primarily to counter known pathogens, and even still, they require long lead times to develop responses. If a successful countermeasure can be created, it must then be mass-produced and stockpiled in preparation for a large-scale emergency. This cycle is time-consuming enough for known and predictable threats. It can be tragically slow when it comes to unfamiliar pathogens, which can spread at pandemic scales before researchers can devise countermeasures.



The Advanced Plant Technologies program envisions plants as discreet, self-sustaining sensors capable of detecting signals of interest and reporting them via remotely monitored, programmed responses to environmental stimuli.

Cognizant of this time-sensitive dynamic, DARPA structures its medical countermeasures R&D to pursue generalizable, threat-agnostic technology platforms that can be tailored on demand, rather than building one-off solutions to address individual threats as they emerge. The flu vaccine must be reformulated every year because the flu virus mutates, so the medical community is always playing catch-up. Imagine if, in preparation for the next flu season — or worse, the next outbreak of Ebola or Zika — an existing therapeutic tool could be quickly and easily customized to fight the specific pathogen, stopping the outbreak in its tracks. This type of technological firebreak could eliminate the national security risks posed by infectious disease.

BTO's Pandemic Prevention Platform (P3) program encapsulates one such approach, with the ambitious goal of halting a pandemic in 60 days or less. It builds on pioneering DARPA research into gene-encoded therapeutics and is designed to achieve higher therapeutic potency and efficiency compared to traditional countermeasures by using the human body as a bioreactor.

When individuals are exposed to pathogens, their bodies respond by producing protective antibodies and other immunological responses. Some of these are strong and highly protective, while others are weak and nearly ineffective. P3 aims to deliver the genetic code for the most potent antibodies directly into a person in need of protection so they can produce their own defense agents *in situ*. This approach circumvents slower, traditional, vaccine-based approaches or biotherapeutic production protocols

involving chicken eggs, animals, or plants. Used in combination with other high-throughput screening methods for identifying anti-virus therapeutics, all that is needed to combat a known, emerging, or engineered virus is the genetic code (its sequence of A, T, G, and C nucleotides) of an effective antibody. By translating the genetic code of that antibody into a therapeutic manufactured directly by cells within the body, researchers can dramatically reduce the time from virus identification to creation of an effective and approved countermeasure, and ultimately open the way for the timely distribution of the countermeasures within the time window before an outbreak can rage into a pandemic.

The P3 program manager, Col. Matt Hepburn, M.D., is now working to translate an initial proof of concept into a deployed capability by overcoming the remaining bottlenecks in the process. "We're working to demonstrate an ability to quickly grow the quantities of virus needed to test and evaluate therapies, identify antibodies, and increase their potency within the first weeks of an outbreak, and scale methods for delivering treatments into humans," Hepburn said.

Moving from Discovery to Real-world Capabilities

Key to BTO's vision is that biotechnology capabilities have a role in the DOD well beyond human health and biowarfare defense. Central to this vision is the increasing interactions among the fields of gene editing, materials science, artificial intelligence, and other specializations.

Long-standing problems of national security are growing in complexity and beginning to push up against the limits of traditional engineering in a variety of operational contexts. Whether on land or at sea, the DOD deploys people and equipment to maintain military overmatch across vast and diverse spaces. And the people carrying out these missions must now



interact more frequently with autonomous or semi-autonomous systems that work at speeds faster than humans can process without assistance.

Biotechnology approaches are beginning to present an intriguing alternative to solutions that were previously based solely in engineering, because unlike electrical and mechanical systems, biological systems are persistent, adaptable, autonomous, power-efficient, and discreet when deployed in their native environments. BTO's team of program managers and leadership – a cadre of life scientists and engineers – has set out to evolve national security capabilities by embracing those unique properties.

Dr. Blake Bextine, an entomologist by training who joined DARPA in 2016, understood that plants could complement traditional military sensors in complex environments to detect signals of interest. “Plants are highly attuned to their surroundings and naturally manifest physiological responses to stimuli such as light and temperature, so DARPA asked if it's possible to reprogram those detection and reporting capabilities for stimuli like pathogens and threat agents,” said Bextine.

He wants to apply the latest advances in genome-editing technologies to develop persistent plant sensors that not only detect and report on signals of interest – that is, chemical, biological, radiological, nuclear, and explosive (CBRNE) threats – but that are also robust enough to survive in place and reproduce. By also engineering into these systems response mechanisms that can be remotely monitored using existing ground-, air-, or space-based hardware, this new plant-based technology ultimately could reduce the risk to troops now tasked with maintaining traditional hardware-based sensors in contested areas.

Many of the researchers working with BTO are not used to thinking about their work in the context of national security. “DARPA has awakened the agriculture and horticulture communities to opportunities that didn't used

The Persistent Aquatic Living Sensors program promises to bring new resources to the challenge of detecting and characterizing underwater vehicles in strategic waters.

to exist,” Bextine said. “Now that we've made these connections to the defense and intelligence worlds, I'm excited to see what else biologists can contribute.”


BTO is putting similar concepts to work in the maritime domain.

“Projecting power and maintaining awareness in the oceans has always been critical for national security, but the vast scale and harsh environment is an endless source of challenges,” said Dr. Lori Adornato, a chemist and ocean scientist who joined DARPA in 2017. “Biology presents new opportunities for DOD to achieve better results with fewer resources.”

Adornato launched the Persistent Aquatic Living Sensors program this year with the vision of tapping into the natural sensing capabilities of marine organisms to help detect, track, and monitor adversary submarines, unmanned underwater vehicles, and divers. In true DARPA form, her concept takes signals once considered background noise and transforms them into a new sensing modality.

“Sonar operators on submarines often hear shrimp and other organisms become active when vessels pass by, and they've worked for years to minimize those signals. DARPA asked, ‘Well, can't we use those signals to our advantage?’,” Adornato explained.

In other domains where DOD operates, it is human biology that could come into play. Service members already control complex military systems in stressful circumstances, and, in the future, the pace and complexity



As the pace and complexity of modern combat increases, warfighters require new tools to help them stay meaningfully engaged. BTO continues to demonstrate new capabilities for intuitively interfacing with complex systems. For example, emerging technologies such as noninvasive neural interfaces could, in the future, allow pilots to control groups of unmanned aircraft by delivering high-level commands and maintaining oversight while the systems adjust to rapidly changing conditions to complete the mission.

of missions will only increase as more semi-autonomous systems come online. BTO is developing technologies that can help personnel operate more collaboratively and intuitively with these systems and make informed decisions at tactically relevant speeds.

Building on a rich history of DARPA-funded breakthroughs, performers supporting BTO's brain research portfolio are delivering increasingly sophisticated tools and methods to allow a person's brain to communicate with external devices [see "Taking Neurotechnology into New Territory"]. DARPA has demonstrated direct neural control of a robotic limb, restoration of touch sensation to a paralyzed patient, improvement in short-term memory, and acceleration of training times.

These previews of the power of neurotechnology have mostly involved surgically implanted devices, and so far, they have been directed at restoring functionality to individuals who have suffered the effects of injury or illness. It will take additional breakthroughs to move neurotechnology into a form factor that makes it widely applicable to DOD operations, including for tactical mission execution and strategic planning.

Dr. Al Emondi, a neuroscientist and electrical engineer, joined BTO last year to take on this challenge. "Our No.1 goal is to develop communications links to the brain that do not require surgery," Emondi said. "A high-performance, noninvasive neural interface would open up possibilities such as immersive training, new forms of interaction with AI systems, improved situational awareness and intelligence analysis, and distributed task management with machines to speed tactical decision-making and free up cognitive function for strategic planning. It's a potentially foundational technology for the next generation of DOD systems."

Building Partnerships, Controlling Risk, and Maximizing Benefit

For advanced biotechnology systems like neural interfaces and gene editors, the path from demonstrating a capability to the end user is neither simple nor direct. BTO structures its programs to systematically reduce various forms of risk to ease that transition. From facilitating regulatory review to attracting additional investment to initiating dialogue on potential legal and ethical concerns, BTO program managers shepherd their technologies to a point where they can thrive after DARPA's involvement ends.

Ethically, DARPA's leadership and program managers understand that any technology can be used for good or ill. Whether it's nuclear energy or the internet, a capability developed for one purpose may easily be put to another in a way that ultimately affects society as a whole. DARPA expects that biological technologies especially will permeate many aspects of both the national security and commercial worlds. In anticipation of future use of the tools and capabilities it develops, BTO emphasizes transparency,



Programs such as Biological Control, Biostasis, BRICS, ElectRx, HAPTIX, and Neural Engineering System Design lay the foundations for powerful new capabilities to preserve troop health, facilitate recovery, and restore function.

data-sharing, and publication of results, and initiates critical discussions concerning the ethical, legal, and social implications (ELSI) of active and proposed research.

At all stages of BTO programs, including initial development of program concepts, program managers consult with independent, unpaid ELSI advisors to assess potential concerns, provide feedback on proposed research plans and results as they become available, and participate in program meetings with researchers. In addition, these advisors help to share DARPA ideas with the larger research community so that others may join the dialogue on technologies that could find broad application.

Operating at the frontiers of science and technology means that BTO and its research partners are among the first to confront the implications of new capabilities, and often there is no clear guidance on how to address concerns. Because it's not an option to stay away from those spaces, as BTO explores this *terra incognita*, it tries to lead by example regarding how to incorporate ethical norms into the research.

Practically, BTO forms interagency partnerships early in the program life cycle to ensure that technology development follows a path that is relevant to regulators and potential end users. Organizations such as the Environmental Protection Agency, the Food and Drug Administration, and the Department of Agriculture lend expertise on drugs, medical devices, and technologies involving plants and animals. Those interactions provide valuable guidance on how to facilitate regulatory review, while the Defense Threat Reduction Agency and the Biomedical Advanced Research and Development Authority provide tangible pathways for moving BTO-developed technologies into operational test and evaluation.

Beyond these intragovernmental relationships, BTO needs the help of companies, universities, venture capitalists, and entrepreneurs to refine technologies and move innovations to service members, patients, and consumers. Because many biotechnologies are so new, despite promising early results, it may still be several years before a mature product is ready for direct integration into DOD operations. In the meantime, the timing could be ripe for a commercial venture to invest additional research and development funds into a consumer application. BTO aims to engage with potential commercial transition partners at the earliest possible opportunities in the technology development cycle to inform early research and increase the chances of creating a product that can move beyond the lab bench.

Looking Over the Horizon

Future BTO programs and focus areas will continue to more closely integrate biotechnology and national security. Through its exploratory investments and technology demonstrations, BTO seeks to establish new and more powerful technical foundations for a wide array of critical DOD missions and challenges. BTO cannot succeed alone, though. The collaborations that have been formed to date will only become more important as engineered biology reveals its full potential, and BTO looks forward to engaging the military services, integrating with the defense-industrial complex and commercial industry, and deepening its connections with non-DOD agencies to scale up the use of biological solutions.

DARPA has opened the door to a vast new range of opportunities for the Department of Defense. It needs to seize the moment, consider missions in the new context of what is becoming possible with biology, and work with collaborators to develop capabilities that used to exist only in science fiction.

After all, biology adapts to thrive, and any system that evolves can become a source of surprise. DARPA and the DOD finally have the tools to work with biology and to harness that potential.

AUTONOMOUS TECHNOLOGY COMETH

The advent of machines that act and decide will benefit from the lessons humanity has learned from itself

By J.R. Wilson

Autonomy. Until the 21st century, it almost always referred to humans, either through individual freedoms or regional/national self-governance, sovereignty, independence, and freedom.

Today, however, autonomy is more likely to raise images of robots, self-driving vehicles, unmanned aerial vehicles (UAVs), and, more recently, their unmanned ground (UGV), water surface (USV), underwater (UUV), and even space counterparts. For DARPA, autonomy is a context-specific descriptor that can vary by domain, according to Timothy Chung, a program manager in the agency's Tactical Technology Office (TTO), where he focuses on autonomous systems.

"For terrestrial robot systems, working in structured environments, even dynamic ones like streets, we're starting to see a maturation, because we have the compute power, different sensing modalities, GPS, and a lot of data to help inform the autonomous behaviors. That also applies to the air, which, despite being 3-D, has fewer obstacles," he said. With the ocean, there are environmental challenges, whether surface or subsurface, for getting air- or terrestrial-grade autonomy into those environments.

"Depending on how you define it, autonomy already is out there," Chung continued. "Think about adaptive cruise control in cars or self-parking vehicles. A smartphone might be a conduit for some autonomy to exist, such as doing machine-learning off-board, then using it for facial recognition or augmented reality, for example, where you might interact with digital characters in a virtual world."

Building on early and dramatic successes of UAV technologies in Operation Desert Storm and their deployment as essential elements of U.S. and coalition operations in Iraq and Afghanistan a decade later, the agency has continued to push UAV technologies forward. At the same time, the commercial and industrial R&D sectors have been building up their own momentum in the UAV space to generate the current worldwide explosion of personal and commercial UAVs.

Pervasive as unmanned systems have become, however, the vast majority of them remain connected to a human being with a joystick, even if that human is safely operating the system from continents away. This connection of one unmanned system to one (and usually more than one) human operator places constraints upon the system, limiting the effectiveness of unmanned systems and increasing operational costs. "The No. 1 manning problem in our Air Force is manning our unmanned platforms," Gen. Philip

M. Breedlove said in November 2011, when he was leading the Air Force's operations, plans, and requirements directorate. The Air Force estimated at the time that each operational MQ-1 Predator required a crew of about 168 personnel, the MQ-9 Reaper approximately 180, and the RQ-4 Global Hawk some 300 people.

One of DARPA's newest programs — the Collaborative Operations in Denied Environment (CODE) program — is designed to extend the capability of the U.S. military's existing UAVs to conduct long-distance engagements of highly mobile ground and maritime targets in contested, denied, and constantly changing battlespaces. If the program proves successful, the resulting collaborative autonomy technology could enable groups of unmanned aircraft to work together under a single person's supervision rather than through the current system of continuous control by a pilot and sensor operator, each supported by numerous analysts. UAS could navigate to their destinations, find, track, and identify targets, and then present recommendations for coordinated actions to a mission supervisor, who would approve or disapprove such team actions and direct any mission changes from hundreds or even thousands of miles away.

Meanwhile, DARPA's OFFensive Swarm-Enabled Tactics (OFFSET) program focuses on the future ability of small infantry units to deploy swarms of 250 or more mini- or micro-UAVs and/or UGVs to help accomplish diverse missions in complex urban environments. Contracts to deliver on the program's Phase 1 challenges — to design, develop, and demonstrate a swarm system architecture to advance the innovation, interaction, and integration of novel swarm tactics — were awarded in February 2018 to Raytheon BBN and Northrop Grumman Mission Systems.

DARPA also is moving forward with a larger portfolio of autonomy projects and programs:

- The Assured Autonomy Program, a new research effort that builds on recent breakthroughs in autonomous cyber systems and formal methods (reliable and telling mathematical models of complex systems), aims to advance how computing systems can learn and evolve to better



In January 2018, DARPA transitioned the Sea Hunter, a vessel that emerged from its Anti-Submarine Warfare (ASW) Continuous Trail Unmanned Vessel (ACTUV) program, to the Office of Naval Research (ONR).

manage variations in the environment and enhance the predictability of autonomous systems, such as driverless vehicles and UAVs.

- The Gremlins program aims to build UAVs that can carry 60-pound payloads up to 300 nautical miles and be launched from and recovered by fighter aircraft, bombers, or C-130 manned transports. The roles of these versatile UAVs could include intelligence, surveillance, and reconnaissance, signals intelligence, and electronic warfare.
- The Fast Lightweight Autonomy (FLA) program, now in Phase 2, aims to explore nontraditional machine-vision and autonomy methods that could empower a new class of algorithms for high-speed navigation in cluttered environments such as streets or inside buildings. FLA aims to develop and demonstrate the capability for autonomous UAVs small enough to fit through windows and fly at speeds up to 45 mph without human control or GPS guidance.
- The Experimental Spaceplane program, with Boeing, aims to build and fly an entirely new class of fully reusable hypersonic craft (Mach

5 or higher) that can be launched into low-Earth orbit 10 times in 10 days; reduce launch costs by 90 percent; and deploy a satellite of up to 3,000 pounds. The goal of the program is to provide the nation with an unprecedented ability to quickly recover from a catastrophic loss of critical military or commercial satellites. Automated flight-termination and other technologies for autonomous flight and operations, including some developed by DARPA's Airborne Launch Assist Space Access (ALASA) program, are part of the program. The first test flights are scheduled for 2020.

"In essence, we're still conducting applications that people can or currently do, and seeing how autonomous systems can supplant the human actors to save lives, mitigate risk, or increase efficiency. There also may be breakthroughs enabling wholly new capabilities, perhaps some not yet envisioned, that humans do not currently do," Chung explained. "Whether subjective or not, we have a baseline of human performance, which tends to be more forgiving for humans than it is for machines. We hold our autonomous systems to a higher bar of performance and, as a result, we need to understand the type of reliability, predictability, and trustworthiness they can have. These are the same traits we would want from a human teammate."

The technical tasks for achieving this kind of autonomy, Chung said, involve advances in human-robot interfaces, techniques of coding the



An artist's conception of manned and unmanned aircraft working cooperatively as envisioned in the Collaborative Operations in Denied Environment (CODE) program. CODE intends to focus on developing and demonstrating improvements in collaborative autonomy: the capability for groups of UAS to work together under a single human commander's supervision.

autonomy in the software, and improving reliability throughout the system, from its chips to its mobility and other performance parameters.

"On a factory floor, we can pretty much trust what robots and automated machinery do," Chung continued. "This is not the case for something like unpiloted transport of passengers or deliveries, where we need software methods to determine and certify that the system is reliable. Part of the solution will reside in software tools for realistic simulation studies of autonomy.

Said Chung: "Those simulation tools will dramatically change how we feel about those autonomous systems in use, either alongside or remote from humans going forward."

Historically, DARPA often has conducted research into technologies that neither are understood nor yet in demand, revisiting those same technology categories periodically as new program managers come aboard with fresh viewpoints and new developments to help move a "fringe" technology closer to being useful. Although the agency continues to pursue such research, the speed and global extent with which technology research and deployment is advancing has spurred DARPA to redouble its own signature knack for pushing the envelope without being deterred by short-term setbacks.

The agency's most coveted sign of success is when it transitions the capabilities its R&D projects support into downstream development, testing, and operational use pathways.

On Jan. 30, 2018, for example, DARPA officially concluded about three years of collaborative development with the Navy by transitioning the Sea Hunter, a demonstration vessel that emerged from its Anti-Submarine Warfare (ASW) Continuous Trail Unmanned Vessel (ACTUV) program, to

the Office of Naval Research (ONR). ONR now is continuing development of the revolutionary prototype vehicle, the first of what ultimately could become an entirely new class of ocean-going vessel – the Medium Displacement Unmanned Surface Vehicle (MDUSV) – that traverses thousands of miles of open seas for months at a time, without a single human crewmember aboard.

"ACTUV's move from DARPA to ONR marks a significant milestone in developing large-scale USV technology and autonomy capabilities," Alexander Walan, a TTO program manager, said at the time in an agency release. "Our collaboration with ONR has brought closer to reality a future fleet in which both manned warships and capable large unmanned vessels complement each other to accomplish diverse, evolving missions."

"ACTUV represents a new vision of naval surface warfare that trades small numbers of very capable, high-value assets for large numbers of commoditized, simpler platforms that are more capable in the aggregate," added TTO Director Fred Kennedy. "The U.S. military has talked about the strategic importance of replacing 'king' and 'queen' pieces on the maritime chessboard with lots of 'pawns,' and ACTUV is a first step toward doing exactly that."

"Sea Hunter serves as a promising test bed for a new kind of mentality in the Navy about what it would take to have an autonomous, long-duration vehicle. But we think commercial fleets also can benefit from those same capabilities," Chung said. "We see the law enforcement community adopt bomb disposal robots that came out of DARPA for the military well over a decade ago."

Most likely far more consequential will be the self-driving car technologies that evolved from DARPA's Grand and Urban Challenges, Chung said, adding that "cellphone assistants, such as Siri on the iPhone, evolved from PAL [Personal Assistant that Learns], a language assistant program that came out of DARPA roughly four years before Siri was launched."

Autonomy manifests itself in many ways and holds promise as an enabling framework for multiple DARPA projects. For example, robotic arms developed to service and repair orbital platforms in space also can be used for such earthbound roles as inspecting oil and gas pipelines.



ABOVE: The Fast Lightweight Autonomy (FLA) program is exploring nontraditional machine-vision and autonomy methods to empower high-speed navigation in cluttered environments for small, autonomous UAVs. RIGHT: The Offensive Swarm-Enabled Tactics (OFFSET) program focuses on a future capability for small infantry units to enable them to deploy and operate 250 or more mini- or micro-UAVs and UGVs in complex urban environments.

|||||

In airborne settings, having a broad understanding of how autonomous UAVs can expand surveillance, communications, and other facets of the battlespace opens the way to better interactions between human pilots and UAVs, perhaps leading to UAVs acting as a combat pilot's "wingman." That would reduce the number of manned aircraft needed on some missions and thereby remove at least some human pilots from harm's way.

Ushering autonomous UAVs into the battlespace is "not just a matter of self-reporting to the human pilots what it is doing, like a sophisticated transponder, but being able to make decision recommendations, alleviating the [pilot's or military air traffic controller's] responsibilities, leaving him to be a tactician or battlespace commander more effectively. That touches on a lot of what DARPA is doing," Chung said. "On the ground, you see . . . Ground X-Vehicle Technologies, which are more about crew augmentation, such as instrumenting a windowless ground vehicle so the crew inside will have a better situational awareness of their environment.

At sea, being able to maneuver, especially underwater, remains an active area for the development of autonomy options. The Cross-Domain Maritime Surveillance and Targeting [CDMaST] program, for example, is trying





The DARPA Subterranean Challenge calls on participants to deliver autonomous vehicles that can perform rapid navigation, search, and mapping of underground environments.

to better monitor and protect the waters, and part of the vision includes autonomous systems. These could provide cuing to other systems – manned or unmanned – to defend and protect those areas. So a distributed network of sensors and autonomous vehicles could help shape the waterscape for maritime warfare.

DARPA's OFFSET program is all about leveraging tactical opportunities that could come with coordinating swarms of air and ground robots – up to 250 of them – in urban areas. Missions could include aiding a small unit of warfighters charged with securing a bridge, or more crucially, aiding in the historically challenging and costly task of clearing a multi-floor building of the enemy. Important to this technology vision is understanding where such swarms excel and where they don't at all.

"The complexity, command, and design space of 250 physically disjointed robots, some air, some ground, with different speeds and capabilities, requires having to break people's familiar and ingrained mentality of how we interact with smaller numbers of vehicles," Chung said. "Think of the number of different tactics we could employ with larger numbers of vehicles. Those 250-unit swarms are not that far away."

The subterranean domain, perhaps the most unsung among all of the warfighting domains, was important in both Vietnam and Iraq. In the spirit of previous DARPA challenges, the agency has issued a Subterranean (SubT) Challenge to both accelerate an entire technology category and catalyze the formation of a research community to do just that. The

capabilities the SubT Challenge calls on participants to deliver include rapidly navigating, searching, and mapping underground environments. In addition to its military applications, the potential payoffs of such technology range from finding trapped miners in a mine collapse to diagnosing specific problems in damaged storm drain systems and utility tunnels to mapping out natural caves.

"Caves are unpredictable and unstructured, but provide a natural test bed for assessing how well the systems we hope the participants deliver meet the various capabilities at the heart of the Challenge," Chung said. "We find, time and again, that the geology and configuration of rock formations, and the location of water tables, can change only a short distance away as the environment changes. We want to inspire competitors to address all three domains [human-made tunnels, underground urban spaces, and natural cave networks] and come up with system solutions – a federated solution of platforms – that can work across unpredictable environments."

Organizers of the SubT Challenge hope the effort will deliver new and powerful technology options in a wide range of arenas, among them mine safety, infrastructure inspections, GPS-free navigation, and communications – often with jamming and out of line of sight. "Many of our systems, even UAVs, tend to be thought of in a two-dimensional environment, where subterranean contexts may have an inclined shaft or a deep, dark hole," said Chung. "Competitors in this Challenge will have to think through how to deal with such encounters."

It will be years before any seeds of technology that emerge from the SubT Challenge can grow into actual new features of the technosphere. Perhaps the autonomous technology most frequently in the news these days is self-driving cars – both their amazing feats on the road and their accidents, some of them tragic. But while this mix of success with some failure has raised questions about whether the technology is sufficiently sophisticated



and reliable to be unleashed on millions of public streets and roads, the technology's track record has convinced Department of Defense (DOD) decision-makers that it's ready for primetime in military applications.

"The questions raised are relevant; these are issues DOD has and will continue to take seriously," Chung said, adding that the DOD has issued directive 3000.09, providing guidance regarding what decisions autonomous vehicles can and cannot make, including ones regarding the use of potentially lethal onboard weapons.

"We have to hold our systems to a high level of scrutiny because just as with driverless cars, lives will be on the line. The same type of reliability we care about in the military is of interest to the commercial market," he added. There also is a confluence of interest between the military and commercial sectors when it comes to interoperability and compatibility. "Both industry and government can step up to help that in the future," Chung said. "The commercial sector has to rely on market demand, but the stars are aligning between the interests of DOD and the needs for greater autonomy and robotic systems on the commercial side, which is a great development."

At the heart of DARPA's autonomy portfolio are various levels of artificial intelligence (AI) — from the level of basic machine learning to that of adaptive systems to a level that is "cognizant" and can continuously learn from experience and from data it acquires on the fly. A current theme of the overall DARPA research portfolio centers on moving up the AI levels and applying these ever more capable AI systems to many technologies, including autonomy.

DARPA's Experimental Spaceplane program to build a new class of hypersonic reusable spacecraft will leverage technologies for autonomous flight and operations.



"Machine learning is still relevant," Chung noted, offering as an example the need of a threat detection system to tell the difference between a person holding a broom and an AK-47. "Adaptive AI is closing the loop around decisions associated with basic machine learning to outputs where autonomous systems will thrive," he explained. "Moving from adaptive to cognitive AI, you get to a framework of not only understanding the environment, but being able to explain and think through scenarios. We're accelerating our efforts toward that."

Chung says he is excited to do his part in accelerating the rate of technology development and maturation in robotics, AI, and autonomy. Driverless cars might be one of the first ubiquitous technologies to arrive, but as he sees it, we are going to witness many more such arrivals of AI-powered robotics and autonomy.

"Human beings are intelligent and autonomous, and we know very well about how many different ways things go right and wrong with people," Chung said. "So we can learn from ourselves what we need to keep an eye on as we usher in artificial forms of intelligence and autonomy."

FIGHTING IN MEGACITIES

DARPA takes on the challenge of warfare in expanding urban settings

By Stew Magnuson

In a high-rise building, cyber attackers are wreaking havoc on the U.S. economy and already cost lives by taking down a power grid. The Pentagon has to take action.

The problem is that the high-rise is located in the middle of a sprawling Middle Eastern city of 12 million inhabitants. It would be easy to send cruise missiles to destroy the building, but most floors are occupied by innocents, including dozens of children.

This is a scenario military thinkers are increasingly worried about: fighting in urban environments, particularly the so-called megacities of the future, defined as having 10 million or more residents.

DARPA is thinking ahead, and has kicked off a series of programs to help warfighters operate in these complex conditions.

“As more populations across the world move to larger and larger cities, we need to understand the three-dimensionality of cities and how to operate in those very crowded, very three-dimensional spaces,” DARPA Director Dr. Steven Walker told reporters recently. “That’s going to become more important in the future.”

Dr. Thomas J. Burns, former director of DARPA’s Strategic Technology Office (STO), said that the scenarios that scare him most are ones involving city-dwelling terrorist cells – whether they are state, non-state actors, or their proxies – which use cyber tools to bring down cities around the world.

“If that becomes a reality, we are going to have to find ways to get in there, find them, and dig them out,” Burns said. “The United States military does not have a doctrine per se that they can follow to conduct warfare in those cities,” he added.

But it will soon need one.

Jean-Charles Ledé, a program manager at DARPA’s Tactical Technology Office (TTO), said: “War is a human affair. You fight where people are. The likelihood of us having to fight in the urban environment is increasing.”

Megacities of the Future

People move to cities to pursue opportunities. While 54 percent of the world’s population lives in urban areas, cities already account for more than 80 percent of global economic activity (as assessed by gross domestic product measures), according to the World Bank. Those numbers will only increase. The number of megacities is expected to rise from 28 in 2014 to 41 by 2030, and by 2050, 66 percent of the world’s population is expected

to live in urban areas, according to a recent report by the United Nations’ Department of Economic and Social Affairs.

Fighting in urban environments is not new for the U.S. military, but the cities of the future may be vastly different from what U.S. forces have seen in the past. They pose all sorts of complications for troops, Ledé said.

Cities are not only growing bigger, they are growing taller. “The verticality element is continuing to grow, and that is a challenge for many reasons: detection, targeting, and prosecution of targets in a 3-D environment is a considerable challenge,” he said.

Fighting in urban areas is marked by corner-to-corner engagements over short distances with equally short timelines, he noted. There are only seconds to discriminate between hostile actors and innocent bystanders.

“We need to see around the corner and observe enemy activities before they target us,” he said. Currently, the way ground forces detect enemies is by being shot at. “Not ideal,” he noted.

C3D – Drone Scouts to the Fore

DARPA’s TTO is wrapping up a program called Centralized Control for Commercial Drones, or C3D, which is using small, inexpensive unmanned aerial systems (UAS) to provide a better vanguard for troops.

The idea is to send these intelligence-gathering systems ahead of troops to make sure routes are clear, set up perimeter overwatch, and maintain surveillance on a building or intersection, said Ledé, who is C3D’s program manager. It is important that the workload is carried out automatically, he added.

Operators submit the mission they want performed through an intuitive interface. The software decides how many drones are needed and how often. “The aircraft are launched automatically, and they are assigned an air corridor to deconflict with one another. Before the first drone has to fly back, the second one is sent,” Ledé said. “The person doesn’t have to do anything.”

When the system detects something – perhaps a person of interest coming out of a building, or carrying a weapon – a warning

is sent to the human operator. The system requires a great deal of autonomy, artificial intelligence, and machine vision to complete its task, Ledé added. The program already is transitioning to the Marine Corps, and has sparked the interest of U.S. Special Operations Command and the U.S. Air Force, which could use it for base defense, Ledé said.

Along with drones that serve as the eyes of ground forces, DARPA's TTO is also working on swarms of ground and air robots that both move ahead of troops and fight alongside them.

OFFSET – A Swarm of Unmanned Allies

The Offensive Swarm-Enabled Tactics (OFFSET) program's goal is to dramatically augment small-unit forces in an urban environment with upwards of 250 small ground and air robots, said its program manager, Dr. Timothy Chung.

"Can we speak a language of swarm tactics – coding as well as operator level – to convey a commander's intent in a highly dynamic urban environment?" he asked.

OFFSET is in the beginning stages of trying to answer this question. It is expected to advance the science of swarm technology.

"Human-swarm teaming has to dramatically change compared to what we do today," Chung said. That may give rise to what would be a new job in the military: the "swarm tactician." This operator will have to go beyond the mouse, keyboard, and screen to control swarms of hundreds of robots, he said.

"We don't exactly know what the best concept of operations is for these swarms," he said. It could entail pushing the robot mass forward without troops, where they serve as a vanguard. Or they could march or fly amongst the warfighters. They could be called on to make first contact with enemies, then be joined by soldiers or Marines, he said.

OFFSET will have three increments. The first is based on Army and Marine Corps tactics, techniques, and procedures. It will first demonstrate isolating an urban objective with 50 robots over a two-square city block for 15 to 30 minutes. The swarm will establish a perimeter, and "put eyes on" potential avenues of attack or reinforcement.



Two 3rd Brigade Combat Team, 4th Infantry Division Personal Security Detail soldiers descend the stairs of a school in Muhalla 511 during a patrol in the Sadr City district of Baghdad on Sept. 27, 2008. DARPA is exploring technologies that will make U.S. warfighters safer and more effective in urban combat.

The second objective will be based on a raid and boost the number of robots to 100. They will be expected to do building ingress, localize targets or signals of interest, and maneuver indoors and outdoors for one to two hours over four square city blocks, Chung said.

The ultimate goal will be to seize urban terrain with 250 ground and air robots. They must capture a city square, bridges, or other objectives and hold onto eight blocks for four to six hours, he said. Along with a physical testbed, the program will have a virtual testbed using synthetic technology.

"It allows us to develop swarm tactics ahead of, and at a pace faster than, technology that currently exists," Chung said. For example, what if the robotic systems had see-through-wall sensors. What if they had unjammable communications? The software developed in the virtual testbed could then be transferred to the real world once these features are mature, he added.

Burns noted that rivals already are employing drones against U.S. forces. They have even used swarm tactics, although not in large numbers. These small aircraft will be particularly hard to defend against in urban environments.

Troops "need eyes 360 degrees around them and above them. You just don't know what's coming at you from above. The Marines and Army have seen a lot of examples in Iraq of ... swarm-like IEDs," or improvised explosive devices, Burns said.

John Waterston, a program manager in the Strategic Technology Office, said, "If we want to have freedom of movement in an urban environment, we are going to have to figure out how to deal with these kinds of threats."

It's relatively easy to detect and defeat a rogue drone out in a field or the middle of the desert, but not so in the cities. "The traditional sensing approaches do not work well in an urban environment," Waterston said. There is lots of radio-frequency interference and clutter. Audio sensors have to deal with noise from cars, and radars with motion from fans in windows or people on the street.

Aerial Dragnet – Above the Streets and Buildings

The Aerial Dragnet program is seeking to deal with the unmanned aerial vehicle threat using a series of sensors.

"It's not really focused on countering these threats; it's really focused on the first part, which is detection and discrimination between unmanned platforms" within the clutter of everyday urban life, Waterston said.

The main impediment is that in a big city, the terrain is almost all non-line-of-sight. That means multi-modal solutions – radar, optical, or acoustic – he said. With high-rise buildings, the aircraft can approach from above as well.

"Sensors can hear the UAS rotors or a radar could bounce off a building to pick them up. You can tell someone is using their weedwacker even though you don't see their weedwacker," Waterston said, to give a sense of the kind of detection he envisions emerging from the Aerial Dragnet program. The catch is that the program seeks technologies that are relatively inexpensive.

Defending Mobile Forces Against Hostile Drones

Even once hostile drones are detected, defeating them in an urban environment is tricky. Ledé is managing the Mobile Force Protection (MFP) program, which seeks to defend high-value convoys against rogue drones.


Today's counter-UAS systems mostly rely on jamming or taking over the UAS' data links. The program will track, ID, and engage the small UAS of



The 3-D environments of megacities pose challenges to troops and technologies, both due to the "vertical terrain" of tall buildings and the lack of unobstructed sight lines.

the future, which may not rely on radio-frequency signals for command and control.

The proximity to civilians and innocent bystanders makes the problem particularly difficult in cities. "We cannot start shining a multi-kilowatt laser with a building in the background," Ledé said. "We have to have counter-UAS systems that will not have hazards for civilian populations."



"Kinder, gentler solutions include nets that entangle the aircraft, or obscuring optics to blind them," Ledé said.

The already-completed first phase of the MFP program sought to stop unmanned aerial vehicles in a relatively open space outside a forward operating base. In the upcoming Phase 2, the program will try, in an urban scenario, to deny the operation of unmanned aerial systems that are no longer using radio frequency signals to communicate. Phase 3 will up the anti-UAS challenge by tasking performers in the program to handle scenarios with multiple drones attacking a friendly convoy on the move.

Beneath the Streets – The SubT Challenge

While cities are growing taller, and ground forces must look to the skies for threats such as rogue drones, there is a whole additional layer to the urban environment below.

"As much as the cities are building upwards, I think they are going to build downwards as well," said Chung. "The subterranean environments pose an increasingly more threatening environment for our warfighters." That could include clandestine tunnels or city infrastructure such as sewers or subway lines, he said.

Chung is managing the DARPA Subterranean (SubT) Challenge, which will distribute prize money to teams that can demonstrate that they can map, navigate, and search underground structures. They will have to find objects quickly, check air for breathability, and determine the navigability of the tunnels, he said.

This artist's conception of the Aerial Dragnet program depicts several small UAVs carrying sensors to form a network that provides wide area surveillance over all low-flying UAVs in an urban setting.

Phase 1 tasks technology developers with demonstrating the ability to map illegal tunnels. Such tunnels have been found under borders for smuggling, and have also been used in war going back centuries. The Viet Cong famously used such tunnels in the Vietnam War, with some being built under Army bases. Phase 2 of the SubT Challenge will include the mass-transit tunnels, sewage systems, and other underground structures found in urban settings. Phase 3 of the competition will be the most challenging. Contestants will survey naturally occurring caves – replete with their complex irregularities – rather than the structures found in cities.

The SubT Challenge will ask teams to map all three together. Similar to the OFFSET program, contestants can write software and try their ideas on a synthetic course rather than a physical course. The winner of the synthetic course will receive \$750,000 and the winner of the physical course will receive \$1 million, Chung said.

Over centuries of warfare, urban combat has traditionally been among the bloodiest and most dangerous of operational scenarios. Today, DARPA research is finding ways to trade technology for blood, providing eyes, ears, and allies for troops fighting the battles of tomorrow in an increasingly perilous urban battlespace.

TAKING NEUROTECHNOLOGY INTO NEW TERRITORY

With new interface technology, the human-machine relationship is deepening

By Justin Sanchez and Robbin Miranda

In a laboratory at the University of Utah, Doug's virtual left hand reached out and touched the virtual door for just a second before he quickly pulled it back. A few moments later, he extended his virtual fingers again and ran them down the simulated wood grain of the door's surface.

"I just felt that door," Doug said with a gasp. "That is so cool."

Doug is one of the first people ever to tangibly interact with a virtual world, using his mind to guide his avatar in a virtual space, and feeling that virtual embodiment's simulated contact directly in his own brain. For Doug, it was as if his own flesh-and-blood hand had touched a real wooden door.

More amazing than that is that Doug is missing his left hand and part of his left arm due to a long-ago accident. A quarter-century later, as a volunteer research participant in DARPA's Hand Proprioception and Touch Interfaces program — nicknamed HAPTIX — Doug felt real touch sensations.

Doug is one of the first testers of an implanted, peripheral nerve interface. When he imagines moving his missing arm and hand, signals from his brain travel down to the peripheral nerves and muscles of his residual limb. The interface reads Doug's neural motion-planning signals, and then uses a set of algorithms by which the interface "learns" to translate the signals into a set of motion controls that a computer relays to a virtual hand.

This interface goes beyond previous one-way systems that enabled Doug to move a virtual or even physical prosthetic limb by thought alone, for with those systems he could not receive tactile feedback. The HAPTIX program closed the loop and made it possible for users of upper-limb prostheses to regain the tactile sensations of pressure and texture. Researchers in the HAPTIX program now are refining the user experience by providing nuanced sensations. The research could prove transformational for amputees, and, as the system becomes better, it might also provide opportunities for able-bodied people to engage with machines and virtual spaces in unprecedented ways.

Just how important is touch? Imagine trying to use a prosthetic hand to handle a delicate lightbulb, applying enough pressure and finesse to screw it in to a socket, but not so much as to shatter the bulb. Without touch


feedback, a prosthetic limb user has to rely solely on visual information, which is a poor way of measuring force, to interact with an external object. With haptic feedback, someone like Doug can manipulate objects and move more confidently and could theoretically complete tasks in the dark or with his eyes closed.

A man named Nathan proved that point. Nathan was paralyzed in an automobile accident in 2004. Since then, he has lived with impaired signaling between his brain and peripheral nervous system, a deficit that precludes candidacy for a peripheral nerve interface. Twelve years after his injury, Nathan took the opportunity to volunteer to be implanted with a direct interface to his central nervous system. He now contributes to DARPA's Revolutionizing Prosthetics program, helping researchers test technology for complex sensorimotor control of prosthetic limbs and other devices.

With electrodes on his motor and somatosensory cortices — the areas of the brain that control movement and touch sensation — Nathan can control a prosthetic arm using his thoughts alone, and he can feel what the arm touches via signals from sensors embedded on the prosthetic fingers. The technology is currently precise enough that Nathan can distinguish contacts with individual fingers. During experiments in 2016 at the University of Pittsburgh Medical Center, a blindfolded Nathan correctly identified which fingers on a prosthetic arm were being pressed by a researcher, and he could even tell when the researcher engaged two fingers simultaneously.

"Sometimes it feels electrical, and sometimes it's pressure, but for the most part, I can tell most of the fingers with definite precision," Nathan told a *Washington Post* reporter for an Oct. 13, 2016 article. "It feels like my fingers are getting touched or pushed."

The concept of "man-computer interaction" began circulating around DARPA in 1960, when computer scientist J.C.R. Licklider published his vision of how computers could one day augment human abilities. As director



of the agency's Information Processing Techniques Office, Licklider shepherded development of a suite of technologies for visualizing, processing, sharing, and interacting with information. Based on the notion that humans and computers could have a symbiotic relationship that produced a sum greater than its parts, these technologies ultimately formed the foundations of today's internet and personal computing experience.

By the 1970s, Licklider's vision inspired DARPA's first research into human-machine interactions facilitated by direct neural interfaces. An early set of experiments explored how well noninvasive sensors could measure responses to sensory stimuli experienced while performing tasks. At the time, the enabling technology for meaningfully interacting with the brain did not yet exist, and so the research results were marginal. But that situation began to change by the late 1990s with the accumulation of advances in information systems, materials science, and sensors for studying brain structure and function at a new level of detail.

By the early 2000s, DARPA began investing heavily in neurotechnology. The agency established the Brain-Machine Interfaces program to record patterns of neural activity in animal models and decode the neural states associated with memory formation, sensory perception, and motor intent.

When the research moved to humans, the ability for a user to directly observe neural decoder outputs in the form of a moving cursor or robotic arm proved critical. That visual feedback allowed the user's brain to adapt — essentially altering its own function to help the neural decoder achieve the task. Subsequent development of more advanced decoders opened the way for iterative co-adaptation between the system's algorithms and the user's neural activity, which further accelerated a user's mind-based motion control. Today, researchers expect that the added ability to convey near-natural touch sensations will further improve feedback-driven learning.

During the in-human studies conducted under the HAPTIX and Revolutionizing Prosthetics programs, study participants were so highly engaged

DARPA has invested more than \$500 million in support of the White House Brain Initiative since it was announced in 2013. This federal investment has accelerated the development of innovative neurotechnologies with the potential to improve human health and change the way people live, work, and play. Among the breakthroughs realized with DARPA funding was development of the CLARITY method by researchers at Stanford University, which allows intact brain tissue to be studied in rich, three-dimensional detail and makes possible better understanding of how brain processes work.

with the work that they effectively became part of the research team. That dynamic made it possible to tailor the system's performance to the needs and wants of the participant. For instance, Jan, a woman living with quadriplegia, quickly achieved the goal of feeding herself a chocolate bar with a prosthetic arm controlled by a direct interface to her central nervous system. Jan next determined to project herself beyond the confines of her wheelchair and to enter a simulated cockpit. Despite being unable to move from the neck down, Jan used her neural interface to fly a virtual aircraft simply by looking at the plane on a monitor and visualizing it moving one direction or another.

"I could raise the nose of the plane up and down. Then I could bank it right or left," Jan explained. "I was lost so quickly in that world because I was up in the clouds, and I was flying. And I was out of my chair. I was out of my broken body. I was flying!"

Nathan was similarly able to extend his abilities. Adding more sensors to his prosthetic arm and hand enabled him to detect infrared (IR) signals. Nathan used his brain signals to move the hand over a surface that emitted invisible IR signals only in a specific location. When the prosthesis crossed the target, the sensors converted the IR signal into electrical



In late 2016, DARPA's Biological Technologies Office, in partnership with the DEKA Research & Development Corp., delivered prosthetic limbs to Dr. Paul Pasquina, chief of the Department of Rehabilitation at the Walter Reed National Military Hospital in Bethesda, Maryland.

pulses delivered to Nathan's somatosensory cortex, enabling him to "feel" infrared radiation. Nathan reported an immediate, touch-like perception of the IR field. Still unknown is whether users' brains, after long-term use of a bidirectional interface with novel sensors like the IR ones, will adapt to a new type of input and ultimately experience a "sixth sense" in a new way.

Prosthetic movement and sensation are currently the best-studied applications for neural interfaces. Through the Neural Engineering System Design (NESD) program, DARPA has even extended its aspirations for higher-resolution upgrades of such systems to potentially restore hearing and vision to people with sensory deficits. What's more, the agency's view of human function expands beyond the motor and sensory domains and into the realm of cognitive function. That is why DARPA's Biological Technologies Office has set out to explore whether implanted neural interfaces might be used to treat individuals with neuropsychiatric and memory disorders.

DARPA's efforts in treating neuropsychiatric dysfunction kicked off in 2013, leveraging the availability of clinical devices for monitoring brain activity that federal regulators already had approved. DARPA-funded researchers recruited individuals with epilepsy or Parkinson's disease, who, as part of their clinical treatments, had electrodes implanted in various regions of their brains. As many as half of these patients also experience symptoms such as anxiety or depression, making them especially suited for DARPA's research.

One of these volunteers, Jane (name changed to protect the participant's privacy), participated in a study funded under the Systems-Based Neurotechnology for Emerging Therapies (SUBNETS) program while

she underwent neurosurgical monitoring for epilepsy at the University of California at San Francisco (UCSF). In addition to epilepsy, Jane was diagnosed with major depressive disorder and bipolar disorder and exhibited symptoms of severe anxiety.

The SUBNETS team fitted Jane with a new type of therapeutic neural interface that records a patient's neural activity across interconnected subnetworks of the brain and delivers targeted, corrective electrical micro-stimulation designed to mitigate unhealthy brain activity. When the system delivered stimulation to a subnetwork of Jane's brain responsible for regulating her emotions, she reported, "All of a sudden ... I have some energy!" When a clinician asked her whether that emotional state was something she would experience on a good day, she affirmed, "This is normal Jane."

That question addressed an important aspect of the SUBNETS approach. The program is not pursuing interventions designed to simply flip an emotional switch from sadness to happiness, but rather to maintain a healthy balance between emotional states by detecting and modulating extremes. An optimal therapeutic intervention would relieve a depressed patient from prolonged sadness or apathy, but still allow the individual to feel a normal range of emotion in response to experiences that typically evoke negative feelings in healthy individuals.

The complexity of the brain makes the process of developing such interventions especially challenging. Neuropsychiatric conditions are often associated with abnormal states across multiple cognitive functions, such as emotion regulation, propensity for risk-taking, and cognitive flexibility, each of which is associated with a distinct subnetwork of the brain. Patients may also fall on opposite ends of the spectrum for these conditions. As such, clinicians must tailor interventions to a patient's symptoms. That is why DARPA required that SUBNETS systems simultaneously record from multiple locations in a patient's brain and interpret neural signals in real time to determine specific locations, parameters, and timing of therapeutic interventions.

A related approach opened the way for closed-loop cognitive prostheses designed to facilitate memory formation and recall. Foundational studies in rodents that began in the early 2000s, led by a research team at Wake Forest University, quickly progressed into non-human primates, and now humans.

The goal was to identify patterns of neuronal activity that predicted successful memory formation, then use artificial stimulation to reproduce those patterns to improve performance during memory tests. The studies relied on algorithms, developed by research partners at the University of Southern California (USC), that characterize how neurons in the hippocampus communicate with one another while a task is performed. The algorithms revealed patterned codes in neuronal activity that corresponded with successful memory formation. When these codes were delivered to the hippocampus during task performance, researchers found that rodents performed with 15 to 20 percent greater accuracy compared to non-stimulated trials or trials in which random patterns of stimulation were delivered.

Inspired by these results, the researchers extended their work to non-human primates under DARPA's Restorative Encoding Memory Integration Neural Device (REMIND) program. REMIND ran from 2009 to 2014 with similarly favorable results.

By 2013, it was time to test the approach in humans. As with SUBNETS, the Restoring Active Memory (RAM) program moved quickly by recruiting volunteers who were already implanted with clinically approved electrodes for monitoring and treating intractable epilepsy. In many cases, epilepsy



Nathan, a participant in DARPA's Revolutionizing Prosthetics program, controls a robotic arm with his brain and feels what the limb touches.

patients also experience memory deficits as a result of their disease.

When the Wake Forest/USC team applied its computational model and closed-loop recording and stimulation approach to the volunteers, they found average performance improvements of 35 percent or more for short-term and long-term memory.

A separate RAM-funded research effort led by a team at the University of Pennsylvania took a different approach to improving memory formation. Rather than focusing on the activity of single neurons in a small area of the brain, the researchers recorded electrical activity generated by small populations of neurons throughout multiple regions of the brain, typically recording from up to 200 sites at once. Their goals were to identify the brain state associated with strong memory performance and deliver stimulation to improve memory function during poor brain states. The Penn team demonstrated that this closed-loop therapeutic approach could facilitate participants' ability to form new memories by way of continuously monitoring brain activity and intervening only when necessary.

If the research bears out, such technologies could find wide application among healthy, sick, and injured people alike. However, the requirement for surgery to implant electrodes currently represents an obstacle too great for all but those individuals with the most

pressing needs. This leaves DARPA with the ethically charged question of how to make the benefits of neurotechnology accessible to the Defense Department's primarily able-bodied population of service members.

DARPA created the RAM Replay program in 2015 to develop interventions involving non-invasive transcranial stimulation or presentation of sensory cues during sleep to boost the process of neural replay associated with successful consolidation of memories. A number of these approaches focused on precisely timing the intervention to the specific phases of slow-wave brain oscillations during sleep that are associated with periods of neural replay—that is, replay of brain activity that occurred during an initial learning experience.

For example, during experiments led by Teledyne Scientific, participants were presented with environmentally realistic sounds while they learned to navigate through a virtual city. Later, while the participants napped, those same sounds were quietly played back to only certain individuals during stages of sleep characterized by slow-wave brain oscillations. After their naps, the participants who received this acoustic intervention—known as targeted memory reactivation (TMR)—could navigate the virtual city roughly 40 percent faster during subsequent trials than control-group participants.

Through the Targeted Neuroplasticity Training (TNT) program, which began in 2016, DARPA invests in both implantable and noninvasive interfaces to improve training outcomes through precisely targeted and timed stimulation of peripheral nerves, which directly modulate the release of various neurotransmitters in the brain. Through these neurochemical mechanisms, peripheral nerve activity can influence the ability of neurons to form new connections with one another—a process called synaptic plasticity—which is crucial to learning new skills and acquiring new



The Restoring Active Memory (RAM) program funds development of closed-loop cognitive prostheses to facilitate memory formation and recall. Through RAM Replay and Targeted Neuroplasticity Training, DARPA is also pursuing noninvasive neurotechnologies to improve training outcomes.

memories. The TNT program funds fundamental research in animals and humans to examine mechanisms of how peripheral nerve stimulation affects plasticity and training outcomes, as well as the long-term safety and efficacy of this approach.

Inspired by the ability of peripheral nerve activity to modulate organ and immune function, DARPA created the Electrical Prescriptions (ElectRx) program in 2014 to explore how nerve stimulation can be applied to preserve human health and to speed recovery. ElectRx researchers are developing new noninvasive and implantable devices for closed-loop sensing and modulation of peripheral nerve activity, and are building functional maps that link specific nerves with physiological responses. If the program is successful, it will open the door to alternative medical interventions that work with far greater precision than current pharmaceutical approaches, and with fewer side effects.

As with all emerging disciplines, progress in neurotechnology is limited by the tools available. Specifically, neural-interface performance is constrained by the number and resolution of channels for recording and stimulating nerves and neurons. Consider that the remarkable achievements described so far have been accomplished using technology that engages up to several hundred neurons at a time. What might be

possible when researchers can interface and intervene with much greater numbers of neurons and with greater precision in the context of detecting, interpreting, and modulating neural function?

The Neuro Function, Activity, Structure, and Technology program, or Neuro-FAST, is creating fundamentally new optical approaches to writing in and reading out information from specific cell types within the brain, each of which plays a distinct role in the biological mechanisms underlying behavioral and cognitive function. Rather than relying on electrodes and electrical stimulation, Neuro-FAST-derived interventions use light to record and control cell activity with single-neuron precision. This “optogenetic” technology has already been demonstrated in animals and is destined for human trials.

The aforementioned NESD program funds multiple approaches to developing high-resolution, implantable neural interfaces – including optical systems – that can measure activity from up to 1 million neurons and modulate the activity of up to 100,000 neurons. NESD researchers are also generating advanced computational approaches for real-time processing of the vast amount of data that neural interfaces will record so that their systems function within available power and bandwidth constraints.

In yet another expansion of its neurotechnology portfolio, the agency this year launched the Next-Generation Nonsurgical Neurotechnology (N²) program to develop a bidirectional neural interface system built primarily around wearable technology. Researchers must overcome the physics of passing signals through the intact skull and brain tissue, but DARPA believes that recent advances in bioengineering, neuroscience, synthetic biology, and nanotechnology can contribute to a portable, precise, high-resolution brain interface. If the program achieves its intended endpoints,



DARPA's Next-Generation Nonsurgical Neurotechnology program aims to develop a high-resolution, portable neural interface system capable of reading from and writing to multiple points in the brain at once. Such a noninvasive system would extend the power of advanced neurotechnology to able-bodied individuals and could support future Department of Defense efforts to improve human-machine teaming.

researchers will demonstrate a defense-relevant task, such as neural control of an unmanned aerial vehicle, using an N³ system.

Although myriad technical barriers remain, the evolution of neurotechnology has followed a trajectory toward more compact, more powerful systems. DARPA's investments have made it possible and prudent to consider what practical applications of advanced neural interfaces imply for our day-to-day lives.

Unlike today's specialized interface systems, which facilitate one or two functions at a time, future, generalizable technology may integrate many functions derived from continuous neural, physiological, and environmental measurements. As the brain and these systems adapt together, an interface might even deliver new function in response to a user's needs and wants and evolve in ways that blur the distinction between internal brain states and the external world. Imagine, for example, an interface that would replay while you sleep the neural codes elicited by the French vocabulary words you studied during the day, pump your IR sensitivity to improve your night vision as you take an evening stroll with your family, or send a wireless signal to your home thermostat when it senses that you are feeling chilly.

At this inflection point in the emergence of neurotechnology, it is imperative to weigh potential risks and trade-offs and to address questions about the ethics of human augmentation. Some societal challenges related to access to technology and the equal spread of benefits are obvious, but even fundamental neurobehavioral questions about how the brain allocates resources across functions also remain to be answered. Could it be, for instance, that enhancing memory impairs problem-solving abilities? That kind of unknown may be acceptable to an individual suffering from

severe functional impairment, but perhaps not to an individual who simply wants to boost normal performance. The transition of neurotechnology from the lab to real-world applications requires such questions to be addressed, which is why DARPA engages independent ethical and legal experts to further the necessary discussions about how the Agency should consider the societal implications of the new neurotechnologies that could emerge from this work.

Tangible technological progress may come sooner than people expect. Many interface technologies are now advancing beyond the proof-of-concept stage and nearing a natural point for commercial transition. The announcement of the White House BRAIN Initiative in 2013 attracted additional investment in the types of foundational technologies DARPA had already been proving out for decades, and since then several companies and venture capitalists have set out to create a neurotechnology industry, often in cooperation with scientists and engineers first funded by DARPA.

At the center of this activity, DARPA's Biological Technologies Office has been adapting its business practices to facilitate technology transition by encouraging commercial involvement from the very outset of new programs. With that input, and with DARPA assuming the early and highest-risk investments, industry partners will be poised to expand and refine the most successful capabilities and package them for end users. This cycle of creation, adaptation, and commercialization directly benefits the Department of Defense by returning proven, field-ready technologies at reduced cost.

The story of neurotechnology is unfolding at an increasing rate and is currently being written by DARPA and other early movers. Even at this nascent stage, the neurotechnology era is pointing toward a future in which humans and machines seamlessly interact to overcome the limitations of the human body and merge human adaptability and creativity with the reliability, ubiquity, and processing power of computers. After nearly six decades, Licklider's vision of human-machine symbiosis is becoming reality in a way that he and his peers from the agency's early days might have only dreamed to be possible.



A PROTEAN TECHNOLOGY

DARPA and the exploration of artificial intelligence

By John Everett


A mere 70 years ago, when early electronic computers ran on vacuum tubes and filled entire rooms, researchers already were striving to enable these machines to think as people do. Only a few years after its start in 1958, DARPA began playing a central role in realizing this ambition by laying some of the groundwork for the field of artificial intelligence (AI). Today, the remarkable pattern-recognition capabilities of deep neural networks, whose circuitry and operation are modeled after those of the brain, have the potential to enable applications ranging from agricultural machines that spot and zap weeds to apps that, with the help of augmented-reality displays on smart phones, enable people to diagnose and solve problems with home appliances.

DARPA is now funding research to enable AI programs to clearly explain the basis of their actions and how they arrive at particular decisions. AI that can explain itself should enable users to trust it, a good thing as increasingly complex AI-driven systems become commonplace. A next challenge for AI researchers will be to emulate human common sense, which is a product of millions of years of human evolution.

DARPA has been on the forefront of establishing the foundations of AI since 1963. J.C.R. Licklider, a program manager in the Information Processing Techniques Office, funded the Project on Machine-Aided

Cognition (MAC) at MIT and a similar project at Stanford to research a wide range of AI topics, such as proving mathematical theorems, natural language understanding, robotics, and chess. Early AI researchers focused on chess because it presents a difficult intellectual challenge for humans, yet the rules are simple enough to describe easily in a computer programming language.

In the 1950s and 1960s, computers were automating boring and laborious tasks, like payroll accounting, or solving complex mathematical equations, such as plotting the trajectories of the Apollo missions to the moon. Not surprisingly, AI researchers ignored the boring applications of





J.C.R. Licklider, the first director of DARPA's Information Processing Techniques Office, funded the Project on Machine-Aided Cognition at MIT, and a similar effort at Stanford to research a range of artificial intelligence topics.

computers and instead conceived of artificial intelligence as computers solving complex mathematical equations, expressed as algorithms. Algorithms are sets of simple instructions that computers execute in sequence to produce results, such as calculating the trajectory of a lunar lander, when it should fire its retro rockets, and for how long.

Despite more than a half-century of trying, we have yet to invent an algorithm that enables computers to think the way people do. Early on, AI researchers discovered that intelligence depends not just on thinking, but also on knowledge.

Consider chess. In the middle of the game, each player must ponder around 35 possible moves. For each of these moves, the player's opponent will have 35 or so countermoves. To determine which move to make, the player must think ahead multiple turns into the future of the game. To think ahead two turns requires consideration of 42,875 moves. To think ahead seven moves would require contemplating 64 billion moves. The IBM Deep Blue supercomputer that beat chess champion Gary Kasparov in 1997 could evaluate 200 million board positions in a second, so looking ahead seven turns would take it a little under six minutes. However, looking ahead nine turns would take it almost two days. Since chess games typically take 50 turns, this brute-force approach of considering all possible moves clearly won't work.


Chess champions use knowledge of the game to ignore most potential moves that would make no sense to execute. The first AI chess programs used heuristics, or rules of thumb, to decide which moves to spend time considering. In the 1960s, this approach enabled Mac Hack VI, a computer program written by Richard Greenblatt, who was working on Project MAC at MIT, to win against a ranked player in tournament play.

As the centrality of knowledge to intelligence became apparent, AI researchers focused on building so-called expert systems. These programs captured the specialized knowledge of experts in rules that they could then apply to situations of interest to generate useful results. If you've ever used a program such as TurboTax to prepare your income tax return, you've

used an expert system. Edward Shortliffe created one of the first expert systems, MYCIN, for his doctoral dissertation at Stanford University in the early 1970s. MYCIN used a set of around 600 rules to diagnose bacterial infections based on input about symptoms and medical tests. It achieved 69 percent accuracy on a set of test cases, which was on par with human experts. Digital Equipment Corporation used an expert system in the early 1980s to configure its computers. Such early successes led to a boom in AI, with the founding of companies such as Teknowledge, Intellicorp, and Inference Corporation. However, it became apparent that expert systems were difficult to update and maintain, and they would give bizarrely wrong answers when confronted with unusual inputs. The hype around AI gave way to disappointment in the late 1980s. Even the term AI fell out of favor and was superseded by terms such as distributed agents, probabilistic reasoning, and neural networks.

Language is so fundamental to our daily experience of the world that early researchers assumed they could write down all necessary knowledge to enable an AI system. After all, we program computers by writing commands in programming languages. Surely, language is the ideal tool for capturing knowledge. For example, an expert system for reasoning about animals could define a bird as an animal that can fly. However, there are many exceptions to this rule, such as penguins, ostriches, baby birds, dead birds, birds with one or more broken wings, and birds with their feet frozen in pond ice. Exceptions to rules crop up everywhere and expert systems do not handle them gracefully.

By the late 1980s, another approach to AI was gaining momentum. Rather than focus on explicitly writing down knowledge, why not try to create machines that learn the way people do? A robot that could learn from



Artificial intelligence has developed in two major waves. The first wave focused on handcrafted knowledge, in which experts characterized their understanding of a particular area, such as income tax return preparation, as a set of rules. The second wave focused on machine-learning, which creates pattern-recognition systems by training on large sets of data. The resulting systems are surprisingly good at recognizing objects, such as faces. DARPA believes that the next major wave of progress will combine techniques from the first and second wave to create systems that can explain their outputs and apply commonsense reasoning to act as problem-solving partners.

people, observations, and experience should be able to get around in the world, stopping to ask for directions or calling for help when necessary. So-called machine-learning approaches try to extract useful knowledge directly from data about the world. Rather than structuring this knowledge as rules, machine-learning systems apply statistical and probabilistic methods to create generalizations from many data points. The resulting systems are not always correct, but then again, neither are people. Being right most of the time is sufficient for many real-world tasks.

Neural networks are an effective machine-learning method. They emulate the behavior of the brain. The human brain consists of a network of interconnected cells called neurons. Electrical signals flow through this network from sense organs to the brain, and from the brain to the muscles. The human brain has something like 100 billion neurons, each of which connects, on average, to 7,000 other neurons, creating trillions of connections. Signals that travel through this network arrive at neurons and stimulate (or inhibit) them. When the total stimulation exceeds the neuron's threshold, the cells start firing, an action that propagates the signal to other neurons.


Researchers have developed artificial neural networks that mimic these aspects of the brain in order to recognize patterns. Instead of writing computer code to program these networks, researchers train them. A common method is called supervised learning, in which researchers collect a large

set of data, such as photos of objects to be recognized, and label them appropriately. For example, to train a network to recognize pets, researchers would collect and label photos of animals such as dogs, cats, and rabbits.


The network consists of layers of nodes arranged in columns and connected left to right by links. On the left side, each input node (think of each node as a neuron) is assigned to a particular part of the photo to be recognized. Each of the links of the network has a weight, which determines how strongly it propagates a signal. Initially these weights, which range from 0 to 1, are set to random values. The training algorithm shows the network a labeled photo. The brightness of each electronic node's assigned piece of the photo determines how intensely that node fires. The signals from the input nodes cause some of the nodes in the layer to their right to fire, and eventually these signals propagate through the network and cause some nodes in the output layer to fire. In the pet example, each output node is assigned a particular type of pet, so there will be one output node for dogs, one for cats, and one for rabbits. The output node that fires most strongly produces the network's answer. The training algorithm compares that answer to the label of the photo. If the answer is wrong, the algorithm works its way backward through the network, making small tweaks to the weights of each link until the answer is correct. This process is called back-propagation.

The training algorithm repeats this process with each labeled photo in the training set, making small adjustments to the weights of the links until the network correctly recognizes the pet in each photo. If that were all the network could do, it wouldn't be very useful. However, it turns out that the trained network will correctly recognize the countless photos of dogs, cats, and rabbits that it hasn't seen before.

By extracting knowledge directly from data, neural networks avoid the need to write down rules that describe the world. This approach makes them better for capturing knowledge that's hard to describe in words. For example, suppose that you need to collect a large set of pictures of refugees. Writing out a set of rules for an expert system to do this would be very difficult. What are the relevant features of such pictures? Many, but



- We are entering a new age of AI applications
- Machine learning is the core technology
- Machine learning models are opaque, non-intuitive, and difficult for people to understand



- Why did you do that?
- Why not something else?
- When do you succeed?
- When do you fail?
- When can I trust you?
- How do I correct an error?

not all, pictures of refugees contain crowds of people, but so do pictures of sporting events, urban areas, and nightclubs. And yet we humans have no problem distinguishing refugees from football fans. A lot of our knowledge is hard to express in words.

Computers became powerful enough to run neural networks in the 1980s, but the networks couldn't be very large, and training was almost as much work as writing rules, since humans have to label each element of the training set. In 1985, DARPA funded two teams under the Autonomous Land Vehicle program to develop self-driving cars. Both teams used neural networks to enable their vehicles to recognize the edges of the road. However, the systems were easily confused by leaves or muddy tire tracks on the road, because the hardware available at the time was not powerful enough. Nonetheless, the program established the scientific and engineering foundations of autonomous vehicles, and some of the researchers went on to NASA to develop the Mars rovers *Sojourner*, *Spirit*, and *Opportunity*. All of these autonomous vehicles operated far longer than specified in their mission plans.

In 2004, DARPA issued a Grand Challenge, with a \$1 million prize awarded to the first autonomous vehicle to cross the finish line of a 142-mile off-road course in the California desert near Barstow. The most capable vehicle traveled less than 8 miles before getting stuck. In 2005, DARPA repeated the challenge on a slightly shorter but more difficult course, and this time five vehicles crossed the finish line. The teams that developed those vehicles used neural networks to enable better detection of the track and to distinguish obstacles such as boulders from shadows. Many of those researchers went on to develop self-driving car technologies for Google, Uber, and other car manufacturers.

Now AI seems to be everywhere. Over the past few years, AI constantly has been in the news, due to rapid advances in face recognition, speech understanding, and self-driving cars. Oddly enough, this wave of rapid progress came about largely because teenagers were eager to play highly realistic video games. Video consists of a series of still pictures flashed on a screen, one after the other, to create the illusion of motion. Realistic video requires the creation and display of lots of high-definition pictures, because they must be displayed at a ferocious rate of at least 60 pictures per second. Video screens consist of a dense rectangular array of tiny dots,


As AI applications become more common, the current limitations of the technology become more apparent. In particular, machine-learning systems cannot explain their outputs.

To address these issues, DARPA is running a program called Explainable AI to develop systems that can produce accurate explanations at the right level for a user. Systems that can explain themselves will enable more effective human/machine partnerships.

called pixels. Each pixel can light up in any one of more than 16 million colors. The processors underlying fast-action video games need to create a constant stream of pictures based on the actions of the player and shunt them onto the screen in quick succession.

Enter the graphics processing unit, or GPU, a computer chip specifically designed for this task. GPUs rapidly process large arrays of numbers that represent the colors of the pixels comprising each picture. In 2009, NVIDIA released powerful new GPUs, and researchers soon discovered that these chips are ideal for training neural networks. This enabled the training of deep neural networks that consist of dozens of layers of neurons. Researchers applied algorithms invented in the 1980s and 1990s to create powerful pattern recognizers. They discovered that the initial layers of a deep network could recognize small features, such as edges, enabling subsequent layers to recognize larger features such as eyes, noses, hubcaps, or fenders. Providing more training data makes all neural networks better, up to a point. Deep neural networks can make use of more data to improve their recognition accuracy well past the point at which other approaches cease improving. This superior performance has made deep networks the mainstay of the current wave of AI applications.

Along with GPUs and clever algorithms, the internet has enabled the collection and labeling of the vast amounts of data required to train deep neural networks. Before automated face recognition was possible, Facebook provided tools for users to label their pictures. Crowdsourcing websites recruit inexpensive labor that AI companies can tap to label pictures. The resulting abundance of training data makes it seem that AI systems with superhuman abilities are about to take over the world.



However, deep neural networks are woefully inefficient learners, requiring millions of images to learn how to detect objects. They are better thought of as statistical pattern recognizers produced by an algorithm that maps the contours of the training data. Give these algorithms enough pictures of dogs and cats, and they will find the differences that distinguish the one from the other, which might be the texture of fur, the shape of the ear, or some feature that conveys a general sense of “dogness” or “catness.”

For some applications, this inefficiency is not an issue. Internet search engines can now find pictures of just about anything, from cats sitting on suitcases to people playing Frisbee on a beach. For applications where training data is scarce, neural networks can generate it. An approach called generative adversarial networks takes a training set of pictures and pits two networks against each other. One tries to generate new pictures that are similar to the training set, and the other tries to detect the generated pictures. Over multiple rounds, the two networks get better at generation and detection, until the pictures produced are novel, yet usefully close to real ones, so that they can be used to augment a training set. Note that no labels are required for this generation phase, as the objective is to generate new pictures, not classify existing ones.

However, we still lack a solid theoretical foundation to explain how neural networks function. Recently, researchers have discovered that making a minuscule change to an image causes wild misclassification. A picture of a panda is suddenly identified as a monkey, even though it still looks like a panda. Understanding these anomalies is essential, as AI systems are increasingly being used to make critical decisions in areas as varied as medical diagnosis and self-driving vehicles.

Another approach to machine learning relies on cues from the environment to reinforce good behavior and suppress bad behavior. For example, the program AlphaGo Zero can teach itself to play the board game *Go* at

Once trained, current machine-learning systems no longer adapt to their environments. DARPA’s Lifelong Learning Machines program is researching ways to enable systems to learn from surprises and adapt to changes in their environments. The Assured Autonomy program is developing approaches to produce mathematical assurance that such systems will operate safely and predictably under a wide range of operating conditions.

championship levels without any human input, other than the rules of the game. It starts by playing against itself, making random moves. It uses the rules of the game to score its results, and these scores reinforce winning tactics. This so-called reinforcement learning can be highly effective in situations where there are clear rewards for effective behavior. However, determining what behavior created the desired result in many real-world situations can be difficult.

As AI increasingly makes its way into industrial settings and consumer products, companies are discovering that its substantial benefits come with costs, in the form of engineering complexity and unique requirements for ongoing maintenance. The computational intensity of AI systems requires racks of servers and networking gear, which must be secured against and continuously monitored for cyber intrusions. The insatiable appetite of these systems for data often makes them dependent on many different enterprise databases, which requires ever-increasing coordination of operations across the organization. And finally, machine-learning systems must be continually retrained to keep them in sync with the world as it continually changes and evolves. This retraining requires a team of skilled data scientists and AI engineers, who tend to be in short supply.

Ultimately, people are still far more effective learners than machines. We can learn from teachers, books, observation, and experience. We can quickly apply what we've learned to new situations, and we learn constantly in daily life. We can also explain our actions, which can be quite helpful during the learning process. In contrast, deep learning systems do all their learning in a training phase, which must be complete before they can reliably recognize things in the world. Trying to learn while doing can create catastrophic forgetting, as backpropagation makes wholesale changes to the link weights between the nodes of the neural network. DARPA's Lifelong Learning Machines program is exploring ways to enable machines to learn while doing without catastrophic forgetting. Such a capability would enable systems to improve on the fly, recover from surprises, and keep them from drifting out of sync with the world.

The knowledge of a trained neural network is contained in the thousands of weights on its links. This encoding prevents neural networks from explaining their results in any meaningful way. DARPA is currently running a program called Explainable AI to develop new machine-learning architectures that can produce accurate explanations of their decisions in a form that makes sense to humans. As AI algorithms become more widely used, reasonable self-explanation will help users understand how these systems work, and how much to trust them in various situations.

The real breakthrough for artificial intelligence will come when researchers figure out a way to learn or otherwise acquire common sense. Without common sense, AI systems will be powerful but limited tools that require human inputs to function. With common sense, an AI could become a partner in problem-solving. Current AI systems today seem superhuman because they can do complex reasoning quickly in narrow specialties. This creates an illusion that they are smarter and more capable than they really are. For example, you can use an internet search engine to find pictures of cats sitting on suitcases. However, no current AI can use the picture to determine if the cat will fit in the suitcase. To the AI, the thing we recognize as a furry animal that purrs, uses a litter box, and ruins the single most expensive piece of furniture in the house with its needle-sharp claws is just a fuzzy two-dimensional texture. The AI has no conception of a three-dimensional cat.

What would happen if somebody put a cat in a suitcase? Suffocation as a possibility leaps to mind, because we learn while growing up to consider the likely consequences of our actions. We also know that a plush toy cat can go in an airtight suitcase with no ill effect, because we categorize things according to their properties, such as being alive, not just by their appearance (in this case, an inanimate toy that looks just like a real cat). No AI system today can do this type of reasoning, which draws on the immense amount of commonsense knowledge we accumulate over our lifetimes.

Commonsense knowledge is so pervasive in our lives that it can be hard to recognize. For example, most people could easily sort through pictures of furniture to find black couches with white legs. As an experiment, try to find such pictures on the internet. In 2018, at least, your search results will contain mostly pictures of black couches and white couches, with pictures of various other couches thrown in just in case. The color of the legs will be whatever is in fashion, because we don't yet understand how to create AI systems that can figure out the parts of objects.

AI systems with common sense ultimately could become partners in problem-solving, rather than just tools. For example, in emergency

situations, people tend to make snap decisions about the cause of the problem and ignore evidence that doesn't support their point of view. The cause of the Three Mile Island accident was a stuck-open valve that allowed cooling water to escape from the reactor containment vessel. The heat of the reactor caused the remaining water to turn to steam, which increased the vessel pressure. The operators decided that the high pressure meant that there was too much water, and made the situation worse by overriding the automatic emergency cooling system. An AI that could understand control-room conversations and vet them against its own models of reactor operation might be able to suggest alternative possibilities before the human operators commit to a particular course of action. To act as a valued partner in such situations, the AI system will need sufficient common sense to know when to speak and what to say, which will require that it have a good idea of what each person in the control room knows. Interrupting to state the obvious would quickly result in its deactivation, particularly under stressful conditions.

DARPA is mounting a major initiative to create the next generation of AI technologies, building on its five decades of AI-technology creation to define and to shape what comes next. DARPA's substantial AI R&D investments will increase to fund efforts in the following areas:

New Capabilities. DARPA programs routinely apply AI technologies to diverse problems, including real-time analysis of sophisticated cyber attacks, detection of fraudulent imagery, human language understanding, biomedical advances, and control of prosthetic limbs. DARPA will advance AI technologies to enable automation of complex business processes, such as the lengthy and painstaking accreditation of software systems required for aviation, critical infrastructure, and military systems. Automating this accreditation process with known AI and other technologies now appears possible, and would enable deployment of safer technologies in less time.

Robust AI. As noted above, the failure modes of AI technologies are poorly understood. The data used to train such systems can be corrupted. The software itself is vulnerable to cyber attacks. DARPA is working to address these shortfalls by developing new theoretical frameworks, backed by extensive experimental evidence, to fortify the AI tools we do develop.


High-Performance AI. In combination with large data sets and software libraries, improvements in computer performance over the last decade have enabled the success of machine learning. More performance at lower electrical power is essential to allow this use of AI for data-center applications and for tactical deployments. DARPA has demonstrated analog processing of AI algorithms that operate a thousand times faster using a thousand times less power compared to state-of-the-art digital processors. New research will investigate AI-specific hardware designs and address the inefficiency of machine learning by drastically reducing requirements for labeled training data.

Next-Generation AI. DARPA has taken the lead in pioneering research to develop the next generation of AI algorithms, which will transform computers from tools into problem-solving partners. New research will enable AI systems to acquire and reason with commonsense knowledge. DARPA R&D produced the first AI successes, such as expert systems and search utilities, and more recently has advanced machine-learning tools and hardware. DARPA is now creating the next wave of AI technologies that will enable the United States to maintain its technological edge in this critical area.

THE NETWORK OF OUR TIMES

A DARPA perspective on the development of the internet

By Henry Kenyon



The beginning of the internet. A sketch of the 1969 ARPANET network.

The history of the internet is inextricably connected with DARPA and the development of new networking and information-sharing technologies. From its humble beginnings as a collection of connected research facilities to today's global (and extra-global) network of billions of devices, a core theme of the internet's evolution has been and continues to be the cutting-edge research conducted under DARPA's auspices. When internet pioneer J.C.R. Licklider described, with intentional grandiloquence, his vision of an "Intergalactic Computer Network" in 1963, he still might not have imagined just how extensively the network whose seeds he planted would grow in the terrestrial realm.

In its scope and ubiquity, the internet has transformed life on a civilization-wide scale. Smartphones and other mobile devices connect bankers in Manhattan and rice farmers in India to a vast sea of information. Using predictive tools that mine this information has allowed child welfare agencies to flag high-risk cases, enabled health care providers to determine which patients are at risk for various conditions and diseases, and helped financial institutions to identify fraud and identity theft, among other benefits. But there are challenges in this new planetary data environment: The same networks of computers and databases and systems and infrastructure that afford society previously unimagined benefits also have been bringing with them new vulnerabilities for violations of privacy, cybercrime, and cyber warfare.

To get a sense of how the networked present came to be as it is now, and to imagine the potential future of an internet

The Interface Message Processor (IMP) was the first packet router for the ARPANET. The first ARPANET transmission was between UCLA and Stanford Research Institute (SRI) on Oct. 29, 1969.

driven by artificial intelligence, virtual/augmented reality, and the Internet of Things (IoT), a visit to the past is a must.

In the Beginning

The roots of the modern internet lie in the groundbreaking work DARPA began in the 1960s (when the Agency was known as the Advanced Research Projects Agency, or ARPA) to create what became the ARPANET. In its earliest form, ARPANET began with four computer nodes, and the first computer-to-computer signal on this nascent network was sent between UCLA and the Stanford Research Institute on Oct. 29, 1969. “To understand how the ARPANET got started, it’s important to realize that at a time when there were very, very few computers in the world, a few visionaries started seeing the idea that eventually [computers] would all want to talk to each other,” said James Hendler, Director of the Rensselaer Institute for Data Exploration and Applications, and who served as chief scientist of DARPA’s former Information Systems Office.

Secure communications and information-sharing between geographically dispersed research facilities was one of ARPANET’s original goals. As more computers became involved in this early computer network, however, engineering problems arose. A key issue was maintaining communications, because if ARPANET behaved like a traditional circuit-based telephone system, failure of a single node could take the entire network down.

What was needed was a means to get messages to their destination in a way that did not require the presence of any single node. This is where the concept of packet switching originated. By moving bits of data that dynamically worked their way through a network to the destination and reassembled themselves there, the problem of data loss if one or more nodes went down could be avoided.

A common communications protocol between computers was also necessary, because the computers involved “were anything but compatible,” noted Vint Cerf, Vice President and Chief Internet Evangelist (this



STEVE JURVETSON PHOTO VIA WIKIMEDIA COMMONS

○ IMP
 □ TIP
 △ PLIP BUS IMP

NOTE: THIS MAP DOES NOT SHOW ANY AS YET EXPERIMENTAL
 SATELLITE CONNECTIONS.

NAMES SHOWN ARE IMP NAMES, NOT IEEE (SSRN) NAMES

The ARPANET in 1977.

The initial implementation of the first TCP/IP protocol occurred at Stanford in 1975. As it tested out over the next few years, the now-famous protocol was being implemented on a growing number of computer operating systems around the world. In January 1983, enough individual networks

To achieve widespread use of the network, DARPA supported work that opened a way for general-purpose machines to do the work of the internet rather than relying on specialized machines. These multipurpose devices became key elements of the first server farms as the internet emerged and then began to rapidly expand. It was these underlying technologies and standards that helped set the stage for the first major internet-based companies, such as AOL and Netscape.



A visualization study of inbound traffic measured in billions of bytes on the NSFNET T1 backbone for the month of September 1991. The traffic volume range is depicted from purple (zero bytes) to white (100 billion bytes). It represents data collected by Merit Network, Inc.

By the early 1980s, the bones of what would become the internet were in place. This new and growing network, with continued support by DARPA and the DOD, became available to the academic community, which led to the launch of other government networks. In 1983, the first part of what would become the National Science Foundation's (NSF) computing backbone, NSFNET, was laid down using an optical fiber network.

NSFNET grew rapidly, connecting thousands and then millions of university computers around the United States by the early 1990s. The success of the network convinced the Department of Energy (DOE) and NASA to establish their own networks using the same protocols, which led to the birth of the DOE's Energy Sciences Network and the NASA Science Internet. These new networks joined NSF's to become the four major networks of the U.S. internet.

Commercial internet service did not exist until 1989. Cerf noted that commercially available software and hardware products available from firms like Proteon, and later Juniper and IBM, were available for dial-up use. "You could buy those [products]. You couldn't buy [internet] service, because it was only available to people who were sponsored by the U.S. government research agencies," Cerf said.

In 1988, Cerf, who was then vice president of the Corporation for National Research Initiatives, received permission from the Federal Networking Council to connect the MCI Mail service, which had been providing commercial email services over phone lines via modem since

1983, to the internet. As president of MCI's Digital Information Services from 1982 to 1986, Cerf led the development of MCI Mail. This served to break the restriction prohibiting commercial traffic from traveling over the government backbone, he explained.

This new commercial, rather than government-based, email service launched in 1989. It opened the floodgates for commercial players, and the major email providers of the time, such as AOL, CompuServe, and Telenet's Telemail, which were providing messaging via telephone modems, soon followed suit. This rapid expansion of commercial internet services prompted the NSF to shut down its dedicated backbone in 1995.

Powering the Growing Web

The 1990s saw the development of web search technology, driven by users' need for more efficient ways to find information as the World Wide Web expanded exponentially. Work in web search technology derived largely from earlier work in information retrieval from text, such as DARPA's 1991 TIPSTER program. These earlier programs helped set the stage for a line of DARPA investments aimed at improving information extraction and search that still continues today. By the late '90s, the need for better machine processing of the information on the growing web, as well as for more precise search capabilities, was becoming critical to the ability of users to find what they were looking for. In the national security context, the challenges to accessing, managing, and sharing data residing on the classified DOD networks were becoming more important to solve. "A lot of information management technologies had been developed by DARPA," Hendler said, but "the question was, how could we apply them to this exponentially growing space?"

Much of DARPA's work before the late 1990s balanced the needs of different government user communities with the challenge of controlling when and where to share — or not to share — information between individuals, agencies, or battlefield commanders. However, as the information needs of warfighters grew, especially for complex operations, DARPA also started working with the DOD to develop better technologies for improving the sharing of information between the different military services and particularly for coalition operations in which the United States had to work with forces from other countries or the U.N.

One particular challenge was that even within the military, let alone on the open web, people needed means to find documents or data that might be relevant for them, even when the terms used were different. “A good example,” said Hendler, “was that what a United States warfighter would refer to as a ‘MIG-25’ might be called a ‘Foxbat’ in a report filed by United Nations troops.” Users needed a means by which search and retrieval wasn’t just limited to the words in a document but could also include knowledge about how entities were related to each other. “We needed a way that users could search on ‘Russian fighter aircraft’ and find MiGs, Foxbats, and all the other things which they might be called,” said Hendler.

To address this problem, DARPA invested in the DARPA Agent Markup Language (DAML) program, which Hendler started up in 1999. A key researcher on the project was MIT professor Tim Berners-Lee, who had created the World Wide Web during his time at CERN in the late 1980s and who coined the term “Semantic Web” in the mid-’90s. It was clear that for these technologies to be used widely by the military, they would have to be supported by commercial web companies, which could provide the scale needed to be used across the many millions of pages on the growing web. The early languages of the Semantic Web, promoted by Berners-Lee, Hendler, and others, eventually diffused throughout the World Wide Web to become commercial standards. For example, one of the best-known uses of the Semantic Web has been in the creation of the “knowledge graph” technologies used today by large internet search and social-networking companies. Semantic Web technologies are also widely used in the health care and life sciences community.

Many of the early concepts explored by DARPA in the mid-1990s also played in the development of command, control, and communications technologies that changed how the DOD operated. The commercial applications of these technologies found their way into products in the 2000s. One example was DARPA's Personal Assistant that Learns (PAL) project, which opened pathways to the development of popular commercial voice-response software applications such as Siri and Alexa. “You start to see a lot of growth as these things went out into the commercial world,” Hendler observed, “but if you trace back a lot of that work to its roots, you’ll find DARPA’s fingerprints.”

Another outgrowth of DARPA's funding early in the new century was the application of machine-learning and artificial intelligence (AI) technologies on the web. The development of machine-learning systems led to the creation of new applications and practices such as data-mining, which opened new avenues for rapid analysis and use of massive data sets across the internet. DARPA was a key developer of these machine-learning technologies, and much of the current boom in big data grew out of DARPA's investments.

The Wireless Internet and Beyond

The packet-switching technology that helped launch ARPANET also led to another networking revolution. Packet-switching-based messages based on TCP/IP standards were originally sent via telephone wires, but the same information subsequently was converted into radio signals that could be transmitted via satellite relay stations. Later, this approach was applied to terrestrial radios, said Randy Katz, Vice Chancellor for Research at the University of California, Berkeley.

In the late 1970s, DARPA launched a program to develop a packet radio networking system. Intended to support DOD requirements for mobile wireless networking, the challenge was to make the radios operate in tactical environments. This presented completely different engineering constraints compared to the fixed ARPANET infrastructure. “The packet radio program of the late 1970s was focused on, maybe not sending as many bits, but in a much more challenged environment, much more dynamic ... because that mirrored the typical battlefield application,” Katz said.

Commercial wireless applications did not become profitable until the late 1980s, with the development of the predecessor of Wi-Fi technology. Advances in miniaturization also helped drive what would become the wireless revolution. Katz noted that in the late 1970s, the technology that would eventually become the basis of cell phone communication filled a large van, largely because of the computational power required to run the software-intensive TCP/IP framework. By the late 1980s, the system was miniaturized down to a small computer with an antenna in a package roughly the size of a pack of cigarettes.

It is best to view the advent of wireless communications in terms of mobile computing, Katz explained. This was driven by the co-evolution of computer miniaturization with the networking technology developed by ARPANET and NSFNET. This allowed computers, and what would become smartphones, to be connected wirelessly using the emerging technology of Wi-Fi.

Katz joined DARPA in the early 1990s to head the Global Mobile Information Systems (GloMo) program, which aimed to accelerate the development of mobile computing (later in the same decade, Katz would become deputy director of DARPA's Computing Systems and Technology Office). GloMo looked at dual-use technologies for digitizing military and civilian communications. Katz noted that beginning in the early 1990s, DARPA became very interested in mobility and wireless networking to support the DOD's efforts to modernize its communications and command and control systems.

Much of DARPA's work at the time in this area was in understanding how mobile wireless systems could be scaled up and operate in a changing environment. These mobile networks also had to be backward compatible with all the information and resources already available on the internet, Katz said. Scientists and engineers also had to ponder how mobile devices could exchange information without a centralized network or set of servers. “All of these ideas really blossomed in the time frame of the late 1990s into the early 2000s,” he said.

Wireless networking and broadband will continue to influence near- and long-term computing for both commercial and military applications, Cerf said. The advent of the first true smartphones in the early 2000s (capable of internet access, GPS location, video recording/cameras, and high-speed wireless network access) and the release of the Apple iPhone in 2007 heralded the beginning of the current era of ubiquitous access



DAVID BERKOWITZ PHOTO VIA WIKIMEDIA COMMONS

ABOVE: The control panel of an LG smart refrigerator at CES 2011. DARPA's Brandeis program aims to enable the safe and predictable sharing of data from the billions of "Internet of Things" (IoT) devices while preserving privacy. RIGHT: A Nest Learning Thermostat, just one node in the IoT.

to data. An infrastructure of blanket Wi-Fi and cellular coverage in most urban areas across the globe has greatly expanded ease of access.


"People can get access to the internet from anywhere they can get an LTE, 3G, or 4G connection," Cerf said. "And because they use their smartphones, smartphones are benefited by the fact that they have access to all the broadband internet. So you get these mutually reinforcing effects starting up in 2007 and continuing until today."

But this access is a double-edged sword. On the one hand, it allows unprecedented access to information. Roughly half the planet's population can access the internet and its sea of data, Cerf said. But the downside to this easy availability is that anything can be put on the network, including malware and misinformation.

"Now we have to figure out how to cope with the fact that this infrastructure can be abused and is being abused, and we need combinations of critical thinking by the users of the net to reject bad information," Cerf said.

Part of the solution could come in the form of technical means to probe and manage broadband wireless communications and networking, possibly supported by regulations to mitigate the international reach of crime and misinformation on the internet, Cerf said.

AI and advanced machine-learning, now appearing in applications like Siri or IBM's Watson, will have an impact on the future internet, Hendler said. Instead of replacing people, he sees a future where smart AI tools work with humans to take advantage of machines' abilities to search and sift through massive amounts of data. A person would make a final decision on a process or subject, and apply intuition and experience backed up by the AI's input. For example, a doctor could use his or her AI tools to search a vast number of medical texts for a specific set of symptoms and use that information to come to a diagnosis. But finding the right balance of human and machine intelligence for particular tasks, whether it be medical diagnoses or



battlefield decisions, will be a challenge. Said Hendler, "When do you decide that maybe the machine is right and when do you decide that maybe the human is right, or when do you decide that together they're right?"

The Internet to Come

Questions surrounding the right balance of human and machine intelligence remain moot for a majority of the world's population, however. While the internet has become a pervasive part of life in the West, roughly half

PHOTO BY RAYSONHO @ OPEN GRID SCHEDULER / GRID ENGINE VIA WIKIMEDIA COMMONS

of the world's citizens still have no access to the internet. One of the more important reasons for that lack is the cost of what today is the infrastructure necessary to its growth. Part of the solution to that problem may lie in DARPA-funded research on wireless and mobile networks that began more than 20 years ago.

In the mid-1990s, DARPA began work to enable secure, dependable battlefield communications. The key issue was that unlike the civilian internet, a military network could not depend on fixed infrastructure such as cell towers, routers, and server farms, partly because these fixed but vital nodes would be prime and easy targets on a battlefield and partly because such infrastructure probably would not be where it was needed. What was needed was a mobile network, and DARPA-backed research that helped create a network of devices, such as laptops and mobile phones, that required no fixed infrastructure. Instead of devices routing information through fixed nodes, each device could send and receive information from one another. Should one device be lost to the network, many others would still be available to send and receive data. Scaled upward from a local area network, the concept paved the way for “mesh networks.”

In a mesh network, rather than traveling to a fixed node such as a router or cell tower, information travels across the mesh until it reaches its recipient. Since every device seeks out other devices to connect to wirelessly, only one device has to be connected to a fixed router.

For the civilian world, this could mean investing less in expensive and vulnerable infrastructure, and instead building a network that is both cheaper and more robust. The cities of Garland, Texas, and Medford, Oregon, in the United States have already invested in mesh networks, for example. Perhaps more importantly, such networks could deliver the internet to large swaths of the world that still lack access.

One other important side benefit of the mesh network is that it enables various devices on the IoT — things like thermostats, home automation systems, vehicles, and wearable devices — to communicate with each other as well as provide connectivity for more devices that might be too far away from an internet source.

By 2017, there were 8.4 billion of these devices connected online in the IoT, according to the PEW Research Center, and numerous studies estimate there will be 20 to 50 billion devices networked online by 2020. This explosion of devices creates both great benefits and great challenges. The IoT depends upon making the data and storage of those connecting devices, from wearable “watches” to home automation systems, accessible to IoT companies and storable in the cloud. This raises issues of privacy and security, as well as citizens’ ownership of their own data. The DARPA program Brandeis is all about trying to break the seemingly incompatible goods of privacy and the ability to tap into the great value of large and diverse databases. The goal of Brandeis, in the word of its managers, “is to enable safe and predictable sharing of data in which privacy is preserved.”

Indeed, these issues are among the greatest concerns for Berners-Lee almost 30 years since his creation of the World Wide Web.

“We must work together with web companies to strike a balance that puts a fair level of data control back in the hands of people, including the development of new technology like personal ‘data pods’ if needed, and exploring alternative revenue models like subscriptions and micropayments,” Berners-Lee wrote in a World Wide Web Foundation blog post in March 2017. “We must fight against government over reach in surveillance laws, including through the courts if necessary. We must




The Disruption Tolerant Network protocols will enable the Solar System Internet, allowing data to be stored in nodes until transmission is successful.



push back against misinformation by encouraging gatekeepers such as Google and Facebook to continue their efforts to combat the problem, while avoiding the creation of any central bodies to decide what is ‘true’ or not. We need more algorithmic transparency to understand how important decisions that affect our lives are being made, and perhaps a set of common principles to be followed. We urgently need to close the ‘internet blind spot’ in the regulation of political campaigning,” he wrote.

With programs such as “Explainable AI,” which seeks to develop AI tools that explain how they arrive at conclusions and make decisions, and the new Media Forensics (MediFor) program, which seeks the capability to automatically detect manipulations of online imagery and to reveal precisely how these manipulations were performed, DARPA is taking steps toward the safeguards Berners-Lee calls for. Though many challenges and issues face the internet, today it has in many ways fulfilled Licklider’s vision of “the main and essential medium of informational interaction for governments, institutions, corporations, and individuals.” Indeed, Licklider’s “intergalactic” network has now actually reached out into space.

One of ARPANET’s lesser-known but potentially most-expansive contributions is interplanetary networking, Cerf explained. Beginning in 2008, researchers at NASA’s Jet Propulsion Laboratory in Pasadena, California, used a protocol called Disruption Tolerant Networking (DTN) to successfully transmit data to far-flung spacecraft in Earth orbit. The goal of DTN is to support communications and data transfer across interplanetary distances. The driver of this effort, according to NASA, is because, “when data are transmitted and received across thousands and even millions of miles, the delay and potential for disruption or data loss is significant.” In 2016, a dedicated DTN link between Earth and the International Space Station went online, making it the very first part of an interplanetary internet, Cerf said.



SPINTRONICS

A DARPA-spurred spin on fundamental electron physics just keeps on giving

By Ivan Amato

On a late summer day in 1995, Stuart Wolf arrived for work at the Advanced Research Projects Agency with a 40-pound, radiation-hardened memory cabinet that engineers at the Naval Center for Space Technology (NCST) had taken out of a satellite they were assembling at the Naval Research Laboratory (NRL). They “put it in a box, and I actually took it to DARPA,” recalled Wolf, a physicist who then was splitting his time as a program manager at ARPA (soon to get a “D” for defense in front of its acronym) and materials scientist at NRL. It helped that Wolf was friends with the NCST director at the time, who let him borrow the 128 kilobyte, quarter-million-dollar memory box.

“I plopped the memory on Gary Denman’s desk,” Wolf said, referring to the ARPA director at the time. “I promised him I would replace that memory with a chip that weighed a fraction of an ounce. He approved \$30 million dollars, at that meeting, to start this program,” Wolf said during an interview earlier this year at DARPA headquarters in Arlington, Virginia. That was the moment when the Magnetic Materials and Devices program went from a glimmer in a program manager’s eye to the first of a series of well-funded DARPA programs that later would become recognized as central to one of the great physics-to-technology arcs of our times.

Within a year of that moment, a synergy of physicists, materials scientists, electronic engineers, computer scientists, and others were bonding into a new scientific research and technology development community spawned from the most fundamental discoveries about electronic behavior with roots extending into the 19th century, noted historian of technology Patrick McCray of the University of California, Santa Barbara. In the context of space, the promise of smaller memory technology potentially also meant lighter and less power-needy technology, all of which could amount to a so-called SWaP victory, in which SWaP translates to “Size, Weight, and Power.” For many DARPA programs and their proposed technologies, the lowest measure of each of those traits often has been the metric of success.

In planning the program’s first-year review, Wolf had a linguistic epiphany that had a rare combination of being scientifically spot-on and public-relations gold. Rather than using the drab name of the program – Magnetic Material and Devices – to identify the review, Wolf glommed onto what then was one of the field’s hottest topics, a set of phenomena known as spin transport electronics. Wolf shortened that quantum-mechanics-based phrase to spintronics, which he used in the name of the program review. Wolf’s neologism caught on, becoming the name that stuck for a broad swath of electron-based behaviors and technologies as well as the field devoted to their study and development.

First for some of the basic physics. In 1857, William Thompson (a Belfast-born physicist known more regally as Lord Kelvin) observed that a

magnetic field could cause changes in an electrical conductor’s resistance, a phenomenon that became known as magnetoresistance. One hundred and fifty years later, in 2007, Albert Fert and Peter Grünberg shared a Nobel Prize in physics for their independent discoveries in 1988 of what became known as giant magnetoresistance (GMR). In their respective laboratories in France and Germany, the researchers both observed unexpectedly large magnetoresistance effects in materials made of alternating layers of magnetic and nonmagnetic metals.

The language of “spin” comes into all of this because magnetism depends on an electronic trait as fundamental as its electrical charge but that the quantum mechanical giants of the early 20th century – among them Paul Dirac and Wolfgang Pauli – thought of more like a spinning top. In the quantum mechanical framework, an electron can spin in one of two directions. Materials like iron and nickel, whose outermost electrons can assume spins primarily oriented in one direction because of a fundamental interaction between them (exchange interaction) are magnetic. Gold and silver and wood and most polymers are nonmagnetic.

Within just a few years of the initial GMR reports in the late 1980s, Stuart Parkin and his colleagues, working at the IBM Almaden Research Center near San Jose, California, had exploited the GMR effect for an invention they referred to as a “spin valve.” And in 1997, IBM rolled out the first use of spin-valve sensors in the read head of its 16 GB Deskstar 16GP Titan hard drive. In subsequent years, the technology would ascend to blockbuster status as it became the workhorse of the multi-billion-dollar data storage industry.

Like Grünberg’s and Fert’s Nobel Prize-winning GMR materials, the IBM team made its spin valves out of a pair of thin magnetic layers flanking a nonmagnetic layer. One of the magnetic layers had a magnetic moment (a measure of the net magnetic effect of the orientations of a material’s numerous electrons) that was pinned in a given direction or “fixed” and the other had a “free” magnetic moment that could be flipped. When the magnetic moments of the outer layers in these GMR structures align in the



Airman 1st Class Nick Fazio verifies the status and configuration of the 96th Communications Squadron's data center equipment at Eglin Air Force Base, Florida. Approximately 40 airmen and civilians are responsible for the connectivity of the base's more than 19,500 user workcenters. DARPA spintronics research enabled transformative technologies that have become the backbone of data storage today.


same direction, electrons move between them with less resistance (an open spin valve) than when the moments of the layers are not aligned (a closed spin valve). Because even subtle magnetic fields, such as those from tiny magnetically-encoded bits on a storage disk, can toggle the magnetic moment of a spin valve's "free" magnetic layer, these GMR structures held great promise for new generations of read-head materials that could read data stored at much higher densities than before. Unlike the exquisitely controlled but impractical fabrication methods that the Nobelists used to make their GMR materials, Parkin and his IBM team from the get-go had in mind GMR materials made with easier sputtering techniques that their manufacturing brethren readily could apply on industrial scales.

In Wolf's early years at ARPA, beginning in late 1993, the agency was managing a well-funded "dual-use" program called the Technology Reinvestment Program (TRP), which the Clinton administration had begun a year earlier to forge stronger collaborations between military and commercial R&D sectors. This controversial policy stance about technology

was a contributing factor that enabled Wolf to easily secure \$5 million of seed money to organize what he called the GMR Consortium. Its members ranged from large and established research organizations such as Honeywell and NRL to nimble startups like Nonvolatile Electronics.

The goal of the Consortium's two-year effort was to identify GMR-based technology trajectories that could make an impact, perhaps in both defense and civilian contexts, and in 1995, the Consortium's verdict was in: GMR research could yield a host of new magnetic field sensors but, more importantly, a new kind of memory — magnetoresistance random access memory (MRAM). For one thing, the magnetic basis — rather than the charge-basis — of the storage technology opened the way to memory that would be retained even when power to a computer or storage device was shut off. For another, MRAM promised instant-on computer operation that would bypass the then standard need to boot-up by shuttling information from a hard drive to the powered-up computer chips. Crucially for national defense purposes, MRAM devices also could be "rad-hard" compared to semiconductor-based RAM chips, and therefore less vulnerable to radiation damage in places like space. On top of all of that, MRAM research looked like a good bet for winning a significant SWaP advance in memory technology.

Jon Slaughter was part of a semiconductor research group at Motorola that got in early on the DARPA push for new MRAM technology. At first, his team focused on a type of GMR structure in which electrons traveled between the two magnetic layers via a conductive pathway. But they soon realized that the way to go, one that IBM already had adopted, was with



A hard-drive read head with a spintronics-based operation that enables it to read data stored at high densities.

so-called tunnel magnetoresistance (TMR). In this process, electrons move between the two magnetic layers by “tunneling” through an intervening insulator layer in that mysterious quantum-mechanical sense. “This is mind-blowing,” Slaughter said, “because the electrons are never in the insulator.”

Mind-blowing yes, and a commercial coup to boot. Slaughter says his team, working on what was referred to within the company as the “Panther Project,” identified a variation on the TMR theme, called toggle switching, as the technology to bet the farm on. In 2004, Motorola spun off its vast Freescale semiconductor operation (including its MRAM research) as an independent entity and two years later Freescale became the first company in the world to sell an MRAM product – a 4 Mb MRAM chip based on toggle switching. Freescale did not have a memory business at the time and was not investing in new markets, Slaughter said in an interview. So he, with help from others in the community, attracted sufficient venture capital to spin out in 2008 what then was the MRAM startup Everspin Technologies. A decade later Slaughter is the director of research and development and, according to the company website, 70 million of its MRAM products are deployed “in data center, cloud storage, energy, industrial, automotive, and transportation markets.”

“People don’t realize it, but MRAM is out there,” Slaughter said. In a poetic twist, the process Everspin relies on for its rad-hard MRAM products depends on rad-hard silicon supplied by Honeywell, which was the other

The magnetic basis – rather than the charge-basis – of the storage technology opened the way to memory that would be retained even when power to a computer or storage device was shut off.

major company – besides IBM and Motorola – that the agency funded under the Magnetic Materials and Devices program.

A few years into that pioneering program, two fundamental physics discoveries expanded the already rich spintronics vista. A team of Japanese researchers found they could use an electric field – rather than a magnetic field – to induce magnetism in a semiconductor material (gallium manganese arsenide), though they needed to chill the material to 150 K (-123 Celsius). In the same 1997-1998 time frame, physicist and materials scientist David Awschalom and colleagues at the University of California Santa Barbara (UCSB) showed that they too could induce magnetism in a semiconductor material (gallium arsenide), but they did so with a laser.

“These two discoveries gave me the idea that there could be a new DARPA program to explore how to exploit these two effects,” Wolf said. That idea became, in 1998, the Spins in Semiconductors (SPINS) program, which

also benefited from the availability of TRP money. One goal of this mostly fundamental-science program was to discover semiconductor materials like gallium manganese arsenide, but in which the magnetism could be induced with an electric field at the merely chilly temperature of 273 K rather than a frigid 150 K. About a third of the program's funds underwrote what then was a still speculative goal of quantum computation based on "qubits," which are physical entities such as isolated ions or well-engineered crystal semiconductor nanostructures that can embody a superimposition of the canonical one and zero states (and perhaps many more states) at the same time. This was driven by the surprising discovery by the Awschalom group at UCSB that light could create and control quantum coherent states of matter (akin to light waves or mechanical waves whose phases can be deliberately arranged) based on electron spins, and that engineers could use these quantum states for novel types of information processing. Moreover, the UCSB group found that it could turn to simple electrical gates to transport these quantum states across hundreds of microns in semiconductors, quite a long distance in these contexts.

Because this search for qubits was occurring within the specific context of spintronics (electronic spin states), however, several DARPA offices banded together and expanded the quest for physical incarnations of qubits by starting up an ambitious \$100 million Quantum Information Science and Technology (QUIST) program. Running from 2001-2005, the program also supported research into new algorithms for computation with qubits.

Awschalom, now at the University of Chicago, embraced these particularly deep and speculative dives into the quantum physics of electron spins. "DARPA was largely responsible for launching the fields of semiconductor spintronics and quantum spintronics, the latter of which has also helped drive the emerging area of quantum computing," he said. "This was high-risk science and, at the time, not obvious that the underlying physics and engineering would work out so well."

"Development in MRAM took a long time," Wolf noted. "It started in 1995, but it did not lead to a commercial technology until 2006." That actually stands well against the several decades it commonly takes to develop scientific discoveries into hold-in-your-hand technology. The pathway to this end result in the MRAM adventure illustrates what, from DARPA's institutional point of view, is an ideal sequence of events – fundamental research, to proof-of-concept technology, to good-to-go technology with national defense consequences. The end result was full-on DARPA, in this case a rad-hard MRAM chip, which was cheaper, more capable, more energy-efficient, and minuscule compared to the previous space-based storage technology.

Although IBM, one of the three major contractors of the original MRAM program, had already been in the spintronics-based data storage business by the mid-1990s with its spin-valve technology, the company continued to push forward storage technology based on other spintronics developments, such as the invention of TMR, the phenomenon that would set Everspin on a pathway to success. Here, a thin insulator replaces the semiconductor between the magnetic layers in a traditional GMR device. This results in a magnetoresistance signal that is 100 times stronger than that of a spin valve. An article in *The New York Times* on Sept. 11, 2007, credited GMR technology up to that time – more than a decade ago – with making "consumer audio and video iPods, as well as Google-style data centers, a reality."


Every new technology harbors the seeds of its own obsolescence. Even during DARPA's early spintronics programs, researchers knew that in time they would eke out as much technological capability from the spintronic materials and structures they were developing as was physically and

"DARPA was largely responsible for launching the fields of semiconductor spintronics and quantum spintronics, the latter of which has also helped drive the emerging area of quantum computing ... This was high-risk science and, at the time, not obvious that the underlying physics and engineering would work out so well."

practically possible. They knew they would be able to engineer finer and finer magnetic bits of memory in GMR structures, but they also knew that as they did so, those bits would become ever more vulnerable to random switching effects due to thermal fluctuations. The size at which this bit-scrambling limitation would thwart further advances in MRAM technology was at about 50 nanometers (nm), or about twice the diameter of a polio virus. For ambitious memory technologists, even that is too big.

This is where the experimental demonstration in the early 2000s by a DARPA-supported group at Cornell University led by Robert Buhrman of yet another spintronic phenomenon, spin torque transfer, came into the picture. Predicted theoretically by other researchers at IBM and Carnegie Mellon, this phenomenon provided the pathway for overcoming what otherwise could have been an impasse to higher storage densities. If MRAM bit sizes were to shrink beyond sub-viral dimensions, Wolf said, "people knew the switching would have to move from magnetic-field switching to spin-torque switching," which is based on electric current conditioned so that the spins of its constituent electrons align. That alignment or lack thereof – states that can be controlled – enables the electric current to mimic the effect of a magnetic field, including the ability to exert mechanical torque. And that meant the current could serve as the switching trigger in GMR structures. This opened up clever pathways around an obstacle that was confronting MRAM designers. As magnetic bits shrank, they had to be made of materials that were harder to switch because of the thermal jumbling problem. By turning to the spin torque phenomenon, engineers would be able to apply sufficient switching forces to flip the smaller but tough-to-switch magnetic bits.


Everspin, the company that made the world's first commercial MRAM chips based in part on technology that emerged from the MRAM program, is now making spin torque transfer MRAM. "So it took several years for MRAM to run out of steam, but then we had a solution to scale beyond that," said Wolf. "This was a major discovery within our spintronics programs." Buhrmans' group also pioneered work in spin torque oscillation, a subtle phenomenon that engineers have yet to exploit but, said Wolf, "may also lead to a revolution in signal processing and computation." The same



“nano-oscillations” could find applications in radiofrequency (RF) filters, minuscule RF sources, and a novel pattern-recognition technique.

Today, the quest to unveil yet more technological magic in electronic spins continues at DARPA. In 2017, Program Manager Rosa Alejandra “Ale” Lukaszew launched her Topological Excitations in Electronics (TEE) program, where the goal is to find yet new approaches to shrink electron-spin-based magnetic domains while maintaining their resistance to random switching due to thermal fluctuations. Lukaszew has challenged researchers (DARPA calls them “performers”) in the program to focus on skyrmions, which are vortex-shaped multi-electron structures that flip states only as a unit. Those traits bestow the skyrmions with more stability compared to simpler electron organizations. Lukaszew envisions that skyrmion-based memory could, in her words, “offer the possibility of 100 times more storage density at 10 times less power, while being equally stable and rad-hard as today’s technology” compared to other spintronics-based memory. That is the sort of technological leap that could put far more processing power into a soldier’s hands without asking him or her to carry more battery weight. She’s talking SWaP here. Said Lukaszew, “I want to create a community of experts that find materials that can do this.”

Making leaps, rather than taking steps, has always been the DARPA ideal. “DARPA was a huge accelerator” for spintronics, Slaughter said. “It came in just at the right time, just when the scientific developments made it sensible” for well-placed technology developers in the semiconductor sector to move forward on new products like MRAM. The ripples of DARPA’s investment and community building in spintronic continue to expand. “Most of the potential has yet to be realized,” Slaughter said, noting as an example that the world’s top semiconductor foundries are planning on embedding MRAM directly onto processor chips to create “systems on a chip.” Said Slaughter, “MRAM’s best days are to come.”



TOP: Wafers bearing spintronics technology, like this colorized specimen from Everspin Technologies, host a multitude of magnetoresistive random access memory (MRAM) dies before these are diced and packaged into individual MRAM chips.

ABOVE: Skyrmions are vortex-shaped multi-electron structures that flip states only as a unit, bestowing them with more stability compared with simpler electron organizations.

TECHNOLOGY THAT MATTERS

Ushering game-changing technology into the real world

By Craig Collins

In September 2015, at the *Wait, What?* forum hosted by DARPA at the America's Center Convention Complex in St. Louis, Missouri, presenters spoke eloquently about emerging technologies and their potential to radically change the world – but one speaker in particular held the audience spellbound. Kenneth Dwyer, a U.S. Army Green Beret, told the story of what happened on Aug. 19, 2006, when his combined Special Forces-Afghan National Army patrol was ambushed by more than 100 heavily armed enemy combatants at a bazaar in a central Afghanistan village. In the back of a Humvee, while Dwyer was helping a turret gunner reload, a rocket-propelled grenade (RPG) detonated near his left hand, blowing it off at the wrist, and shrapnel ripped into his face and neck, perforating his trachea and esophagus. Dwyer lost his left eye, and many of his teeth, but the injury that threatened to kill him was high on his right arm, where shrapnel had ripped open his brachial artery.

Dwyer credited his survival first and foremost to his fellow service members: “If not for the courage and heroism, the training and the capability, of my team sergeant and my 18 Delta Special Forces medic being able to fight their way through about 75 meters of uncovered terrain through heavy enemy fire to get to my position,” he said, “then I wouldn’t be here today.”

But the story Dwyer wanted to tell wasn’t about himself, he said. “It was more importantly about my recovery, and how science and technology developments from 10 to 15 to 20 years ago not only saved my life, but allowed me to continue to serve my country.”

Dwyer showed the audience a picture of himself in his hospital bed after he’d been sent home to Walter Reed Army Medical Center, an image of a young bearded man with a bandaged left eye and arm. He’s breathing through a tracheal tube just above his collarbone. His face and his right arm are red, scuffed, and scabby. But his torso is immaculate; it looks as if he’s wearing a vest made of unblemished skin.

“Everything that was covered by my body armor,” he said, “everything inside my arms, from my neck down to my abdomen – and you can’t see it, but everything above my eyebrows, my head – are completely clean.”

The helmet and vest that saved Dwyer’s brain and vital organs from an RPG are the fruits of DARPA-sponsored research that began in the 1990s and continues today in the agency’s Soldier Protection Systems program, which is driven to innovative extremes by the increasing lethality of enemy ballistic and blast weapons.

Dwyer’s own resilience was the key factor in his decision to return to service, but during his presentation, he gave credit to a man in the audience: Fred Downs Jr., a veteran who’d lost his own left arm to an antipersonnel mine in Vietnam at the age of 23. Downs, who at the time was director of the VA’s Prosthetics and Sensory Aids Service, went to visit Dwyer personally at the hospital. Said Dwyer, “I remember being motivated by the things Mr. Downs had been able to accomplish in his life regardless of the loss of limb.”

Dwyer’s eventual return to jumping out of airplanes and firing weapons was enabled by prosthetics that provided him with dexterity far beyond the static stump-and-socket limbs of the past. When Downs stood to acknowledge Dwyer’s thanks in St. Louis, he showed off his own prosthetic arm, the LUKE arm, originally developed under the sponsorship of DARPA’s Revolutionizing Prosthetics program, launched in 2006 to improve quality of life for the growing number of Iraq and Afghanistan veterans who were surviving their battlefield injuries limb amputations. A modular, battery-powered robotic arm, the LUKE system enables precise arm and hand movements, allowing users to turn a key or use power tools.

After years of testing and optimization in a DARPA/Veterans Administration collaboration, the LUKE system was cleared for commercial-scale manufacture in 2016 by the Food and Drug Administration. In June 2017, the first two LUKE arm systems were prescribed to veterans, one of whom was Downs. Meanwhile, DARPA’s work on prosthetics continues through several programs aimed at adding more functionality to robotic arms, enabling a sense of touch and connecting directly with the brain’s motor



ABOVE: Lt. Col. Kenneth Dwyer, garrison commander, explains what three things are needed for recovery during a mortal-injury forum hosted by the 188th Combined Arms Training Brigade at Fort Stewart in Georgia. The three elements needed, according to Dwyer, are a sense of humor, being surrounded by positive people, and finding a purpose in life and living it every day. Dwyer credits DARPA-funded research for saving his life.

RIGHT: Fred Downs Jr., a Vietnam veteran and former executive with the Department of Veterans Affairs, shakes hands using a technologically innovative prosthetic arm, known as the "LUKE" Arm, developed by DEKA Research & Development Corp.



cortex to control movements. "The technology today," Dwyer said, "is going to push forward to the point where . . . people will be able to do what we've never imagined, even after injuries."

Speeding Up the Refresh Rate

The technologies that saved Dwyer's life and aided his recovery were nurtured from concept to capability by DARPA program managers. Since its inception 60 years ago, DARPA's mission has been to ensure the United States is the nation imposing technological surprises on its adversaries, rather than the other way around. For several reasons, meeting that mission — transitioning technology out of the laboratory and into the field, where it can make a difference for people like Dwyer — is more of a challenge today than in the past.

More often now, in the Information Age, technologies spread rapidly into the hands of those who, in the past, would have been late adopters. Today, peer and non-peer competitors around the globe have the ability to create and deploy capabilities that narrow our technical superiority. According to Scott Reed, a special assistant to DARPA's director who specializes in transitioning DARPA-developed technology, "DOD is not the driver it once was in terms of information technology. So we now need to become more agile to maintain relevance and superiority inside a much more competitive space. Adversaries have easy access to commercial overhead imagery, when not that long ago it was a capability that was only held by the highest order of nations. Now you can buy it off of Google."

Tom Masiello is director of DARPA's Adaptive Capabilities Office (ACO), a stand-alone office that works with Reed and other members of the director's staff to transition new technologies toward operational military

capabilities. Avoiding surprise in the warfighting context, Masiello said, is getting more difficult. "Sometimes Americans view it as a birthright that we will automatically have technological superiority over potential adversaries," said Masiello. "But that's not necessarily the case . . . Part of the problem is it just takes too long for us to field advanced technology into military capability. We are not lacking in innovation. We are not lacking in cutting-edge research or technological breakthroughs. The problem is transitioning that high technology into military capability in the field."

One of the factors that can bog down technological progress and keep innovations from being fielded is the risk-averse posture of the military acquisition process. If, for example, you're building a \$13 billion aircraft carrier, a program whose primary metrics are that it be delivered at cost and on time, there isn't much incentive to risk using a lot of new technology.

The Pentagon's budgeting process, in which high-level decision-makers project where they want to spend money over the next five years, often lacks the agility to adapt to the accelerating rate of technological change. "Often times," said John Murphy, a former Navy pilot who is now DARPA's Operational Liaison for Technology Transition and Warfighter Engagement, "the DARPA technology is such that just about the time that we're getting ready to prove it, we're looking for somebody to begin budgeting for this technology — but there's no room in the budget, because the budget was created two, three, and four years ago."

DARPA has developed several strategies to overcome such bureaucratic snags. The agency always has considered program managers — with the many contacts in the military-technology community they often bring with them — to be the primary agents responsible for transition. It always has been a tall order — the tasks of building service acceptance of new technologies, and the associated changes to culture, doctrine, and operations, are arduous. Furthering the cause are transition-devoted professionals, such as Murphy, Masiello, and Reed, who support DARPA's R&D-centered program offices as they pursue technical breakthroughs. Adding to the collective of transition talent are uniformed military liaison officers (LNOs) whose job is to identify ways in which those breakthroughs can meet the needs of the service branches of which they have on-ground, at-sea, and in-air knowledge and experience.

Enduring connections between DARPA and the service branches expose program managers to warfighting challenges in different domains, while




Former Deputy Defense Secretary Bob Work poses for a photo in 2016 after a christening ceremony in San Diego, California, for a technology-demonstration vessel designed, developed, and built through DARPA's ACTUV (Anti-Submarine Warfare Continuous Trail Unmanned Vessel) program. In 2018, DARPA transferred the system to the Office of Naval Research for further development.

allowing the services to track the progress of relevant and potentially useful technology development efforts. DARPA has intermittently sponsored specific initiatives or events to encourage relationships and sharing of knowledge among DARPA researchers and their customers. But according to Masiello, these programs, such as a "PM Boot Camp," which connected new program managers with partners in the services, are merely formal structures for carrying out what DARPA does every day. "Our bread and butter is those day-to-day engagements, building those relationships," he said. "We have formal engagements at the four-star and director levels. But we also foster lower-level informal relationships to help us identify the technology gaps or requirements – the needs of the services and the combatant command – which then enable us to seize opportunities where DARPA technology can meet those gaps or requirements."

These working relationships, Murphy added, make it more likely that military budgeteers will be aware of new technologies and their development projections far enough in advance that they can be sensitive to the potential for a DARPA project to transition into a fielded capability. That way, he explained, "When the DARPA program is coming to an end, there is ... some money to keep the program going and keep the program team together, so that it continues on and then you move forward" into a transition phase.

Murphy himself recently led such a transition in October 2016 when a prototype of the world's largest uncrewed naval ship, *Sea Hunter*, arrived under its own power at the U.S. Navy's Space and Naval Warfare Systems Center Pacific. DARPA launched this effort in 2010 to develop what's now known as the Anti-Submarine Warfare Continuous Trail Unmanned Vessel (ACTUV). At the time, an autonomous submarine hunter, operating without remote guidance, seemed a risky investment of the Navy's limited research funds. Murphy likens the Navy's reticence to his own skepticism in the 1990s about the value of the unmanned aerial vehicles (UAVs) DARPA was developing. "It did require a paradigm shift for traditional aviators to recognize the value of a UAV," he said, "and I think the same thing has been true for operators of the surface fleet ... [autonomous] submarine-hunting was a mission area that people thought was just a leap too far."




As the military's high-risk, high-reward research agency, with years of experience in unmanned vehicle development, DARPA decided to assume the risk that others were reluctant to take on. Over time, said Murphy, as the ACTUV program matured and unmanned vehicles of all types began to be deployed and used on the ground, in the air, and under water, "people ... began to see that you could do many things with this unmanned surface vessel, as opposed to filling it with human beings and putting it in harm's way."

In 2014, DARPA signed a Memorandum of Agreement with the Office of Naval Research (ONR) to jointly fund the ACTUV prototype, a 132-foot-long, 140-ton trimaran whose performance in a series of at-sea evaluations indicated it could perform an expanded set of roles beyond antisubmarine warfare (ASW). In 2018, DARPA transitioned the prototype system, with the name of *Sea Hunter*, to ONR for further development. For the Navy, the payoff has been significant: After allowing its R&D partner, DARPA, to assume some of the early risk, the Navy has acquired a revolutionary prototype that could become an entirely new class of naval vessel, capable of traveling thousands of miles over the open ocean, for months at a time, without a single crewmember aboard.

"DARPA develops technology to the point where we prove it's operationally possible," said Reed. "We'll say, 'Hey, the technology can satisfy this operational concept,' and then it will get developed and made more mature somewhere else. And we'll move on to do the next hard problem."

The ACTUV partnership fits this description. Congress, duly impressed by the *Sea Hunter*'s capabilities, has approved funds for the Navy to buy another one and to conduct further evaluations. These revolutionary vessels could be the vanguard of a new era in naval warfare, in which the world's oceans are patrolled by fleets of autonomous surface ships assigned to



specific missions: hunting submarines, relaying communications, clearing minefields, or conducting surveillance.

The speed and urgency of today's technological change compel DARPA program managers, in partnership with their service counterparts, to close near-term technology gaps that weaken national security. And while these relationships are often the key to planning, launching, and aiming a DARPA program squarely at fielded results, a program sometimes proves important enough to national security that extra measures are taken to insure it crosses the finish line. One approach on these occasions has been to establish a joint office devoted entirely to the development of an application, as DARPA and other DOD entities have done with past acquisitions or projects. A recent example of this specialized-office approach, the Long-Range Anti-Ship Missile (LRASM) Deployment Office, or LDO, transitioned a suite of technologies into a capability that was honored with the 2017 Defense Laureate award by the editorial staff of *Aviation Week* in recognition of excellence in the technology's sophistication and the execution of the LRASM program.

The road to establishing LDO began in 2008, when the U.S. military was five years into grueling combat operations in the Middle East. While other nations were developing supersonic anti-ship missiles with increasing strike reach as well as sophisticated electronic anti-access/area denial (A2/AD) countermeasures (involving long-range strike, sensors, guidance, and other military technologies), the American arsenal was falling behind. The former commander of the U.S. Pacific Fleet, Adm. Robert F. Willard, identified an Urgent Operational Need — an anti-ship missile less dependent on traditional intelligence, surveillance, and reconnaissance (ISR) platforms, and with a longer range than existing missiles in this category. Willard reached out to DARPA leadership to accelerate LRASM toward an operational and fielded status.

This heat map of human-trafficking activity across the world is one of the tools that is part of DARPA's Memex program, designed to help law enforcement officers and others do investigations online and hunt down human traffickers.

The weapon developed over the next few years was based on the airframe of Lockheed's Joint Air-to-Surface Standoff Missile (JASSM). Successive flight tests that began in 2013 demonstrated LRASM's ability to navigate independently, including traversing pre-established waypoints before transitioning to autonomous guidance; identify a moving naval target from among a group and hit that target; and detect, track, and avoid an obstacle deliberately placed in its path. Demonstrations of the LRASM have been launched from B-1 bombers, an F/A-18 Hornet, and from shipboard vertical launching systems.

The key to developing these capabilities so rapidly, Reed said, was to structure a process that more closely resembled commercial technical development — a process driven by time rather than technical milestones or specifications. "Apple is going to put out a new iPhone every two years," he said. "They're going to deliver. So they make choices, in terms of tech development, based on being able to meet that schedule." The LRASM demonstration was built around a similar concept: maximizing, rather than perfecting, the technical integration of new technologies into the JASSM platform, up to a certain date. "They had a date on which they were going to deliver that capability," said Reed. "And that was much more of a driving factor than in a traditional DOD acquisition program, which is performance-parameter based."

After demonstrating the system's core capabilities in these first flight tests, LRASM became a formal U.S. Navy program of record in early 2015.

The LDO transitioned the program to the Naval Air Systems Command (NAVAIR) in 2016.

The Military and Beyond

A significant challenge to understanding the diverse ways in which DARPA technologies are transitioned is the fact that every program is different. “There are probably at least 15 different models of how you might transition a particular technology,” said Murphy. “Sometimes transitions involve very small components. In some cases, DARPA creates a new material, or a new process that dramatically changes the structural integrity of certain materials. That’s typically a very, very different sort of transition, probably more likely to be taken up and developed by a commercial customer, rather than directly to a military customer.”

Larger-scale, higher-level military “program of record” transitions, such as ACTUV and LRASM, are successful development efforts that receive approval for follow-on funding by the government, to get them to the finish line of actual and acquired warfighting capabilities. Responsibility for this later-stage development, as with ACTUV, often transfers to another service branch’s research operation. These transitions can begin with a “pull” on the part of DARPA researchers (i.e., demonstrating ACTUV’s technology to the point that it proved alluring enough to the Navy to form an R&D partnership) or a “push” from the service (i.e., the statement of Urgent Operational Need that led to LRASM’s development).

In less formal transition types, some DARPA technologies are handed off to other federal or state government agencies, which have recognized enough promise in the technologies to take them over for next-stage development and even implementation. In the lead-up to the 2014 Super Bowl at the New Jersey Meadowlands, for example, the Federal Bureau of Investigation invited DARPA to sit in on the meetings of a joint task force composed of federal, state, and local law enforcement officials to strategize about how to combat the surge in sex trafficking that’s often tied to major sporting events. These agencies approached DARPA because of its expertise in deep data analysis, as demonstrated by its XDATA program, which produced tools for processing and analyzing large, imperfect, and incomplete datasets that helped several federal agencies identify previously invisible routes of terrorist financing. But according to Wade Shen, a program manager in DARPA’s Information Innovation Office, it didn’t take long to discover that the challenge of finding and prosecuting human traffickers went beyond data analysis.

He and others scanned the Deep Web — an enormous part of the World Wide Web that isn’t discoverable by standard search engines and which is populated with sites that use password protection or encryption — for prostitution ads. “There were thousands and thousands of ads every day posted on these [Deep Web] forums,” said Shen. Even with the XDATA tools in hand, “it was hard to figure out whether you were looking for an ad for a private contractor versus somebody who is trafficking in large numbers of people,” Shen said. “We realized we had more than just a data analytics problem.” It was a problem in searching, acquiring, and filtering data — a problem, in this Deep Web case, that commercial search engines weren’t equipped to solve.

Shen managed the follow-on DARPA program, Memex, through most of its life cycle. Within six months, the team had fielded an initial capability for users in the New York/New Jersey metropolitan region. By the time the program had concluded in 2018, more than 40 law enforcement agencies, several of them in overseas countries, were using Memex software — an



The Long Range Anti-Ship Missile (LRASM) was rapidly developed through, in part, the use of a specialized joint office devoted to the transition of a suite of technologies into a new capability. After a rapid period of development, LRASM transitioned to Naval Air Systems Command (NAVAIR) in 2016.

open-source, downloadable set of tools for searching and analyzing data from the Deep Web.

One of the presenters at the 2015 *Wait, What?* forum in St. Louis was John Temple, chief of the Human Trafficking Response Unit for the New York County District Attorney’s Office. The Memex software, Temple said, allowed his investigators to thwart the tactics often used by traffickers — having young women hold up signs with contact information in a photograph, for example, rather than including the information in searchable text — and build a case that led to the conviction of a man who had been prostituting young girls for more than five years.

Memex is just one example of a technology transition that extends beyond the U.S. military. Some years ago, when the U.S. Air Force Space Command articulated a need to rapidly detect small objects in space — both the proliferating microsatellites, or CubeSats, being launched by the United States and other entities, and the increasing accumulation of space debris orbiting the earth — DARPA launched a program, the Space Surveillance Telescope (SST), which would be designed to overcome the small field of view and other limitations of conventional telescopes. Operational since 2011, the SST now offers viewers a “windshield” view, capable of detecting 10,000 softball-sized objects, in a field of view the size of the continental United States, within seconds.

Dr. Lindsay Millard, an expert in space-based sensing who joined DARPA in 2014 as a program manager in the Tactical Technology Office, describes the SST as a suite of technologies — mirrors, lenses, small-object detection algorithms, and faster search operations, among others — that ushered in new capability. “What the technology allowed you to do,” she said, “was maintain a very, very



wide field of view, but at the same time a very high resolution. The curved focal plane that DARPA and the MIT Lincoln Laboratory developed enabled that.”

By the time DARPA formally transferred the SST to the Air Force in 2016, the National Aeronautics and Space Administration (NASA) already had begun capitalizing on the SST’s capability to help provide warning of asteroids and other near-Earth objects. NASA has used the SST to observe millions of asteroids, discovering 3,600 new ones in the process, along with 68 near-Earth objects — making the telescope the most prolific tool for asteroid observation in the world. While it was designed specifically to track threats to the Air Force’s space-based assets, said Millard, “I think once a technology becomes available and people see what it can do, you end up finding a lot of other uses that you didn’t anticipate.”

The main lever that tipped the SST from a DARPA-led design and construction program to ownership and operation by the Air Force, Millard said, was a 2013 Memorandum of Understanding between the U.S. Secretary of Defense and the Australian Minister of Defense to locate the telescope in Australia, part of a joint initiative to increase space situational awareness over the Southern Hemisphere. The SST is currently being upgraded and prepared for installation at the Harold E. Holt Naval Communication Station near Exmouth, Western Australia.

As the SST and other programs illustrate, there is no fixed formula for transitioning a DARPA technology into the field. The mode of transition is often effectively discovered as program managers, DARPA transition professionals, liaisons, service branch officials, and the investigators whose work is solicited or supported by DARPA, collectively move a technology forward. In many transitions, such as the recent transfer of the LUKE arm to the commercial sector, DARPA will develop technologies to a point where the risk becomes low enough for civilian entities to take over development. In these cases, private capital and commerce can drive further advances and efficiencies that may, down the road, facilitate incorporation into military systems. Often, the interplay between the military and civilian researchers is almost symbiotic, with each influencing the other as a technology matures. Much of the technology that the world takes for granted today — the internet, the graphical user interface, the computer mouse, and handheld GPS receivers, for example — can be traced, at least in part, to DARPA-funded projects of decades past.

DARPA’s Space Surveillance Telescope (SST) program was designed to rapidly detect small objects in space. Transferred to the Air Force in 2016, SST has also been used by NASA to discover more than 3,600 new asteroids, including 68 near-Earth objects. SST will be upgraded and installed at a new facility in Australia in the near future.

As Masiello pointed out, tactical grassroots programs such as Remote Advise and Assist (RAA) — which led to a situational awareness and decision-assistance tool for ground troops — in contrast to strategic big-picture programs such as ACTUV and LRASM, illustrate DARPA’s ability to receive and act on requests for capabilities that originate anywhere within the military: from both the bottom and the top. [For more on RAA, see “DARPA at the Tactical Edge”] “Along with the service LNOs, we have insight across the entire portfolio, and are able to match technology to requirements,” he said. “On the one hand, with LRASM, we had a four-star admiral issuing an Urgent Operational Need statement that echoed throughout the department, and DARPA responded to that. On the other hand, we had lower-level discussions with warfighters in the field, to determine what they needed from the RAA, and we responded quickly to address those requirements.”

It’s easy to measure whether a DARPA transition can be called a success, said Murphy: “When the last DARPA dollar is spent and somebody else is investing significantly — in further research, utilization, buying down the technological risk, or making it more user-friendly — then we consider that a transition success. Sometimes it ends with people sort of liking it and thinking it could be useful someday, or thinking we’re pushing the technology envelope too far.” Programs that end in this way, understandably, don’t make much noise, but they’re evidence that DARPA is staying true to its mission of pursuing the limits of what technology can do.

Whether they begin at the bottom or the top, and whether they transition or not, all DARPA programs have one thing in common: Even before they’re launched, their program managers develop strategies for transitioning anticipated results into the real world. From their first day on the job, DARPA program managers envision that the new extraordinary technology ideas they expect to make real ultimately will make a difference.

DARPA'S 60-YEAR SPACE ADVENTURE

A maker of the Space Age continues its expansive take on space technology

By Edward Goldstein

Four months after the Soviet Union shocked the United States on Oct. 4, 1957, by launching Sputnik 1, the first Earth-orbiting satellite, and three months after they doubled down on this Cold War victory in the space race by launching Sputnik 2 (carrying the dog Laika), and eight days after the United States finally got into this highest-stakes Cold War game when the Army Ballistic Missile Agency satellite Explorer 1 made it into orbit, our country established its first real space agency. It was not NASA. It was ARPA, the Advanced Research Projects Agency, which was established via a Department of Defense (DOD) directive on Feb. 7, 1958,¹ five months before NASA was established.

“We realized at that [first Sputnik] moment, if somebody launched something that could go all the way around the world, they could launch something that could go half-way across the world,” said former DARPA Director Anthony “Tony” Tether. “And Russia was a nuclear power. That changed the whole dynamic of our national security. President [Dwight D.] Eisenhower said, ‘How the heck did a third-world country beat us into space?’ He investigated and found that in the Defense Department, each of the services had a space program, but they were a low priority. There was no sense of urgency. So, he created ARPA to be the agency, whose charter was very simple: to prevent technological surprise.”

A Legacy of Progress

As specified in DOD Directive 5105.15, ARPA also had the charge to expand the frontiers of technology and science. By implication, therefore, the new agency had the job of bringing technological surprise to the other guy. Throughout its history, DARPA has done just that, sometimes with ambitious and audacious space projects that have led to tremendous technological advances in communications, weather forecasting, early warning for missile defense, reconnaissance, and geolocation. One of DARPA's strengths over the years is that it has used its position as a technology driver, rather than as an operational space agency, to develop new systems and concepts that are in turn picked up by other organizations, in both the military and civilian sectors. “We want to do things that we know will have a receiver at the far end,” said Jerry Krassner, senior adviser to DARPA's Tactical Technology Office. “We're not just here to throw something over the roof and hope that they will take it. Hopefully, our transition partners are with us along the way. We're very integrated in a coordinated way with our mission partners in the services for that reason.” In government-speak, this act of handing off a technology development effort often is referred

to as “transitioning” and those receiving the technology are thought of as transition partners.


Eighteen Defining Months

This organizationally defining role already was in ARPA's future during its critical first 18 months of operations. On Jan. 9, 1958, in his State of the Union Address to Congress, Eisenhower laid out the justification for ARPA's founding: “In recognition of the need for single control in some of our more advanced development projects, the Secretary of Defense has already decided to concentrate into one organization all the anti-missile and satellite technology undertaken within the Department of Defense.” Initially, ARPA had authority over the nation's launch vehicles, including the concept for a large booster, which would eventually become the Saturn V rocket that would launch Apollo astronauts on their voyages to the moon. The young agency also oversaw several satellite projects in development. These included Transit, the first global satellite navigation system and forerunner of the Global Positioning System, and TIROS, the world's first weather satellite.

“In 1958, everything was up for grabs, and [the U.S.] could well have ended up with an Air Force-dominated space program instead of having a NASA, or having ARPA become the [initial] national space program and try to coordinate between the services,” said Michael J. Neufeld, senior curator of the Smithsonian National Air and Space Museum's Space History Department. But events in 1958 and 1959 dictated dramatic change in how the nation's space programs would be conducted. First came Eisenhower's insistence that America have a civil space agency, NASA.

NASA served two key purposes, according to Neufeld. “One was internal,” he said. “This interservice rivalry [between the Air Force and Army

An Atlas-B missile being prepared to launch the SCORE satellite, the country's first communications satellite, from Cape Canaveral, Florida.





DOE PHOTO

over roles and responsibilities for space programs] was very alienating to a lot of people. So, how do we coordinate that? And number two was the external image of the United States in the Cold War. The Soviet propaganda was relentless that we were the militarists, we were the imperialists, we were the aggressors. The optics of a civilian space agency looked good." Defense Secretary Neil McElroy further pursued the issue of coordinating space projects in September 1959 when, influenced in part by ARPA's first chief scientist, Herbert York, he decided to force the issue of why the Army was managing a major launch vehicle program in Huntsville, Alabama, led by none other than German rocket scientist Wernher von Braun, father of the Nazis' V-2 ballistic missile.

"The net result was to clarify things," said Neufeld. "The Army transferred its major assets to NASA and the Air Force became the dominant player in military space programs. I think that's when ARPA began to transition into something else, into being the small agency for exotic ideas."

New Satellite Capabilities and America's Monster Rocket

Amid all this organizational turmoil in the infancy of the Space Age, ARPA served an important function by providing a greater focus to existing space efforts — with the services' help, funding and managing pioneering satellite projects — and influencing a key change in the nation's launch-vehicle program that proved critical to America's quest in the 1960s to land men on the moon.

In July 1958, ARPA assumed control of a project being pursued by the Army's Signal Research and Development Laboratory (SRDL) at Fort Monmouth, New Jersey, to develop the world's first communications satellite, SCORE (Signal Communications by Orbital Relay Equipment). And less



LEFT: Scientists prepare the first navigation satellite, Transit, for launch. Piggybacking on the spiral-adorned Transit was the Naval Research Laboratory's Solrad 1, a satellite with the dual purpose of measuring solar radiation and listening in on Soviet radar emissions. ABOVE: In the days before clean rooms, technicians work on the nation's first weather satellite, the Television Infrared Observation Satellite (TIROS).

than five months from project go-ahead, on Dec. 18, 1958, an Atlas rocket boosted SCORE into orbit and along with it the tape-recorded voice of Eisenhower conveying his holiday-time message — "America's wish for peace on Earth and good will toward men everywhere" — to ground stations around the globe. This success helped provide a model for ARPA's management of space projects, as Owen Brown, Fred Kennedy, and Wade Pulliam wrote in their 50th anniversary recounting of the agency's space efforts, *DARPA's Space History*:

"As with the other projects it began, ARPA took an agile management role in the SCORE system development, with the Air Force being handed the responsibility for the Atlas, built by General Dynamics, and SRDL for the payload, built by RCA. This approach of maintaining low overhead was, and still is, a vital part of the success of the agency, especially during the early years when so many space systems were being developed and demonstrated. In what would become standard procedure, the agency concentrated on developing and fostering innovative concepts, as well as making key project funding decisions, while leaving the burden of facility and workforce maintenance, system development and production, and operations to industry, research labs, and the military services. At first this decision was made mostly for the bureaucratic reason of limiting the objections of the services at having their space programs transferred to ARPA. In the long run, however, it has served the agency well in keeping it



LEFT: Engineers prepare one of the Vela Hotel series of satellites, whose development was supervised by DARPA, for on-orbit nuclear test detection. BELOW: The Saturn I SA-4 flight launched from NASA's Kennedy Space Center in 1963. SA-4 was an uncrewed test flight of the Saturn I booster. The Saturn IB served as a test-bed rocket for the larger and more powerful Saturn V that would eventually carry the first humans to the moon.

nimble, capable of changing directions quickly and able to manage many efforts rather than just a few."²


More successes were to come quickly. For example, ARPA was involved in the funding of CORONA, the world's first photo-reconnaissance satellite, managed by the CIA and Air Force. Though riddled by a series of early failures, the ultimately successful sequence of spy flights helped America's leadership understand that Soviet intercontinental missile capabilities were less extensive than feared. Because the CORONA satellites had a very narrow time window in which their high-resolution cameras would take images of strategic targets inside the Soviet Union — images that were recovered in the upper atmosphere by high-flying aircraft capturing parachuting payloads containing film canisters — planners needed to know in advance if cloud conditions





LEFT: Two soldiers test early models of GPS manpack receivers in 1978.

BELOW LEFT: The Orbital Express Program's NextSat "service" satellite, seen from ASTRO, a robotic repair and maintenance spacecraft.



would allow for each reconnaissance mission to go forward. Hence, ARPA, in July 1958, tasked RCA with developing the first weather satellite, TIROS (Television Infrared Observation Satellite). Launched by NASA in April 1960, TIROS quickly demonstrated the capability of satellites to assist in weather forecasting and provide advanced warnings of severe storms in coastal areas. As Brown, Kennedy, and Pulliam wrote, "Beside the development of the internet (and maybe satellite communications), few other ARPA programs have had such a profound impact on the entire world."³

The concept of using signals from satellites to improve the geo-localization accuracy (within 0.1 nautical mile) of Polaris missile-armed submarines prompted ARPA's work on the world's first navigation satellite, Transit, with the Johns Hopkins Applied Physics Laboratory (APL). This satellite launched in April 1960, just 18 months after ARPA's project approval, based on the idea

of APL researchers William Guier and George Weiffenbach, that a location on Earth could be calculated with knowledge of the precise orbital trajectory of a satellite. The Transit system significantly enhanced the confidence of war planners that the U.S. Navy's ballistic missile submarines, one of three legs of the U.S. nuclear triad, would hit their assigned targets. The Transit system remained in operation until 1996, when the current Global Positioning System (GPS) took over, providing position data to military and civilian ships, enabling the land surveying work of the Defense Mapping Agency, and ultimately transforming how much of the world's population gets itself from Point A to Point B.

The next major ARPA space project, Vela Hotel, had profound consequences for reducing the risk of nuclear war and for protecting millions of people from the radiation fallout from atmospheric tests of nuclear weapons. Eisenhower was motivated to pursue a nuclear test ban agreement with the Soviet Union by the negative international publicity from a 1954 thermonuclear test that exposed Marshall Island natives to high levels of radiation as well as his hopes for a reduction in tensions with the country's Cold War nemesis. Toward that end, the president directed ARPA in 1959 to develop the technologies needed to detect nuclear tests. ARPA's project Vela ("Watchman" in Spanish) consisted of space (Vela Hotel), ground (Vela Sierra), and seismic (Vela Uniform) assets that would detect and monitor nuclear tests conducted in space, the atmosphere, or underground.

"Remember, the general charter of ARPA was ... to prevent technological surprise," noted Tether. "We were trying to negotiate the atmospheric test ban treaty and we realized we had no way to verify whether someone was violating it." On May 1, 1960, the ante upped dramatically when the Soviet Union shot down a U-2 spy plane piloted by Francis Gary Powers. "Only because of Vela Hotel did we sign up for the atmospheric test ban treaty in April 1963," Tether said. "That was a big deal that gave life back to ARPA." The six pairs of Vela satellites placed in orbit from 1963 to 1970 not only helped verify the treaty, but their gamma-ray sensors used to detect a nuclear test in space helped start the field of gamma-ray astronomy, which is critical for the identification of supernovae and measurements of the early universe just after the Big Bang.

In its infancy, ARPA also had an important role for the direction of the U.S. civil space program. For a short while in 1958, ARPA had control of the

U.S. AIR FORCE PHOTO

DARPA IMAGE

U.S. launch vehicle program. “At this point of time, the biggest obvious lack of United States capability was a large booster,” said Neufeld. Despite concerns regarding why the Army had a large space effort underway in Huntsville, Alabama, ARPA played an earlier role in booster development there by supporting, in the summer of 1958, the development of what became known as the Saturn 1. Still led by von Braun, the Army engineering team used the funding to develop the concept of grouping some boosters based on those used in Redstone and Jupiter rockets into a competent booster with eight engines, Neufeld said. This action helped to accelerate development for the Saturn V booster and, thereby, opened the way for the Apollo program. ARPA also funded support of ground testing and the study of launch facilities until booster development was transferred to NASA in 1959.

Navigation Satellites and Cigarette Packs

On the heels of ARPA’s successful navigation satellite program, Transit, the Navy’s development of the Timation satellites, which proved the feasibility of placing accurate clocks in space, among other parallel developments, gave the impetus to a joint service agreement in 1973 to create a Defense Navigation Satellite System, based on a constellation of NAVSTAR-GPS (Navigation System using Timing and Ranging-Global Positioning System) satellites. The original case for the Navstar-GPS was based on providing greater location accuracy to the nation’s strategic bombers and ground- and sea-launched missiles, but it was not a huge leap to imagine expanding the system to provide precision navigation support to U.S. forces on the ground. ARPA had a critical role in developing the receiver sets that allowed our troops to make full use of this game-changing technology.

“I used to go to a lot of military exercises where they would have war games when I was running the Strategic Technology Office back from 1982 to 1986,” recalled Tether. “I came back from one, and I told the guys what I saw. And I said, ‘I found out one thing. The issue is not where the enemy is. We have all kinds of sensors that know where the enemy is. The real problem is we don’t know where our own guys are.’ At that time, we were just starting to have a few GPS satellites up, and the Army had made a GPS utilization capability, but it was a big backpack, it took 20 minutes to get a location, and only the person with the backpack knew wherever the heck he was. At a 20-minute rate, it wasn’t very tactically helpful. So, I told my guys, ‘Hey, this is a problem. What we really want to do is get a GPS capability that every soldier can carry with him in his pocket.’”

This is where advances in the minuscule technology of microelectronics comes into the big-tech story of GPS, Tether explained. Sherman Karp, who was a program manager in DARPA’s Strategic Technology Office, had an interest in the semiconductor material gallium arsenide (GaAs), which had advantages over silicon for applications that might require powerful, compact radio sources. To promote this exotic material, Karp knew he had to find an application for it. Tether recalls a meeting in which Karp said, “Look, Tony, I can make that [GPS receiver] backpack be really small.”... And I said, “Well, how small can you make it?” And he said, “How small do you want it to be?” Tether pulled out a pack of Marlboro cigarettes he had in his pocket and challenged Karp to make the receiver small enough to fit into the pack.

A week or so later, Karp admitted that he couldn’t deliver on that demand, but he had a counter-offer. “We had a secretary named Gilda, who smoked Virginia Slims, whose pack was one and a half times taller than

the Marlboro pack and a bit more compact,” and Karp promised that he could make a receiver of that size, Tether recalled. Karp followed through. By leveraging the lower-noise circuitry gallium arsenide made possible, GPS receivers made with the materials could detect satellites at a greater distance. The firm Rockwell Collins manufactured the first chip sets for the miniaturized “Virginia Slim” GPS receivers. “When you pick up your phone that has a GPS capability, it all goes back to this one guy, Sherman Karp, who said, ‘I can make this thing as small as you want it to be,’” said Tether.

For Tether, another favorite eureka moment came with the development of Orbital Express, a joint effort with NASA’s Marshall Space Flight Center to demonstrate autonomous robotic servicing and refueling of satellites. “The problem we have with spacecraft is once they are up there, it’s hard to do anything with them,” he said. “If something breaks down, you’re stuck; you lose the whole thing.” DARPA’s answer was to “create a spacecraft that could literally go up and repair a satellite on orbit, either change out the batteries or put something new into it,” Tether said. “And out of that demonstration, called Orbital Express (2007), started a whole industry.”

Orbital Express consisted of two spacecraft: Autonomous Space Transfer and Robotic Orbiter (ASTRO), and NextSat, the client or “servicee.” After launch on March 8, 2007, the mission spent 135 days on orbit, conducting seven separate rendezvous and docking procedures from distances up to 400 kilometers. While docked, ASTRO and NextSat carried out 14 refuelings, six battery transfers, and a flight computer changeout, proving the feasibility and utility of unmanned repair, refueling, and maintenance of satellites in Earth orbit.

“No other DARPA program in the recent past may have more profound impact on the future of space access and infrastructure than Orbital Express,” Brown, Kennedy, and Pulliam emphasized in *DARPA’s Space History*. “With Orbital Express, DARPA offered a new way of thinking about the design and operation of future space systems: not only can serviceable satellites offer unmatched capabilities, they also provide decision-makers and warfighters with the ability to change or modify these capabilities at any time in their life cycle, as well as the ability to continue to perform the intended mission despite changes to the operating environment.” Today, several companies are building unmanned spacecraft that will service and repair satellites on orbit.

DARPA and the Commercial Space Revolution

DARPA’s insistence on thinking of better ways to conduct space missions has helped inform its approach to its current activities. Born out of a frustration with large defense satellite projects for missile warning and weather forecasting going way over budget and well beyond schedule, DARPA officials like Fred Kennedy, the Director of the Tactical Technology Office, are leading the charge for dramatic changes in the way we approach the space component of the national security mission. “The surface symptoms of the problem are the extreme costs and long schedules we’ve been enduring to put some of our systems up,” said Kennedy. “This is driving risk-aversion into the culture to a level that is unsupportable, which is feeding our cost and schedule problems, and what’s worse is killing us in terms of our ability to innovate and respond to emerging threats. It’s opening up both our acquisition and tactical decision loops such that our adversaries now can build countermeasures more quickly than we actually field systems.” [See “Enterprise Disruption,” p.68]



DARPA's Launch Challenge tasks competitors with making two launches from different locations in a matter of days.

The solution, says Kennedy, is to “upend the current space enterprise, the order of battle, the architecture such that we will incentivize risk-taking, innovation, and closing the decision loop rather than the reverse.” Specifically, he sees DARPA's future space activities being built on a foundation of “commoditized” and “productized” spacecraft and rapid response launch capabilities. In that sense, DARPA is continuing a technology trajectory with roots, including ones that took hold in the early 2000s in the Rapid Access to Space program. That effort included the first government contract for the commercial space pioneer, Space Technology Corporation (SpaceX), to launch its Falcon 1, which the agency was eyeing as a potential pathway for getting small payloads into space both quickly and cheaply. The two launches executed under DARPA contract provided hard lessons — one launch failed at 25 seconds and the second would have reached orbit had the rocket's second stage not cut off prematurely — and valuable operational analyses that SpaceX subsequently leveraged for a Falcon 1 success on its fourth try and for its subsequent development of its workhorse Falcon 9.

More than ever now, DARPA is embracing the goal of access to space with unprecedented ease, versatility, scale, and affordability. “I want to build simple, very cheap things that I can mass produce,” said Kennedy. “And I think we can do that because now we have the commercial vendors — SpaceX, OneWeb, Telesat, Boeing, Samsung — who are off and running trying to figure out how to do mass production of small satellites. I think the commercial sector is actually going to get out in front of us on this one, and show us how to do the Model T of spacecraft. What we need to do is figure out how to build a good enough payload that we can mount on a good enough [satellite] bus and go do the missions that we do today.”

In a bid to perpetuate DARPA's culture of driving technological disruption, the agency is pushing several ongoing DARPA space initiatives:

Blackjack. This program is designed to develop space technologies that demonstrate an extensive smallsat constellation in low Earth orbit

(LEO). Blackjack, said Krassner of DARPA's Tactical Technology Office, “is designed to take advantage of the emerging commercial LEO constellations. The question is: Can the military adopt these commodity buses to put military payloads on them and operate them as a distributed network? The advantage would include cost savings from using these commoditized buses. They would also provide resilience, since everything we now put on a big platform we could have instead in a distributed architecture. There would be many more targets an adversary would have to eliminate. Systems coming out of a successful Blackjack would provide the opportunity to refresh and update technology on a much more frequent basis. We think this is a potential disruptive architecture concept for national security space.” Added Kennedy, “Blackjack gets us to a different world where we are no longer risk-averse. If we can get the payload community to come along so that they understand they will have to produce mass-reproducible systems, I think we will have entered a new era. Rather than biting our nails every time we are sitting there for launch, we should be saying, ‘I can handle losing six of 12 or 15 of these. I'll just build and launch a few more.’”

DARPA Launch Challenge. In 2004, DARPA kicked off the idea of crowdsourcing technology development by offering “Grand Challenges” to all comers. Autonomous vehicle technology — that is, driverless cars — was the focus of DARPA's first Grand Challenges. Seeking to tap the rapid pace of development in the commercial launch industry, DARPA this year floated an offer of more than \$10 million in prize money for a competitor delivering a sufficiently flexible and responsive launch capability. “Right now, getting to space is a bottleneck, and oftentimes small payloads have to piggyback on large payloads and rockets or ride-share,” observed Krassner. “That means they are going whenever the big guy is ready to go and wherever he wants to go. But with small rockets, you could basically buy your own rocket to go on your own schedule. It's kind of like the difference between a bus and a taxicab.” Under the DARPA Launch Challenge, teams will receive exact details on the payload in the days before each of the challenge's two launch events, with only a few weeks' notice about the location of the first launch site. The winning team will have to

DARPA's Hallmark program aims to provide U.S. senior leadership the ability to manage U.S. space assets in real time, as this artist's depiction might suggest.

successfully deliver their payload to LEO from two launch sites within a matter of days.

If the dual aims of Blackjack (satellite proliferation) and the Launch Challenge (rapid response launch) succeed, said Kennedy, "you would have an incredible deterrent feature against any kind of bad behavior, because you have thousands of assets. They are low-value assets, which means there is no single asset that makes sense to shoot at. You would have to shoot them all. At the same time, you would have a reconstitution capability, so that even if someone decided they want to take potshots at you, you respond by simply putting more of these low-value assets up. I'm trying to force the potential adversary into the realization that it's not worth bothering to take the shot in the first place."

Robotic Servicing of Geosynchronous Satellites (RSGS). Building on the success of Orbital Express, DARPA is pursuing technologies for servicing spacecraft in geosynchronous (GEO) orbits through a mix of highly automated and remotely operated (from Earth) robotic systems.

Experimental Spaceplane (XSP) Program. The U.S. aerospace community previously has undertaken expensive and unsuccessful efforts to develop a spaceplane that can be essentially launched on demand to LEO. Now, DARPA leadership judges that the relevant technologies have matured sufficiently to try again. Said Krassner: "It might launch vertically, it might launch horizontally. It does whatever it is supposed to do in space, deploys a satellite or whatever, and then comes back and lands at an airport like an airplane and gets refurbished just like an airplane at the gate." The program's goal is to design an X-Plane capable of 10 flights in 10 days and that can be transitioned to the Air Force, Navy, and commercial sector.

Hallmark. This battle management tool is designed to provide U.S. senior leadership the ability to effectively manage space assets in real time. "Think about a control center where the general is in charge and he gets all this space situational awareness data, but he has no real way to make heads or tails of it," said Krassner. "The Hallmark program will provide visualization tools and decision support tools to help him make informed decisions based on the various data sources that he has available."

Radar Net. This program is developing lightweight, low-power, wide-band capability for radio frequency (RF) communications and remote sensing for a space-based platform. "A number of organizations including commercial ones are developing synthetic aperture radar satellites" that could provide a basis for the goals of Radar Net, noted Krassner. "The question is, since synthetic aperture radar has so many applications and benefits, is there a way we can leverage advanced technologies to provide a large synthetic aperture radar in a small package, and use a small launch vehicle to launch a small payload, which is much cheaper."

Planar Imager. This program involves space-based electro-optical sensors for intelligence, surveillance, and reconnaissance to replace conventional telescopes. "If you look at things like the Hubble Space Telescope, a large fraction of the weight of the telescope is in the mirror," explained Krassner. "And of course, in space, weight and size translate into cost, launch vehicle constraints, etc. If we can take the mirror out



and replace it with, in effect, a flat plastic component that might be meters across but very thin, and it doesn't take years to grind [like a conventional telescope mirror], we would be able to put up a much larger telescope at a much quicker pace at much less cost. That combination is potentially disruptive. So, the question is whether the technology in advanced materials can support these objectives."

Of Starships and Big Dreams

Because of DARPA's penchant to dream big, not all of its concepts have come to fruition. In the late 1950s and early 1960s, for example, ARPA funded Project Orion, an audacious concept to develop an interplanetary spacecraft propelled by a series of nuclear bomb explosions, which never got off the ground due to concerns about fallout from its propulsion. And in 2011, DARPA joined with NASA's Ames Research Center to provide grant funding for researchers to create a business plan that can last 100 years to help foster the research needed for interstellar travel. The one-year project resulted in a private organization under the same name that is now attempting to carry on this work.

"Technology moves fast," said Krassner. "It's not a bad thing for DARPA to be looking at the long term ... And where we might be thinking in a Wall Street quarter-to-quarter basis, every now and then it's good to step away and take a long view that helps get you some out-of-the-box thinking. And that's what we are all about: Out-of-the-box thinking, disruption, and speed."

"DARPA is refocusing on space," Krassner noted. "The pendulum has swung over the years as to how much we are doing in space, and we now are ratcheting up our space investments as a result of advanced technology opportunities and the growing importance of space for military and civilian applications. DARPA is the right place to try things, because we move fast, we do not have lifetime programs, and we are very cognizant as we start things about who might be a transition partner. Space is on the upswing here."

"The future of space at DARPA is looking bright over the next 20 years, because we are about to upend the space status quo entirely," added Kennedy. "We are going to do that by taking away all the incentivization that currently exists for risk aversion. We are going to change the architecture. We are going to change all the assumptions on which we base our space capabilities. We are going to flip it on its head. The big win is culture change, and if we can affect that throughout all the systems we build, everything will be different. We are in the midst of some major upheavals that are going to change everything."

1. In 1972, ARPA would become DARPA when it added a D (for Defense) to its name.
2. Brown, Owen, Kennedy, Fred and Pulliam, Wade, "DARPA's Space History.
3. Ibid.

ARTIFICIAL INTELLIGENCE TO ACCELERATE SCIENCE

How *in-silico* minds will open new pathways to discovery


By Paul Cohen

The technology of science is as old as science itself. Astrolabes and microscopes, image-displaying tachistoscopes and gene sequencers, mass spectrometers and atomic clocks, molecular tweezers, and the gene-editing tool CRISPR have all accelerated science by making the invisible visible and the uncontrolled do our bidding. Compared with these physical technologies, information technologies for science do different work: They store and manipulate data, they automate tedious tasks, they facilitate replicability and sharing, they learn, and one day they will assist scientists as ably as a good graduate student.

Science has been a frontier for artificial intelligence (AI) since the 1970s. The prevailing view was that AI should do what our most notable scientists do, so autonomous scientific discovery was the highest prize, and the prevailing methods of AI – various forms of generate-and-test – were presented as theories of the scientific discovery process¹. For example, the DARPA-sponsored DENDRAL project generated graph models of molecules in organic chemistry and tested them against mass spectrometer data, and a later program called CONGEN was able to discover all chemical graphs that satisfy empirical constraints². At roughly the same time, other researchers applied rudimentary data-mining heuristics to re-discover Kepler's Law and other physical laws in data³.

If these early efforts failed to convince us that AI could “do science,” it was probably because they did so little of what scientists do: They didn’t read the literature, go to seminars, discuss theories with colleagues, prepare samples, design and run experiments, clean noisy data, or test hypotheses. They focused on the “aha moment” of discovery, not on the daily work of science.

The Big Mechanism program aims to develop technology to read research abstracts and papers to extract pieces of causal mechanisms, assemble these pieces into more complete causal models, and reason over these models to produce explanations. The domain of the program is cancer biology with an emphasis on signaling pathways.






Even now, the heroic theory of science, which holds that scientific discovery is the product of individual genius, influences discussions of AI approaches to science. For example, in 2014 Hiroaki Kitano, whose several research affiliations include president and CEO of Sony Computer Science Laboratories, proposed “a new grand challenge for AI: to develop an AI system that can make major scientific discoveries in biomedical sciences and that is worthy of a Nobel Prize.”⁴ In fact, to date no AI system can recognize a significant scientific result as such, and all “scientific discovery” systems are carefully managed by humans.


Instead of asking how humans can build AI machines capable of making scientific discoveries, it might be more productive to ask how machines can facilitate scientific discovery by humans⁵. This was the premise of DARPA's Big Mechanism and World Modelers programs, which I had designed in my recent stint as a DARPA program manager to investigate if and how AI can accelerate science as well as create new opportunities for AI research.

The Big Mechanism program was designed to develop technology to help humans build causal models of complicated systems. The program focused on the complicated molecular interactions in cells that, when they go wrong, result in cancer. Cell-signaling pathways are sequences of protein-protein interactions that transmit information to the cell nucleus and determine cell fate. The literature on cell signaling is vast, and each paper describes just a few signaling interactions. So the Big Mechanism program developed technology to read the literature, assemble individual results into entire pathways, and help scientists explain the effects of drugs on pathways.

In an experiment in 2016, machines were able to explain 25 known drug-pathway interactions. Given a previously published model of 336 relevant genes (each of which encodes a protein) the machines used natural-language-understanding technologies to discover and read 95,000 journal articles, from which they extracted nearly a million causal assertions about protein-protein interactions. These were filtered and assembled into a single plausible signaling model that not only simulated the dynamics of protein concentrations but also explained all the drug-protein interactions.



TOP: Joshua Lederberg inaugurates the Stanford University Medical Experimental Computer (SUMEX), a computer designed to encourage the application of artificial intelligence in medicine and to provide a computer facility for sharing and interaction among researchers nationwide. Lederberg was one of those who worked on the DENDRAL project, an early effort to use artificial intelligence to further the scientific process. ABOVE: Cover sheet of the Proposal for Continuation of the Stanford Artificial Intelligence Project and the Heuristic DENDRAL Project (June 1969).



Using cloud-computing clusters, the whole process took less than a day.

These results showed that machine reading and model-building can accelerate science in the sense that no human can do what the machines did: No one can read 95,000 journal articles or process a million assertions into a causal model that explains empirical results, even if they devoted years to the task. The machines worked alone, but researchers already are demonstrating more interactive versions of the technologies⁶.

Perhaps the most significant contribution of the Big Mechanism program has been to subvert the dominant paradigm of Big Data with its emphasis on exploitable correlations. All other things being equal, scientists prefer causal, explanatory models to opaque, predictive models based on correlations. The Big Mechanism program – its name a poke in the eye of Big Data – showed that machines could build causal, mechanistic models of cellular processes such as tumorigenesis. One day, perhaps, these technologies will lead to machine-assisted hypotheses of how to interrupt or enhance cellular processes.

While causal knowledge is the highest prize in science, it is quite difficult to extract from data. Contrary to popular belief, it is possible to find causal relationships in correlational data, but the algorithms for doing so are computationally expensive and leave residual uncertainty about whether one thing truly causes another. In contrast, it is not very difficult to find assertions of causality in text. The linguistic constructions can be arcane (e.g., “mitogens stimulate cell division by inhibiting intracellular negative controls”), but researchers have been making steady progress toward extracting causality from text⁷. Eventually machines will read all of the causal assertions in the scientific literature and check them against each other and against available data. Imagine machines that read millions of papers and find odd results that don’t fit the zeitgeist or don’t accord with data. Are they due to fraud or incompetence, or are they showing us something new and unexpected? Human scientists ask themselves these questions as they plod slowly through vast literatures; imagine how machines might accelerate the process.

The Big Mechanism program has a successor at DARPA called World

The Australian government analyzed land use by integrating nine component models, but this effort took approximately one person-century to accomplish. Intelligent machine assistance could radically shorten the timeline for such a project.


Modelers. Here the machine’s task is not to build models *de novo*, but to support humans as they assemble huge, complicated workflows of many extant models. The need for “mega-models” becomes palpable when we realize that many consequential problems, such as radicalization and intolerance, and food and energy insecurity, involve interacting systems.

Consider the food insecurity example. Food insecurity has many causes, from poor soil to political instability, from the El Niño cycle to economic migration. Scientists have been developing models of these individual elements for decades, but superhuman effort is required to link them together into “mega models”: huge, complicated workflows within a common software environment. For example, the Australian government analyzed land use by integrating nine component models – of energy, water, and markets, among others – but this splendid effort took roughly one person-century of work⁸. To accelerate this kind of analysis through intelligent machine assistance, DARPA created the World Modelers program.

These examples suggest that AI will accelerate science in several ways, such as reading and assembling fragmentary results spread widely over literatures, integrating legacy models in common computational frameworks, automating *in silico* experiments, and even designing experiments and controlling the robots that carry them out. Much of this is “good old-fashioned AI,” not contemporary data science. At present, big data and machine learning play roles such as finding associations that might be causal (e.g., associations between genes and phenotypes) and learning computationally efficient approximations to expensive legacy models. But science depends on theories and data, and, importantly, on what people assert about theories and data in published literature. This suggests that future data-science technologies should expand their scope to embrace the interplay of theories, data, and literature.

For all its promise, AI has yet to recreate even the intellectual functions of a good research assistant. Nor is it likely to unless it tackles some problems that are getting in the way. As we review some of these, it might seem incredible that they haven’t been solved, but AI is a field in which seemingly easy things can be very difficult.

What are all those things and processes? Scientists refer to things and processes by names such as “p53” and “desertification.” In general, one thing can have many names and many things can have the same name, and even if things and names were in one-to-one correspondence, machines won’t necessarily know anything about what a name denotes. For example, “EBF” and “breast feeding” are names for very similar processes (“EBF” stands for exclusive breast feeding), but machines can’t know this unless they have access to dictionaries or ontologies that map names to formal descriptions of things. One might hope that someone has specified that “breast feeding” means, well, breast feeding, while EBF means breast feeding exclusive of other kinds of feeding. Ontologies record this kind of information and some fields, such as biology, have excellent



DARPA's World Modelers program intends to support people in assembling large, complicated workflows of many models in order to more quickly gain insight into problems such as food and energy insecurity.

the models well enough to use the output of one model as input to another, either directly or following some transformation. As noted, this understanding might require more semantic depth than machines have, but even when semantic issues are solved, gritty engineering issues remain: If the models are linked, meaning that feedback loops exist between the processes they represent, then they should run in a single computational environment. But this can be difficult if they run at different time scales or require very different amounts of computation. The technology of scientific workflows is progressing rapidly, but it isn't yet possible for machines, rather than humans, to build

ontologies for genes, proteins, drugs, and so on, but in general, scientific fields are poorly ontologized. In short, machines don't know what scientists are talking about.

Coreference. Humans refer to the same things and processes in different ways, and machines have trouble figuring out coreferences, or repeated references to the same thing or processes. A paragraph might begin "The phosphorylation of ERK" and end with "active ERK," but how are machines to know that phosphorylated ERK is often referred to as active ERK? How can machines know that "the first horse across the line" is "the winner"? The technology for dealing with coreference is improving, but every missed coreference is a lost opportunity for a machine to extend what it knows about a thing or process.

Semantic depth. Even when machines know about things and processes, they generally don't know much. Consequently, they can't answer questions that require semantic depth or nuance. For example, suppose it is important to estimate the per-capita income in a neighborhood, but there is no data source that provides this information. A good research assistant would think, hmmm, let's use sales records to find home prices and drone-based imagery to identify the numbers of cars in driveways, and let's put it all together with public tax records to estimate household income. Machines can't do this task (i.e., they can't invent proxies for missing data) unless they know what "income" means and know that other data are proportional to income. The techniques for developing and exploiting meta-data (the things we would like machines to know about data) are improving, but they have yet to incorporate semantically deep-enough nuance for machines to invent proxies.

Gritty Engineering. Models often have parameters that represent local conditions. Crop models, for example, need data about soil quality, sunshine, water, and other factors. A good research assistant might integrate a crop model with a soil model, a hydrological model, and weather model. The challenges would include understanding the parameters of

complicated workflows of many computational models.

This sample of challenges should not discourage anyone from developing AI technology to enhance the work of scientists. Indeed, accelerating science is so important that it should motivate basic AI research on these and other challenges, as happened in the Big Mechanism and World Modelers programs mentioned earlier.

The challenges of our century are systemic, but humans have difficulty modeling and managing systems. Whether we're modeling the molecular signaling pathways in cancer, the diverse factors contributing to food insecurity, or policies for land use in Australia, we find ourselves struggling with complexity. We have no choice about whether to recruit AI technology to scientific research: We must do it because we can't understand complicated, interacting systems without help. It is an added benefit that the vision of AI-accelerated science will drive AI research itself for years to come.

1. Langley, P., Simon, H.A., Bradshaw, G.L., & Zytkow, J.M. (1993). *Scientific Discovery*. The MIT Press.
2. <https://profiles.nlm.nih.gov/BB/A/L/A/F/>
3. Op. Cit. Langley, et al,
4. Kitano, H. "Artificial Intelligence to Win the Nobel Prize and Beyond: Creating the Engine for Scientific Discovery." *AI Magazine*, Spring, 2016.
5. Yolanda Gil, Mark Greaves, James Hendler, Haym Hirsh. "Amplify scientific discovery with artificial intelligence." *SCIENCE* 10 OCT 2014 : 171-172
6. Benjamin M. Gyori, John A. Bachman, Kartik Subramanian, Jeremy L. Muhlich, Lucian Galescu, and Peter K. Sorger. 11/2017. "From word models to executable models of signaling networks using automated assembly." *Molecular Systems Biology*, 13, 11, Pp. 954. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5731347/>
7. Gus Hahn-Powell, Dane Bell, Marco A Valenzuela-Escárcega, Mihai Surdeanu. "This before that: Causal precedence in the biomedical domain." <https://arxiv.org/pdf/1606.08089.pdf>
8. Australian National Outlook. https://www.csiro.au/~media/Major-initiatives/Australian-National-Outlook/CSIRO-MAIN_REPORT-National_Outlook_2015.pdf?la=en&hash=B7430AC9FD3C6665A53B9BDD1B2ED5356CB2FA3B

UNIVERSITIES IN SERVICE TO NATIONAL SECURITY

DARPA's call to academia

By William Regli

America's higher educational institutions and university research centers have been the cauldrons in which many of DARPA's innovations have been brewed. Over the past six decades since DARPA's beginnings, the agency has provided generations of academic scientists and engineers opportunities to put their talents to use in service of national security. This partnership between DARPA and academia has become more important as the threats to our society have become more diverse, complex, and sophisticated.

There is a long list of game-changing innovations that were both DARPA-sponsored and had university researchers in leading roles: the internet, self-driving cars, neuro-prosthetics, and speech technologies, to name just a few.

Consider the generation of electrical engineers and computer scientists who conducted their Ph.D. research under DARPA support during the Strategic Computing era of the 1980s as well as the countless other students who were influenced by these efforts. The first wave of artificial intelligence (AI) technologies came to life on machines developed by DARPA-sponsored researchers and ran on chips prototyped using the Metal Oxide Silicon Implementation Service (MOSIS), a DARPA-founded fabrication service that opened innovation in this space to players who otherwise might have been precluded.

Later, partnerships with universities advanced the Semantic Web — think here about the structure and tagging protocols of web content — and ignited our aspirations for self-driving vehicles. Materials scientists can tell a similar story, going back to DARPA's establishment in the 1960s of university materials research centers — dubbed at the time as Interdisciplinary Laboratories (IDLs). These centers laid the foundations for students who worked on projects ranging from high-temperature ceramics to stealth technologies to metamaterials with unusual optical properties. More recently, DARPA has been among those institutions that


have nurtured university scientists looking at quantum computing and sensing, the limits of photon detection and imaging, and new paradigms in microelectronics to power the post-Moore's-law world.

A generation ago, I was one of these students. My own first experience with DARPA was as a graduate student in computer science at the University of Maryland at College Park. I was spending the summer of 1992 working at the National Institute of Standards and Technology (NIST), and I was attached to a group supporting a new DARPA program — called Manufacturing Automation and Design Engineering, or “MADE” for short — that was just being kicked off. The group I was working for, part of the NIST Manufacturing Engineering Laboratory, had expertise in 3-D modeling and computer-aided design, and we were part of MADE's government team.

At that moment in my early career, the kickoff meeting at DARPA was unlike any event I had ever attended. Consisting of mostly academic researchers, there were about 40 people representing a strange intellectual cross-section of university life. We had theoreticians and geometers, roboticists and materials scientists, artificial intelligence researchers and entrepreneurs. There were major defense contractors among us as well. Presentations ranged from challenges in the design of missile seekers to a strange new manufacturing process that worked by building up layers of metal, like a printer would, but in 3-D. One of the attendees,

Presentations ranged from challenges in the design of missile seekers to a strange new manufacturing process that worked by building up layers of metal, like a printer would, but in 3-D.

A member of the University of Nevada, Las Vegas Team DRC-Hubo tests the team's robot before the DARPA Robotics Challenge (DRC) Finals, a competition of robots and their human supervisors held in 2015.



We are accelerating into a new era of scientific discovery and engineering innovation, and the technologies being unleashed are proving (as usual) to be simultaneously strange, wonderful, and disturbing.

upon hearing the talk about the magical “3-D metal printer,” floated an idea that we could scan physical objects using computer vision and then “beam them,” *Star Trek*-like, over the internet to other places. “That’s a crazy idea!” someone said, dismissively. Yes, crazy indeed. Hmmm, I thought, maybe I’d found my people.

The most exciting of DARPA programs often resemble the Manhattan Project or Apollo program, albeit at a much smaller scale. That’s what I loved about the DARPA I experienced with the MADE program: There was a sense of shared mission, a respect for the technological possibilities, and a commitment to understand how to transform these possibilities into capabilities for the country. What we learn from looking at DARPA’s history with universities is that national security problems do not care what academic department you are from, what conferences you attend, or what the reviewers of your last journal article said. These challenges simply demand that the brightest minds work together from problem to solution. Scientists like Ernest Lawrence, Robert Oppenheimer, Enrico Fermi, and Richard Feynman saw their academic careers as intimately connected to the essential challenges of their era. DARPA has inspired countless engineers and scientists to embrace this same kind of national service.

Fast forward from 1992 to 2014 and I found myself one of those “people” on the technical staff at DARPA, part of the agency’s leadership team, and deputy director of the Defense Sciences Office, where I served until 2017. Now in 2018, we are accelerating into a new era of scientific discovery and engineering innovation, and the technologies being unleashed are proving (as usual) to be simultaneously strange, wonderful, and disturbing. The pace of change is far beyond anything our nation has ever experienced, and the democratization of technologies means that innovation and disruptions can appear from anywhere. While the Manhattan Project required that we establish Los Alamos and a dozen other facilities to carry out the isotope separation, metallurgy, bomb assembly, and other specialized science and engineering functions it would take to build the world’s first nuclear weapons, today globally disruptive innovation emerges from small labs, hackers working with commercial products, anonymous groups on the Dark Web, or even individuals working in their garages. One of the most visible examples of disruptive technology in recent years has been the emergence of the blockchain and associated crypto-currencies whose development trace back to just one scientific paper of uncertain provenance, published by an individual author most likely working under a pseudonym. Indeed, this is both strange and disturbing, perhaps wonderful as well.

DARPA, and the nation, needs the participation of American higher education’s deepest thinkers and most ambitious scientists, engineers, mathematicians, computer scientists, modelers, and theorists, lest we risk losing our technological, economic, and geopolitical preeminence. But many of these strange and disturbing challenges outpace typical government response times, universities’ hiring cycles, and the professional tempo of most faculty and student researchers. Over the past

10 years alone, we have witnessed the emergence of smart phones. CRISPR/Cas9 gene-editing technology, commercial drones, bio-hacking, quantum computation, deep learning, and the rise of social media. The world is rewriting itself on internet timescales and we struggle to keep pace.

Further complicating matters for the current generation is that, in Thomas Friedman’s words, the world is now flat. Universities representing the United States have been vastly successful in promoting our educational model across the globe. As a result, rather than the United States as the final destination for the best and brightest from around the world who receive their educations at our universities, students now find opportunities for research and scholarship in their home countries and elsewhere. While this democratization of innovation and science has been a boon to education and human well-being on the planet, innovation can now come from everywhere and anywhere.

Creating and preventing strategic surprise – an oft-used mantra at DARPA – in this climate requires a new pact between government and academia. DARPA’s call to academia is one that invites university researchers to direct their scholarship, in the spirit of those of previous generations, toward the national security problems of this era. These problems are pushing us toward the unexplored spaces between disciplines as well as to create entirely new disciplines. Because of this, forward-thinking university scholars are realizing that the ways to evaluate success, impact, and societal value are evolving. Lastly, universities are realizing they have new responsibilities to encourage research that is impactful and relevant to national security. However, such work requires universities to find new mechanisms to support restricted research and scholarship that may touch on sensitive matters – ensuring both academic freedom and faculty, students, and national interests are protected. None of this is easy.

This new era of transdisciplinary challenges is unpredictable and unfamiliar compared to traditional academic routes of advancement. As the philosopher of science Thomas Kuhn pointed out in 1962, if we are going to foment scientific revolutions, we are going to need to encourage scientists and engineers to venture into the “crisis” areas in which existing methods are inadequate and anomalies we cannot now explain are bountiful. These explorations should reveal entirely new scientific territories.

As we move forward into DARPA’s seventh decade, the partnership between the agency and universities will continue to open new foundations for our national security and economic well-being. Let us use this sexagenial moment to rethink, refresh, and reinvigorate the historically productive relationship between academia and those in government tasked with tackling some of the most difficult national threats and exciting challenges. This is one of those rare and precious intersections at which national security revolutions emerge. So, to the academic research community: DARPA needs you. The country, for the sake of its future national well-being and security, is counting on you. No pressure. You got this.



• | **60**
YEARS
1958 — 2018