



2025 IEEE 75th Electronic Components and Technology Conference

**Media Contacts:**

Gary Dagastine  
Dagastine & Co. PR  
(518) 785-2724  
[gdagastine@nycap.rr.com](mailto:gdagastine@nycap.rr.com)

Chris Burke  
BtB Marketing Communications  
[chris.burke@btbmarketing.com](mailto:chris.burke@btbmarketing.com)

**For Immediate Release**

*This Tip Sheet highlights selected papers. . .*

***2025 IEEE Electronic Components and Technology Conference (ECTC)  
To Showcase the Latest Electronic Packaging Technologies, including:***

- *Hybrid bonding*
- *Heterogeneous integration*
- *Co-packaged optics*
- *Reliability considerations for advanced packages*
- *Image sensors with built-in artificial intelligence functionality*
- *Flexible systems*

DALLAS, TX (April 24, 2025) – The 75<sup>th</sup> annual [IEEE Electronic Components and Technology Conference \(ECTC\)](#), the world’s leading forum for unveiling, discussing and exhibiting the latest advances in microelectronics packaging and component science and technology, will take place May 27-30, 2025 at the Gaylord Texan Resort & Convention Center here. More than 2,000 scientists, engineers and businesspeople from 20+ countries are expected to attend, along with more than 135 exhibitors.

“Each year the ECTC conference looks forward and explores how we can make new, better, faster, smaller, more reliable and more efficient electronic technologies for our increasingly complex world,” said Przemyslaw Gromala, ECTC 2025 Program Chair and Sr. Expert/Simulation Team Leader at Robert Bosch GmbH. “That makes the ECTC conference a leading technology indicator because the breakthroughs needed to build future electronic products are unveiled in our technical program. This year is no exception, as detailed below.”

“But also this year, our 75<sup>th</sup> annual meeting, we look back to acknowledge the skills, accomplishments, and contributions of the thousands of people who have participated over the years, whose work has helped to make electronics one of the most important, vibrant and exciting industries in the world,” he said.

This Tip Sheet describes noteworthy talks from the ECTC 2025 technical program:

**A) Hybrid Bonding**

**A Better Hybrid Bonding Process:** Rapid advancements in AI and high-performance computing (HPC) have made high bandwidth memory (HBM) and more efficient logic/memory device interconnections

critical. Traditional packaging approaches use solder bump technology to form interconnections between dies and memory components. But as the input/output (I/O) count increases and fine-pitch requirements become more urgent, the reliability and yield challenges associated with solder bump technology have increased. That's why hybrid bonding, which enables the simultaneous bonding of dielectric and metal pads, has emerged as a promising solution. It can accommodate fine-pitch interconnections with superior reliability. However, issues with commonly used oxide-based hybrid bonding approaches pose a significant barrier to its adoption in mass production. The main issues are the high cost of the multiple steps of the CMP process used, and the low tolerance of surface topography.

At ECTC, a National Yang Ming Chiao Tung Univ.-led team will describe a different hybrid bonding approach – Cu/polymer hybrid bonding. They say it offers both lower cost and a broader process window, and achieves results similar to oxide-based approaches. They demonstrated a 30-second, 150-200°C Cu/polymer hybrid bonding process with a) a low-cost, wide process window (surface roughness 2~20 nm) and b) high-throughput (< 30-second bonding duration) for high I/O applications. (**Paper 32.3, “Ultra-Fast Cu/Polymer Hybrid Bonding with Electroless Passivation Layer for Cost-Effective High I/O Interconnection Stacking,”** Yu-Lun Liu et al, National Yang Ming Chiao Tung Univ./ Tokyo Ohka Kogyo Co. LTD)

**Fluidic Self-Alignment for Hybrid Bonding:** Die-to-wafer hybrid bonding (D2W HB) has many advantages over traditional solder bonding, but it is costly because attaching chiplets to a wafer at extremely tight tolerances is slow and complicated. That will only get worse at smaller HB pitches and/or as the number of chiplets per stack increases. Intel researchers will describe the latest developments in the use of fluidic self-alignment to increase die-to-wafer hybrid bonding throughput. It makes use of capillary action for fast, precise alignment of a chiplet to a bottom wafer. First, top dies and bottom wafer are lithographically patterned to create liquid-confinement features around the bonding areas. Then, liquid is dispensed to form a droplet in the center of each bonding area on the bottom wafer. The top dies are transferred onto the droplets, which then spread out to the confines of the patterned area on the bottom surface. Capillary forces from the liquid's high surface tension align the dies to the lithographic patterns on the bottom wafer. The liquid is subsequently evaporated.

The researchers say this fluidic approach does not negatively impact the electrical resistance of the interface compared to standard hybrid bonding, and that it is expected to increase throughput by more than 10X versus current hybrid bonders, significantly easing cost pressures. Intel initially unveiled the process at ECTC last year in a proof-of-concept study, and this year they will describe developments since then, including 1) a new process flow with plasma dicing; 2) changes to wafer processing to ensure self-aligned feature compatibility with the HB Cu pads; 3) assembly process optimization for yield and throughput improvements, and 4) fabrication and initial testing of an electrical test vehicle with HB self-alignment features. They will also describe next steps, including the transfer of some of the assembly flow from the laboratory to an advanced packaging fab, and the development of a first-of-a-kind bonder to leverage the advantages of fluidic self-alignment. (**Paper 22.1, “Hybrid Bonding with Fluidic Self Alignment: Process Optimization and Electrical Test Vehicle Fabrication,”** F. Eid et al, Intel)

**Optimized Self-Formed Barrier Layer for Scaled Hybrid Bonds:** At ECTC, IMEC researchers will describe the use of a self-formed MnOx barrier system in a 400nm pitch wafer-to-wafer hybrid bonding technology, in which a 60nm CuMn PVD seed is deposited directly after the hybrid pad etch. This technique makes it possible to omit a conventional Ta barrier deposition step, opening up the possibility to further reduce hybrid pitch pad scaling. By carefully tuning certain critical process steps, among them the hybrid pad copper polish, high electrical yield on large hybrid bonded daisy chain structures was achieved. The researchers will discuss the results of their studies of dielectric breakdown, copper diffusion, Cu bulge-out and electromigration mechanisms of the self-formed barrier. (**Paper 9.4, “Self-formed Barrier using Cu-Mn Alloy Seed applied to a 400nm Pitch Wafer-to-Wafer Hybrid Bonding Technology,”** S. Van Huylenbroeck et al, IMEC)

**Progress in Flexible Hybrid Electronics:** Flexible hybrid electronics (FHE) is an emerging technology that integrates rigid semiconductor components like sensors, microcontrollers and communication modules onto a flexible substrate. It is a significant advancement over traditional flexible electronics because it enables systems that are not only lightweight and flexible, but also capable of delivering high performance. However, ensuring the mechanical reliability of flexible systems under constant movement with repeated bending, twisting and deformation is critical in applications such as soft robotics, flexible displays and wearable systems. Key to that ability are flexible/bendable redistribution layers (RDLs), which facilitate the distribution of electrical signals and power across the device while maintaining flexibility. They also contribute to the miniaturization of electronic devices because RDL flexibility allows for innovative designs that can conform to the contours of the human body or other surfaces. However, single-layer RDL designs lack the capability to connect multi-terminal functional components, while two-dimensional serpentine RDLs require horizontal space to accommodate bending-induced deformation.

At ECTC, Tohoku Univ. researchers will describe a novel methodology for interconnecting FHE functional components using a double-layer RDL, based on fan-out wafer-level packaging (FOWLP) techniques. They conducted a feasibility study of their approach, in which after 100 25mm-radius bending cycles, the resistance of the first- and second-layer RDL increased by only 9.9% and 16.1%, respectively, indicating relatively high mechanical reliability. The reliability can be further enhanced by a controlled passivation layer on the top of the second RDL. (Paper 29.5, “*Design and Fabrication of Bendable Double-layer RDL Metallization Based on FOWLP for In-Mold Flexible Hybrid Electronics (iFHE)*,” C. Liu et al, Tohoku Univ.)

## B) Heterogeneous Integration

**300mm RF Interposer with Integrated Digital and mmWave Interconnects:** SiGe and III-V technologies have better power efficiency and gain at mmWave frequencies than scaled CMOS nodes, while CMOS offers the best power efficiency for digital applications. Therefore, for high-performance RF communications and sensing applications, heterogeneous integration of chiplets brings many advantages. A key requirement, though, is the ability to integrate mmWave and digital interconnects on a single carrier for optimum performance and cost-effectiveness.

At ECTC, IMEC researchers will describe how they did just that. They built a state-of-the-art 300mm RF Si interposer platform with integrated interconnects, optimized for heterogeneous integration of chiplets in mmWave and sub-THz applications. Compared to earlier work, it demonstrated reduced losses up to sub-THz frequencies (0.23 dB/mm at 140 GHz, 0.5 dB/mm at 220 GHz, and 0.73 dB/mm at 325 GHz for RDL2), which is state-of-the-art performance. The digital links use Cu damascene BEOL layers, while the mmWave links employ transmission lines on a low-loss polymer, enabling integration of RF-to-sub-THz CMOS and III/V chiplets. Also, high-quality passive components were incorporated on the interposer to reduce active chip area, cutting costs and providing compact, low-loss RF interconnects. Built on a 300mm wafer, the researchers say this innovation is scalable for cost-effective, high-volume manufacturing. (Paper 5.6, “*RF Si Interposer Platform for Chiplet-based Heterogeneous Systems*,” X, Sun et al, IMEC)

**A Novel Architecture for Signal & Power Optimization for AI/HPC:** The demand for extremely high data bandwidth (~terabyte-per-second) and extensive input/output connectivity has grown significantly for artificial intelligence (AI) and high-performance computing (HPC) applications. To meet it, advanced packaging techniques have become essential and silicon interposers, known for their high-density interconnects, have been widely adopted. But they bring challenges related to the complexity of their redistribution layers (RDLs) and the inherent limitations of through-silicon vias (TSVs). Adding additional RDL layers can improve power delivery network performance, but the increased thickness can induce wafer warpage due to stress accumulation. Similarly, although high-aspect-ratio TSVs improve

integration density, they introduce issues such as potential void-free filling, stress management problems and increased fabrication complexity.

To overcome these challenges, at ECTC a HANA Micron-led team will propose a heterogeneous integrated chip (HIC™) for 2.xD packaging for wide I/O applications. It uses bridge dies and copper posts instead of a conventional silicon interposer. This design integrates signal and power routing within the bridge die structure and copper posts, thereby eliminating the need for a dedicated frontside RDL and C4 bumps. As a result, routing path lengths are reduced, parasitic RC delays are lowered, and power delivery network formation is simplified, improving both signal and power integrity. The proposed design supports up to 6.4 Gbps per lane in a staggered configuration, suitable for HBM3 high-speed connectivity, and the researchers will highlight its potential as a scalable, cost-effective solution for next-generation high-speed packaging applications. (IP session #38, “*Signal and Power Integrity Optimization Using Novel Bridge Die and Copper Post Interconnect Design in 2.xD Packaging for Wide I/O Applications*,” Y. Na et al, HANA Micron Inc./ Myongji University/ Swevenz Inc.)

### C) Photonics & Co-Packaged Optics

**Reducing the Cost & Complexity of Co-Packaged Optics:** Large-scale data transfers are becoming a critical bottleneck in high-performance computing (HPC) and data center architectures. While total throughput of the electrical interconnects has increased in recent years, to further increase it – while also addressing bandwidth density/energy efficiency challenges – co-packaged optics leveraging silicon photonics (SiPh) has emerged as a promising solution. In a typical co-packaged optics system, optical-electrical chiplets integrate the CMOS chip and SiPh photonic integrated circuit, while external laser sources supply the necessary light for signal transmission through a fiber connector attached to the chiplets. The current standard approach relies on polarization maintaining fiber (PMF) arrays between the external laser sources and the chiplets to ensure stable polarization states, required for optimal performance of the SiPh-based photonic devices. But while PMFs provide robust polarization control, they introduce significant drawbacks including high-precision alignment requirements, increased manufacturing complexity, and substantial packaging costs.

At ECTC, NVIDIA researchers will describe an alternative approach that replaces PMFs with single-mode fibers (SMFs) to reduce cost and complexity without compromising system performance. They will present experimental results from a test chip demonstrating the viability of this solution, including a detailed cost-performance analysis comparing PMF- and SMF-based architectures. (Paper 2.6, “*All-SMF Arrays for Co-Packaged Optics: Optimizing Cost, Complexity and Performance*,” N. Mehta et al, NVIDIA)

**Co-Packaged Optics Enable Novel Multi-Chip Package:** New, fast-growing applications like AI, machine learning, self-driving cars and others are driving demands for complex computing systems that deliver high performance at lower bandwidths, with uncompromised reliability and high energy efficiency at lower costs. To meet these requirements, co-packaged optics (CPO) have gained prominence in recent years for their ability to enable heterogenous integration, so that different dies and components can be packaged together to achieve the desired functionality. Intel researchers will describe development of a proof-of-concept multi-chip package (MCP) with a novel open cavity bridge substrate. It consists of a compute tile connected through Intel’s Embedded Silicon Bridge (EMIB) technology to two data converter dies on one side, and three electronic IC-photonic IC (EIC-PIC) stacks on the other side with three corresponding fiber array units (56 couplers, 127um pitch) via fiber-based butt-coupling technology. (Butt-coupling is a packaging technique used to connect optical components directly to each other by aligning their optical interfaces so light passes seamlessly between them at high optical coupling efficiency and with minimal loss.) The device demonstrated good bump alignment, no voiding or delamination and minimum warpage for successful fiber attach. The researchers say future CPO research

will need to focus on system-level reliability and on architectures optimized for yield and throughput. (Paper 9.1, “*Demonstration of Co-Packaged Optics Assembly for Fiber-Based Optical Interconnect*,” M. Baker et al, Intel)

## D) Reliability

**Machine Learning for Better Thermal Management:** Thermal management is one of the most critical issues impacting the performance and reliability of next-generation electronics packaging technologies. That’s because new architectural features such as backside power delivery networks (BSPDNs) and 3D IC configurations introduce new components within the thermal pathway, creating additional heat dissipation challenges. To optimize thermal management, better estimation of chip and package temperatures is required. But while analytical approaches to estimate BEOL interconnect-layer thermal resistance have been developed, they are inadequate.

At ECTC, IBM researchers will describe how they used a series of finite element modeling (FEM) simulations to train a machine learning (ML) model to rapidly predict the thermal resistance of BEOL stacks in a test chip. BEOL layout design, heights and material information were used as inputs. The researchers say the ML model predicted the thermal resistance of BEOL stacks with a mean absolute percentage error (MAPE) of less than 15%. That is a remarkable accuracy improvement versus the traditional 1-D heat conduction analytical model used for comparison, which showed a MAPE of 300%. The machine learning approach opens the possibility of more accurate predictions of hotspots in advanced packaging architectures, thus helping to increase their performance and reliability. (IP session #37, “*Fast And Accurate Machine Learning Prediction of Back-End-Of-Line Thermal Resistances in Backside Power Delivery and Chiplet Architectures*,” P. Chowdhury et al, IBM)

**A Tool for Optimized Cooling of High-Performance Systems:** Cooling the high-performance computing systems in AI datacenters, 5G radio access networks (RAN), and edge compute nodes is a growing challenge. That’s because the multi-chip modules (MCMs) used in these systems, which integrate CPUs, GPUs, and high-bandwidth memory (HBM) in 2.5D/3D architectures, already draw more than 1000 watts, and that will increase going forward.

A Nokia Bell Labs team will describe a high-power thermal test vehicle (TTV) they built to assess the efficiency of various novel cooling solutions. At its core is a multi-chip module (MCM) consisting of six emulated logic dies and 12 emulated HBM stacks. These are mounted on a ceramic interposer capable of dissipating more than 2600 W, with localized hotspots reaching 314 W/cm<sup>2</sup> (and 817 W/cm<sup>2</sup> over the heater resistor area). The MCM has 264 platinum-based resistive temperature detectors for high-resolution thermal monitoring. It enables independent power control across 60 heater regions, facilitating the creation of diverse heat maps of discretized regions across the MCM surface. This capability enables the study and optimization of various novel cooling technologies. In addition, the TTV was integrated into a pumped two-phase refrigerant cooling system, which successfully handled heat fluxes up to 314 W/cm<sup>2</sup> over a 46 mm<sup>2</sup> hotspot on each logic die. By reducing peak chip temperatures as much as 25°C, it highlights the superior thermal performance of two-phase versus single-phase cooling. (IP session #40, “*Experimental Demonstration of High-Power Thermal Test Vehicle using Two-Phase Cooling for AI Datacenters, 5G RAN, and Edge Compute Nodes*,” Yang Liu et al, Nokia Bell Labs)

**Better Predictions of Thin-Films’ Plasticity and Fracture Behavior:** Brittle thin films such as silicon dioxide, silicon nitride and others are widely used in semiconductor devices, where their mechanical reliability is critical. It’s difficult to accurately characterize the plasticity and fracture properties of thin films at sub-micron scales, though, because of the limitations of conventional mechanical testing methods. Yet understanding these properties is essential because they directly influence film deformation and crack behavior during the manufacturing process. These, in turn, affect the performance and reliability of the microelectronics and memory devices made from these films.

At ECTC, a University of Singapore/Micron Technology team will describe a novel methodology for predicting thin-film plasticity and fracture properties by integrating finite-element analyses (FEA) simulations with experimental nanoindentation data. (Nanoindentation is where a controlled force is applied to a material's surface using a sharp indenter, and the resulting deformation is then analyzed.) The experimental results show a strong correlation between the simulations and the observed fracture behavior, validating the methodology. The researchers say this framework offers a more cohesive, easier-to-implement alternative to traditional modeling methods, and it can be extended to also predict the behavior of ductile thin films; the adhesion strength between layers; and the porosity of low-k dielectric materials, broadening its applicability. (**Paper 24.5, “Predictive Modelling of Thin Films Properties Using An Experimental and Simulation-Based Approach,”** D.K. Lim et al, Nat'l Univ. of Singapore/Micron Technology)

**Microbumps and Electromigration:** As the electronics industry pushes the boundaries of miniaturization to the angstrom level, reliability is a growing concern. One area of focus is the solder microbump, used to interconnect chips within multichip modules, and to connect chip packages to printed circuit boards (PCBs). As microbumps shrink and the current density flowing through them increases, a phenomenon called electromigration (EM) is more likely to damage them. EM is where the flow of electric current actually causes the conductor material to move, creating voids or short circuits within it. A better understanding is needed of how EM occurs in materials commonly used to make ultrasmall microbumps, and how it can be prevented.

At ECTC, TSMC researchers will describe studies they conducted to do that. They investigated the EM performance of microbumps made from Cu/Sn and Cu/Ni/Sn/Cu in a full, complete intermetallic compound (IMC) structure. IMC is the interface layer where the solder and the base metal meet. It is formed during the soldering process. The researchers will report that full IMC microbumps are robust enough to withstand EM. Further, when EM does occur its main failure mode is failure of the aluminum (Al) traces, caused by the electromigration of Al atoms. This finding underscores the importance of forming full IMC joints to enhance the reliability of microbumps in demanding applications such as high-performance computing. (**Paper 28.5, “The Influence of Full IMC Structure on Micro-Bump Electromigration Performance,”** Chung-Yu Chiu et al, TSMC)

**Shock Testing Flip-Chip Interconnects:** Increasingly sophisticated electronic systems are being more widely deployed in harsh, real-world applications such as automotive, avionics, industrial production and others. That means thermal/mechanical issues like heat, vibration and shock are more critical because of their potential reliability impacts. But existing shock/vibration test methods and equipment are falling short because they were designed for less complex electronic systems. At ECTC, TU Dresden researchers will describe better ways to simulate and perform shock testing under real-world conditions. They will unveil a new test vehicle for shock and other mechanical loads, which can test multiple components simultaneously. In the initial design, their test board consisted of nine flip-chips (2.9 x 4.9mm<sup>2</sup>) mounted in a 3x3 matrix configuration. They were tested until failure under shock loads of 130g amplitude/2ms pulse width, and at -40°C, room temperature, and 125°C. Failure analysis after the tests showed crack initiation, propagation and even some complete cracks near the solder joints on both the flip-chip and board sides of the joints. The researchers say these results will serve as a basis for better shock profiles, measurement practices and fixture methods for future testing, and for testing with a mixture of loads as well. (**Paper 18.7, “Isothermal Shock Testing on Flip-Chip Interconnects,”** M. Häusler et al, TU Dresden)

## E) Assembly and Manufacturing Technology

**Warpage Control for Ultra-Large Packages:** Chip-on-wafer-on-substrate (CoWoS) technology is a 2.5D/3D platform used to integrate SoC (system-on-a-chip) and HBM (high-bandwidth memory) dies on an interposer. As AI and high-performance computing requirements push the limits of bandwidth and

performance, the sizes of both the integrated components and the interposers have been increasing. That means warpage control is becoming an urgent challenge, because it impacts joint quality, assembly yields, functional testing and board-level assembly. TSMC's CoWoS-R platform uses an organic polymer-based interposer that incorporates a redistribution layer (RDL). At ECTC, TSMC researchers will report on their studies of warpage in ultra-large (110x110mm<sup>2</sup>) CoWoS-R interposers which integrate 4 SoCs and 12 HBMs. They will present the results of their analyses, which took into account substrate size, CoW effects, stress buffering films, heat sink placement and others, and will identify remaining critical challenges going forward. (Paper 4.1, "*Package Warpage Reduction for Large CoWoS-R Packages*," Yu-Hsiang Hu et al, TSMC)

**Face-Down Chip-on-Wafer Bonding:** System-in-package (SiP) solutions for mobile and wearable devices must be thin and have a small overall footprint, but conventional SiP devices using organic substrates and solder interconnects are reaching their size limits. Also, managing wafer warpage during processing is critical due to mismatched coefficients of thermal expansion among the various system components. At ECTC, a Murata-led team in WOW Alliance of Science Tokyo will propose a Face-Down Chip-on-Wafer (FD COW) process to address these challenges. It makes use of a waffle wafer stack structure that enables bumpless via-last interconnects. Waffle wafers have a patterned grid-like structure on their surface which enables them to be very thin; their reduced volume mitigates warpage. For FD COW bonding to become practical, the adhesive coating used must be enhanced and uniform distribution of that adhesive must be ensured on the trench surfaces of the waffle wafer. It is also essential to optimize chip bondability and to minimize misalignment under high-speed bonding conditions. The researchers will describe an innovative selective adhesive coating method, and evaluate its compatibility with high-speed chip bonding for bumpless FD COW using a single waffle wafer. Through the development of this process, >30,000 chips of several different sizes were bonded at 40 µm spacing, achieving >20X faster bonding speed without any chip-detachment failures. The researchers say the face-down COW process is highly promising for high-speed, high-density heterogeneous bonding applications. (Paper 1.2, "*Face-Down Bonding and Heterogeneous Chiplet Integration by Using Bumpless Chip-on-Wafer (COW) with Waffle Wafer Technology*," Y. Satake et al, Murata/Institute of Science Tokyo/Panasonic Connect Co. Ltd)

**Laser-Assisted Bonding with Compression:** To make high-bandwidth memory (HBM), memory dies with through-silicon vias are vertically stacked in 2.5D or 3D packages, and physically/electrically interconnected by solder joints. But with increasing I/O density these joints are getting much smaller, and conventional soldering technology can no longer guarantee reliable interconnections because of warpage, fume generation, residual flux, and incomplete underfill coverage. Laser-assisted bonding (LAB) has attracted attention for fine-pitch interconnections because it minimizes thermally induced deformation of chips and substrates, and laser-assisted bonding with compression (LABC) improves alignment precision by applying constant pressure on the chips during bonding.

At ECTC last year, ETRI reported a chip-on wafer technology using LABC with a laser non-conductive film (NCF). This year, an ETRI-led team will report on a localized formation of laser NCF on the surface of 10 µm diameter solder bumps and its application on 20 µm pitch interconnection. They carefully dipped the micro-bumps in the laser NCF to selectively coat only the surface of the SnAg solders. Then, the die with the laser NCF-coated micro-bumps was bonded to the substrate using LABC. Fine-pitch interconnections were achieved with no residue and fume generation. (Paper #31.5, "*Localized Formation of Laser Non-Conductive Film (NCF) on 10 µm Diameter Bumps, Applied to 20 µm Pitch Chiplet Chip-on-Wafer (CoW) Bonding*," J. Shin et al, ETRI/Hanbat National Univ.)

## F) Noteworthy Papers on Diverse Topics

**Adding AI to CMOS Image Sensors:** Edge computing using image sensors is expected to become increasingly important in the future, as edge data is fed into AI chips to process image data for face and object recognition. With that in mind, Sony researchers will describe a novel process that makes it possible to integrate an AI chip containing a built-in deep neural network (DNN), into the bottom wafer of a conventional two-wafer stacked CMOS image sensor. The wafer-on-wafer-on-wafer (WoWoW) process makes use of 6 $\mu$ m face-to-face and face-to-back Cu-Cu connections, and 6 $\mu$ m through-silicon vias (TSVs). The researchers say the imaging characteristics of conventional two-wafer-stacked image sensors are maintained, and the pixel density in the top chip was able to be increased to its theoretical limit. The device demonstrated high dynamic range, crucial for image quality, in both bright and dark environments. (Paper 14.4, “*Development of A Novel WoWoW Process for 1/1.3-inch 50 Megapixel Three-Wafer-Stacked CMOS Image Sensor with DNN Circuits,*” K. Shimizu et al, Sony)

**A Design Approach to Optimize Hybrid Voltage Regulators:** Dense, power-efficient 48V-to-1V voltage regulators, embedded near points-of-load rather than mounted on a printed circuit board, are needed to support the high-power, high-current density needs of next-generation high-performance computing systems. These integrated power delivery solutions, or hybrid voltage regulators (HVRs), must support the dual requirements of increased functionality in a small form factor, and reduced energy and power consumption. However, the ad hoc simulation-based design approaches used thus far with HVRs fundamentally limit the optimization of HVR circuit architectures and sizing.

To address this, University of Illinois researchers will present a generalized HVR architecture, analytical model, and design optimization methodology at ECTC. They compared their analytical power loss calculations to simulation results from commercially available simulation software, and saw more than 92.87% accuracy across all the HVR architectures studied. They studied tradeoffs of power efficiency vs. density based on various proposed models, as well as approaches to enhance the power efficiency of existing HVRs. As a result, they will propose at ECTC a design methodology that considers a) the number of components and sharing of inductors among switching branches, and b) the balancing of various types of HVR power losses through the use of optimal design parameters in a multi-step optimization process that has shown over 8% reduction in total power loss. (IP session #41, “*Hybrid Voltage Regulators for High Performance Computing: Analytical Models and Design Methodology,*” S. Abdelzaher et al, University of Illinois)

###