DARPATech, DARPA's 25th Systems and Technology Symposium
August 7, 2007
Anaheim, California
Teleprompter Script for Dr. John Zolper, Director, Microsystems Technology Office

To the Limits and Beyond: The Future of Component Technology

» **JOHN ZOLPER:**

It's good to be here at DARPA Tech to tell everyone about the breakthroughs in component technology that DARPA and MTO
are pioneering,
just as we have
for the past 50 years.

In 1958, the same year DARPA was founded,
the famous physicist Richard Feynman helped launch modern device scaling with a vision:
To build components from the bottom up, one atom at a time.

He offered a $1000 prize for the first person to demonstrate a working motor smaller than the head of a pin.

Since that time,
we've seen fantastic advances in microsystem component technologies:
First the *discrete* transistor.

Then the integrated *circuit*.

And more recently, integrated electronic, photonic, and micromechanical *systems on a chip.*

This progress is the result of a revolution in our ability to make things small by exploiting device scaling.

We've observed Moore's Law by packing more and more transistors onto a microprocessor chip to drive the information revolution;
and we've fabricated millions of tiny mechanical mirrors on a silicon wafer for projection television.

The fact is, we have succeeded in making things small --
but there are many challenges ahead to continuing scaling.

We must find new ways to optimize component performance.

Good afternoon,
I am John Zolper,
the Director of the Microsystems Technology Office.

I'm here to solicit your help in engineering the world of tomorrow to deliver component technology into DoD platforms.

This *could* be a world in which hand-launched UAVs will carry the same computational power as today's room-sized supercomputers.

This *could* be a world in which the bulky collection of sensors and processors in a DARPA Grand Challenge vehicle will fit on a single printed circuit board.

It's a world that *could be*, but only with your help in driving component performance to the fundamental limits.

For a variety of reasons the first fifty years of modern component technology actually exploited **top-down** fabrication to miniaturize components by geometric scaling.

But as we confront the critical challenges of power dissipation,
system complexity, and device density,
the rules of the scaling game are changing!

As a result,
MTO is moving closer to Feynman's original vision of **bottom-up** assembly and atomic-scale fabrication precision.

We are moving into new regimes of device operation where new challenges exist and new opportunities lie.

Later in the week,
the Strategic Technology Office will tell you how they are looking to exploit advances in communications components to enable a truly network-centric force.

The Information Exploitation Office will describe its vision for exploiting advances in sensing components to sense everything, everywhere, all the time.

But these systems offices are working with today's components, developed under past MTO programs.

What will tomorrow bring?

MTO's technical innovations of today
truly are --
what's inside the box -- on tomorrow's DoD platforms.

Let me walk you through an example showing how MTO is setting out to revolutionize a familiar microsystem,

the integrated microprocessor,
by overcoming technical challenges to push processor performance closer to the fundamental limits.

Let’s first look at transistor operation.

Scaling transistors down to feature sizes less than 1/1000th the thickness of a human hair has allowed us to fabricate over a billion transistors on a single chip.

However, this same scaling has produced a significant obstacle in the amount of heat dissipated on chip – enough to heat your chip as hot as your stove!

As we’ll see,
the seemingly simple issue of reducing microsystem power is a multifaceted challenge and is driving technological advances in many different areas.

To start, we know that digital circuits dissipate power every time a gate toggles on and off.

We know that the dissipated power turns into heat.

This active power dissipation per device goes as the node capacitance times the voltage squared times the switching frequency.

Historically, gains in processor performance were achieved by increasing the operating frequency by geometric scaling of the transistor while offsetting power increases through voltage scaling.

Simply put, whenever the frequency was turned up, the voltage was turned down.

Remember when computer vendors kept trying to sell you microprocessors with higher GigaHertz ratings?

You don't hear that anymore.

Manufacturers have hit a limit in their ability to turn the voltage down -- and the frequency, **up**!

No more ads for CPUs with higher clock frequencies!

Now industry wants to sell you multi-core processors that deliver performance increases through architectural tricks that spread out the sources of heat dissipation.

We'll come back to this in a few minutes, but for now it's important to note that we would still benefit from higher switching frequencies if only we could achieve them without overheating the chip!

This begs a question: Have we reached a fundamental limit in the switching frequency of a transistor due to the amount of power that must be dissipated in a switching device?

We don't think so.

The next speaker,
Mike Fritze, will tell you about his vision to beat the active power problem by developing new techniques to turn the voltage down further while still maintaining
high performance.

He believes this can be achieved by developing a new switch technology that exploits quantum mechanical tunneling in place of the classical thermionic transport used in today's transistors.

Mike believes he can exploit quantum mechanics to extend the fundamental performance of digital transistors.

But active computation power is only one source of heat that limits performance gains.

There is also a contribution to heat dissipation from leakage currents that flow even when the transistors are normally off.

Unfortunately this standby power threatens to stop device scaling as it becomes increasingly more significant as transistors become smaller.

But MTO has an idea about how to eliminate this standby power.

Amit Lal wants to solve this problem by taking us "back to the future" by building all-mechanical computers that turn off perfectly with zero standby power dissipation!

His vision is to produce a mechanical computer using highly-integrated, nano-mechanical switches running at multi-GigaHertz speeds.

Now that I have told you how we are looking at reducing heat dissipation at the transistor level,
we must consider how transistors communicate with one another, especially across large chips.

As we push for improved computational power efficiency, we encounter limits in the power efficiency of data communication across microchips.

This brings us back to the new multicore processor architectures I introduced earlier.

Today's manufacturers advertise dual- and quad-core processors, designed in-part to spread out the sources of heat dissipation.

The industry trend is toward the integration of hundreds of cores on a single chip.

But what limits the actual performance that can be realized with such an architecture?

The answer is the connections between the cores, the on-chip communications network.

We want data rates to increase with the number of logic elements on a chip.

The problem is that electrical communications burn a lot of power, and so if we stick with metal wires we are forced to use relatively low-bandwidth interconnects.

The impact can be dramatic.

Imagine having a supercomputer with many computational nodes
all connected via
dial-up modems.

The computational resources exist intrinsically, but actual system performance is limited by the interconnect bandwidth.

The processors cannot run fully loaded with data, because there is an information bottleneck in the on-chip communications network.

As a result, overall computational performance suffers.

Are we stuck?

Will microprocessor performance stop advancing because the multi-cores can not communicate fast enough?

Jag Shah doesn’t think so.

Jag will tell you how he is going to eliminate this communications bottleneck, just as we have done around the world by building optical networks, Jag wants to build a power-efficient *optical* network on-chip!

The technical challenges lie in developing photonic components -- modulators, detectors, and waveguides -- with a small enough footprint, and low enough power,
to be integrated with
high-performance electronics at the necessary densities.

Jag will discuss some of the fundamental limits imposed on these devices and show you that there is plenty of room for progress.

DARPA brought you the technology for optical telecommunications and the internet.

We are setting out to do this again at the dimension of a single integrated-circuit chip.

Let me turn now to how we will design and optimize ever more complex components in the future.

In order to achieve the ultimate scaling limits of all integrated microsystems, we must broadly address the system-level impacts of simultaneously increasing device densities and decreasing device sizes.

Microsystems are becoming so complex that emergent phenomena are likely to appear, such that a microsystem cannot be defined simply as the sum of its parts.

In addition, as individual components become sufficiently small and isolated from their environments, their behavior becomes strongly influenced by the effects of quantum coherence and quantum confinement, we move squarely into the realm of quantum mechanics.

To operate at these extremes, we need to build design tools which appropriately account for these effects.

Dennis Healy will describe how he intends to deliver new tools to enable breakthroughs in component performance even in the quantum regime.

He believes that by exploiting advances in mathematics concurrently with device improvements, he can accelerate new hardware and software co-optimization strategies.

Dennis also believes his new tool set will allow us to design new devices that exploit quantum effects in new ways.

So after MTO revolutionizes microprocessor technology and delivers new quantum devices, what other fundamental limits do we seek to approach?

Well often an explicit defense need serves as the best motivator for the development of a new component technology which dramatically changes the technical landscape.

In the 80’s and 90’s DARPA developed Microwave Monolithic Integrated Circuits -- MMICs -- for DoD radar systems such as the electrically scanned array radars that are found today on F22 fighters and Navy Destroyers.

This same development subsequently enabled today's cellular wireless technology, and the breadth of commercial applications for integrated microwave circuits has been extraordinary.

But there are ways to improve microwave technology.

MTO expects a dramatic improvement in wireless communication and stand-off rf sensing
as we drive up the operating frequencies and functionality of
smart rf arrays.

Mark Rosker will describe the fundamental challenges that must be overcome to develop rf arrays that can fit on any platform,
that can deliver multiple beams with high power, that can operate with high efficiency, and
that can be low cost!

Mark will address his vision for pushing the operating frequency of analog electronics into the THz range, and packing microwave technology into highly integrated three-dimensional arrays.

But there is more excitement coming from MTO!

MTO is investing in many other component focus areas to support our mission partners and provide the technologies needed inside the box.

Here are some examples.

Focal Plane Arrays:
Let's imagine a single scope which can image in low-light or daylight,

detect body heat, or
see through smoke and clouds.

The ability to image and detect photons across a broad range of the electromagnetic spectrum continues to be a critical enabler for the warfighter.

MTO, with your help,
is pushing the limits of photon detection from short-wavelength ultraviolet all the way to millimeter wave,
with high dynamic range and diffraction-limited resolution.

Lasers: MTO believes that compact, ultra-efficient laser designators, directed energy weapons, and highly integrated optical communications systems **all** are within reach because of MTO's efforts on photon generation.

We are pushing towards the fundamental efficiency limits for
diode-pump lasers,
mid-wave IR lasers, and ultraviolet lasers.

And we are seeking to deliver efficiency with superior beam quality and robust operation!

Bring us ideas on how to deliver these results.

Energy: When it comes to the need for fully autonomous, high-endurance sensors, or the simple notion of removing the battery burden from the soldier deployed in the field, MTO plays a critical role.

MTO is pushing the frontiers of Microsystems for energy efficiency, generation, conversion, and storage.

We're pursuing energy harvesting and the exploitation of radioisotopes in a deployable package.

We are striving to reach the fundamental limits of efficiency at the junction level and to match energy source capabilities to microsystem component demands,
but we need your ideas to reach these goals.

But even with such component strategies in place, and even when the resultant microprocessors and microsystems are delivering superior performance,
our job is not done.

We are not content simply to develop technologies!

We want to ensure that they remain available and trustworthy for the DoD.

Microsystems technologies play an important role in the most sensitive defense applications.

That's why MTO is looking to approaches for the production and evaluation of what we call,
Trusted Components.

Dean Collins will tell you about emerging techniques to measure whether components contain exactly what the system designer specified – Nothing More and Nothing Less.

All these -- and more -- breakthrough efforts are being pursued by an entire team of DARPA Program managers working within MTO.

As you hear from the current program managers, ask yourself, is this for me?

Should I join MTO to push the limits of technology and deliver new capabilities to the warfighter?

Please stop by the MTO booth and talk to MTO program managers to learn more about the technological issues we are addressing.

This is a great challenge; if you want to be part of this adventure, I would like to hear from you.

Thank you for your attention, and remember as MTO continues to succeed in making things small our technology
“is what’s Inside the Box!” of DoD platforms.

Let’s get the MTO ball rolling with
Dr. Michael Fritze,
who will describe how he is pushing the limits of digital electronics.